ISSN 1472-2739 (on-line) 1472-2747 (printed) **1167**

Algebraic & Geometric Topology
Volume 3 (2003) 1167–1224
Published: 12 December 2003

On a theorem of Kontsevich

James Conant
Karen Vogtmann

Abstract In [12] and [13], M. Kontsevich introduced graph homology as a tool to compute the homology of three infinite dimensional Lie algebras, associated to the three operads "commutative," "associative" and "Lie." We generalize his theorem to all cyclic operads, in the process giving a more careful treatment of the construction than in Kontsevich's original papers. We also give a more explicit treatment of the isomorphisms of graph homologies with the homology of moduli space and $Out(F_r)$ outlined by Kontsevich. In [4] we defined a Lie bracket and cobracket on the commutative graph complex, which was extended in [3] to the case of all cyclic operads. These operations form a Lie bi-algebra on a natural subcomplex. We show that in the associative and Lie cases the subcomplex on which the bi-algebra structure exists carries all of the homology, and we explain why the subcomplex in the commutative case does not.

AMS Classification 18D50; 57M27, 32D15, 17B65

Keywords Cyclic operads, graph complexes, moduli space, outer space

1 Introduction

In the papers [12] and [13] M. Kontsevich sketched an elegant theory which relates the homology of certain infinite-dimensional Lie algebras to various invariants in low-dimensional topology and group theory. The infinite-dimensional Lie algebras arise as generalizations of the Lie algebra of polynomial functions on $\mathbb{R}^{2n}$ under the classical Poisson bracket or, equivalently, the Lie algebra of polynomial vector fields on $\mathbb{R}^{2n}$ under the Lie bracket. One thinks of $\mathbb{R}^{2n}$ as a symplectic manifold, and notes that these Lie algebras each contain a copy of the symplectic Lie algebra $\mathfrak{sp}(2n)$. The relation with topology and group theory is established by interpreting the $\mathfrak{sp}(2n)$-invariants in the exterior algebra of the Lie algebra in terms of graphs, and then exploiting both new and established connections between graphs and areas of low-dimensional topology and group theory. These connections include the construction of 3-manifold and knot invariants via data associated to trivalent graphs, the study of automorphism groups of free groups using the space of marked metric graphs (Outer space), and the use of ribbon graph spaces to study mapping class groups of punctured surfaces.

This paper is the outcome of a seminar held at Cornell, organized by the second author, devoted to understanding Kontsevich's theory. Kontsevich describes three variations of his theory, in the commutative, associative and Lie "worlds." Kontsevich's papers skip many definitions and details, and, as we discovered, have a gap in the proof of the main theorem relating symplectic invariants and graph homology. In this paper we explain Kontsevich's theorem carefully, in the more general setting of cyclic operads. In particular, we adapt a fix that Kontsevich communicated to us for the commutative case to the general case. We then specialize to the Lie, associative and commutative operads, which are the three worlds which Kontsevich considered in his original papers. Using a filtration of Outer space indicated by Kontsevich, we show that the primitive part of the homology of the Lie graph complex is the direct sum of the cohomologies of $Out(F_r)$, and the primitive part of the homology of the associative graph complex is the direct sum of the cohomologies of moduli spaces (or equivalently mapping class groups) of punctured surfaces. We then recall the Lie bracket and cobracket which we defined on the commutative graph complex in [4], and which was extended in [3] to the case of all cyclic operads. These operations form a bi-algebra structure on a natural subcomplex. We show that in the associative and Lie cases the subcomplex on which the bi-algebra structure exists carries all of the homology, and we explain why the subcomplex in the commutative case does not.

Many people contributed to this project. Participants in the Cornell seminar included David Brown, Dan Ciobutaru, Ferenc Gerlits, Matt Horak, Swapneel Mahajan and Fernando Schwarz. We are particularly indebted to Ferenc Gerlits and Swapneel Mahajan, who continued to work with us on understanding these papers after the scheduled seminar was over. Mahajan has written a separate exposition of some of the material here, with more information about the relation with classical symplectic geometry, using what he calls *reversible* operads, which are closely related to cyclic operads. His treatment of the theorem relating graphs and Lie algebra invariants uses Kontsevich's original fix, and therefore does not work for every cyclic operad. In particular, it cannot handle the Lie case. A succinct outline of Kontsevich's theory can be found in the thesis of Gerlits, which also includes a careful study of the Euler characteristic using Feynman integrals. We hope that these different expositions, with different emphases, will help to make Kontsevich's beautiful theory more broadly accessible.

The present paper is organized as follows. In Section 2 we develop the theory for general cyclic operads. After reviewing the definition of cyclic operad we define a chain complex parameterized by graphs whose vertices are "colored" by operad elements. This chain complex was introduced in a more general setting by Getzler and Kapranov in [8], and was studied by Markl in [16]. We then construct a functor from cyclic operads to symplectic Lie algebras, as the direct limit of functors indexed by natural numbers. We then show how to use invariant theory of the symplectic Lie algebra to define a map from the Chevalley-Eilenberg complex of the Lie algebra to the above chain complex of graphs which is an isomorphism on homology.

In section 3, we specialize to the Lie operad. By using a filtration of Outer space indicated by Kontsevich, we prove that the primitive (connected) part of the graph complex computes the cohomology of the groups $Out(F_r)$ of outer automorphisms of finitely-generated free groups. We also prove that inclusion of the subcomplex spanned by 1-particle irreducible graphs (i.e. graphs with no separating edges) is an isomorphism on homology. This is of interest because, as was shown in [4, 3] this subcomplex carries a graded Lie bi-algebra structure; this implies, among other things, that the homology of the groups $Out(F_r)$ is the primitive part of a differential graded Hopf algebra. In section 4 we note how the theory applies in the case of the associative operad. In this case, the primitive part of the graph complex is shown to compute the cohomology of mapping class groups of punctured surfaces. The proof proceeds by restricting the filtration of Outer space to the "ribbon graph" subcomplexes, on which mapping class groups of punctured surfaces act. As in the Lie case, we show that the subcomplex of 1-particle irreducible graphs carries all of the homology, so that the direct sum of homologies of mapping class groups is the primitive part of a differential graded Hopf algebra. Finally, in section 5 we reconsider the commutative case, which was the focus of our paper [4]. We give a geometric description of the primitive part of commutative graph homology, as the relative homology of a completion of Outer space modulo the subcomplex at infinity, with certain twisted coefficients. This relative homology measures the difference between the relative quotient of Outer space by $Out(F_r)$ and the quotient by the subgroup $SOut(F_r)$ of outer automorphisms which map to $SL(r,\mathbb{Z})$ under the natural map $Out(F_r) \to GL(r,\mathbb{Z})$. Using this geometric description of graph homology, we explain why the one-particle irreducible subcomplex of the graph complex does not have the same homology as the full complex, unlike the Lie and associative cases.

Acknowledgments: In addition to the seminar participants mentioned above, we would like to thank Martin Markl, Steve Shnider and Jim Stasheff for pointing out several typos and unclear statements. Finally, the first author was partially supported by NSF VIGRE grant DMS-9983660, and the second author was partially supported by NSF grant DMS-0204185.

2 The general case - Cyclic Operads

2.1 Review of cyclic operads

Throughout the paper we will work in the category of real vector spaces. Thus an *operad* $\mathcal{O}$ is a collection of real vector spaces $\mathcal{O}[m]$, $m \geq 1$, with

1) a composition law:
$$\gamma \colon \mathcal{O}[m] \otimes \mathcal{O}[i_1] \otimes \ldots \otimes \mathcal{O}[i_m] \to \mathcal{O}[i_1 + \ldots + i_m],$$
2) a right action of the symmetric group Σ_m on $\mathcal{O}[m]$, and
3) a unit $1_{\mathcal{O}} \in \mathcal{O}[1]$

satisfying appropriate axioms governing the unit, associativity, and Σ_m-equivariance of the composition law (see [17] for a complete list of axioms). Intuitively, an element of $\mathcal{O}[m]$ is an object with m numbered input slots and an ouput slot. The symmetric group acts by permuting the numbering of the input slots. The composition $\gamma(o \otimes o_1 \otimes \ldots \otimes o_m)$ plugs the output of $o_j \in \mathcal{O}[i_j]$ into the jth input slot of o, renumbering the input slots of the result consistently (see Figure 1). If $o_i = 1_{\mathcal{O}}$ for $i \neq k$ and $o_k = o'$, we call the result the kth

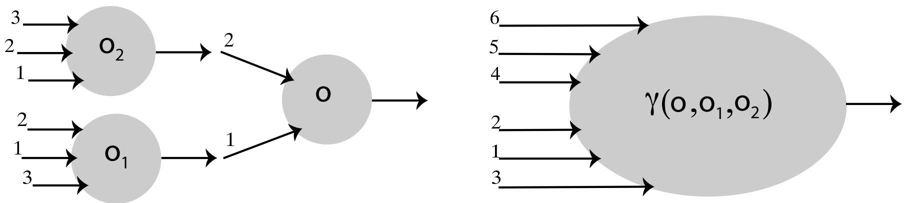

Figure 1: Operad composition

composition of o' with o.

A *cyclic operad* is an operad where the action of the symmetric group Σ_m extends to an action of Σ_{m+1} in a way compatible with the axioms. This concept was introduced in [7] (see also [17]). The intuitive idea is that in a cyclic operad, the output slot can also serve as an input slot, and any input slot can serve as the output slot. Thus, if we number the output slot as well (say with 0), then Σ_{m+1} acts by permuting the numbers on all input/output slots. Modulo the Σ_{m+1} action we can compose two elements using one input/output slot of each.

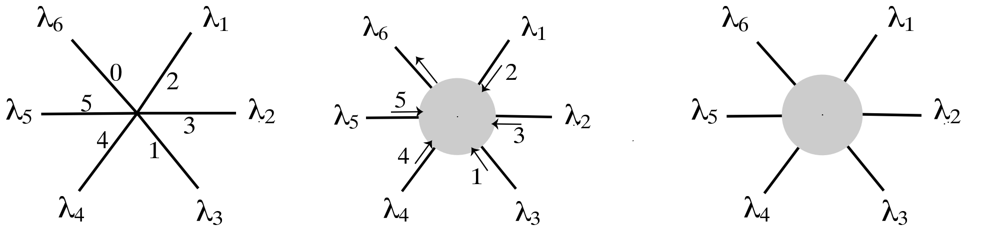

(a) Labeled 6-star

(b) Superimposing an operad element and an m-star

(c) $\mathcal{O}$-spider

Figure 2: "Coloring" an m-star with an element of $\mathcal{O}$ to make an $\mathcal{O}$-spider

We can bring the actual situation closer to the intuition as follows. For each integer $m \geq 2$, let $*_m$ be the *m-star*, the unrooted tree with one internal vertex and m leaves $\lambda_1, \dots, \lambda_m$. A *labeling* of $*_m$ is, by definition, a bijection from the set of leaves to the numbers $0, \dots, m-1$. We represent a labeling L by placing the number $L(\lambda_i)$ on the leaf λ_i, close to the internal vertex of $*_m$. Figure 2a shows a labeled 6-star.

The symmetric group Σ_m acts on the set of labelings of $*_m$, and we make the following definition.

Definition 1 Let $\mathcal{O}$ be a cyclic operad, $m \geq 2$ an integer, and $\mathcal{L}$ the set of labelings of the m-star. The space $\mathcal{OS}[m]$ of *$\mathcal{O}$-spiders with m legs* is defined to be the space of coinvariants

$$\mathcal{OS}[m] = (\bigoplus_{\mathcal{L}} \mathcal{O}[m-1])_{\Sigma_m},$$

where $(\cdot)_L : \mathcal{O}[m-1] \to \bigoplus_L \mathcal{O}[m-1]$ is the natural inclusion into the L summand, and Σ_m acts by $\sigma \cdot (o_L) = (\sigma \cdot o)_{\sigma \cdot L}$.

Every element of $\mathcal{OS}[m]$ is an equivalence class $[o_L]$ for some $o \in \mathcal{O}[m-1]$ and labeling L of $*_m$. To see this, note that $[(o_1)_{L_1} + (o_2)_{L_2}] = [(o_1)_{L_1}] + [(o_2)_{L_2}] = [(o_1)_{L_1}] + [(\sigma \cdot o_2)_{L_1}] = [(o_1 + \sigma \cdot o_2)_{L_1}]$ for $\sigma \in \Sigma_m$ with $\sigma \cdot L_2 = L_1$. We can think of o as sitting on top of $*_m$ so that the labeling of $*_m$ corresponds to the numbering of the input/output slots (Figure 2b).

Modding out by the action of Σ_m erases the labeling and the distinction between input and output slots (Figure 2c). The picture explains the arachnoid terminology for elements of $\mathcal{OS}[m]$.

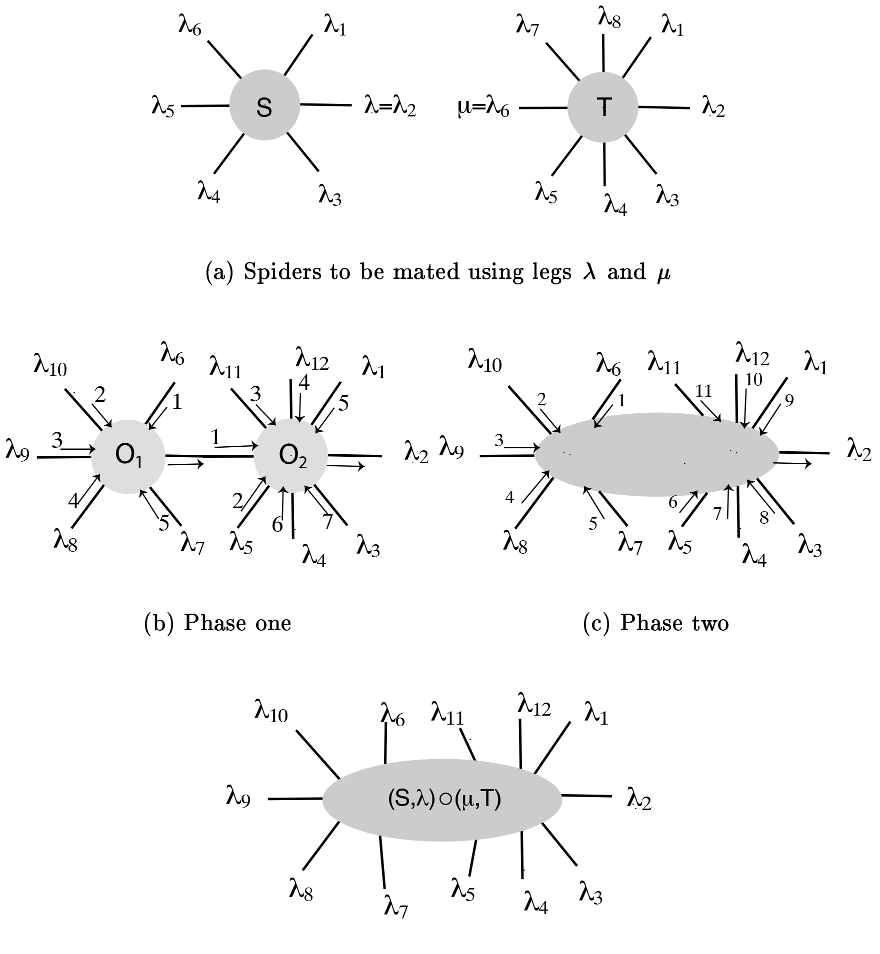

(a) Spiders to be mated using legs λ and μ

(b) Phase one

(c) Phase two

(d) Mated spiders

Figure 3: Mating spiders

The composition law in $\mathcal{O}$ transforms to a *mating law* in $\mathcal{OS} = \bigoplus_{m\geq 2} \mathcal{OS}[m]$, as follows. Consider two $\mathcal{O}$-spiders, S and T, and any pair of legs λ of S and

μ of T (Figure 3a).

Choose representatives $(o_1)_{L_1}$ for S and $(o_2)_{L_2}$ for T such that $L_1(\lambda) = 0$ (so that λ corresponds to the output slot of o_1) and $L_2(\mu) = 1$ (so μ corresponds to the first input slot of o_2). Connect the λ, μ legs together. Rename the spider legs (other than λ and μ) so that the remaining legs of S are inserted, in order, into the ordered set of legs of T, at the slot formerly occupied by μ. (Figure 3b).

Now contract the edge formed by λ and μ to get an underlying m-star, and compose the two operad elements along the corresponding input/output slots to obtain $\gamma(o_2 \otimes o_1 \otimes 1_{\mathcal{O}} \otimes \cdots \otimes 1_{\mathcal{O}})_L$, where L is the induced labeling (Figure 3c).

The resulting equivalence class under the symmetric group action will be denoted by $(S, \lambda) \circ (\mu, T)$ (Figure 3d).

2.2 Examples

We describe the three operads we will be focusing on: the commutative, Lie and associative operads. Figure 4 shows examples of spiders in these operads. There are many other cyclic operads, for example the endomorphism operad and the Poisson operad. See [7] or [17]. It is also worthwhile to note at this point that cyclic operads form a category, and there are obvious morphisms from the Lie operad to the associative operad, and from the associative to the commutative operad. See [17].

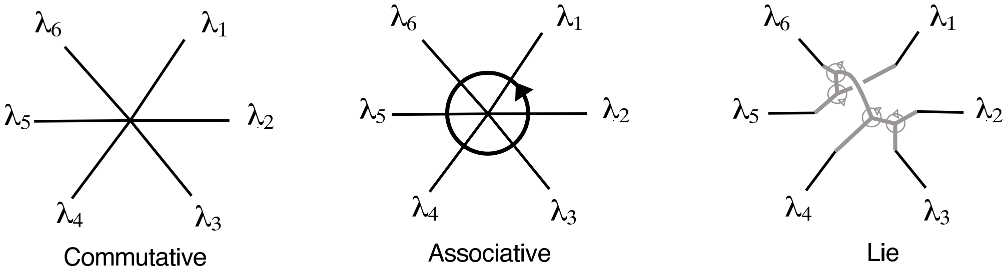

Figure 4: Three types of spiders. In the Lie case, the picture is modulo IHX and AS relations.

2.2.1 The commutative operad

In the commutative operad, each $\mathcal{O}[m]$ is 1-dimensional, with trivial Σ_m action. The composition law is given by the canonical isomorphism (i.e. multiplication)

$\mathbb{R}^{\otimes k} \cong \mathbb{R}$. An $\mathcal{O}$-spider in this case is a copy of $*_m$, weighted by a real number. Mating is done by joining legs λ and μ of two stars to form an edge, then contracting that edge to form a new star and multiplying the weights.

2.2.2 The associative operad

In the associative operad, each $\mathcal{O}[m]$ is spanned by rooted planar trees with one internal vertex and m numbered leaves. The planar embedding that each such tree comes with is equivalent to a prescribed left-to-right ordering of the leaves. The symmetric group acts by permuting the numbers of the leaves. To compose two trees, we attach the root of the first tree to a leaf of the second tree, then collapse the internal edge we just created. We order the leaves of the result so that the leaves of the first tree are inserted, in order, at the position of the chosen leaf in the second tree. A basis element of the space $\mathcal{OS}$ of $\mathcal{O}$-spiders is a copy of $*_m$ with a fixed *cyclic* ordering of the legs. To mate two basic $\mathcal{O}$-spiders using legs λ and μ, we join λ and μ to form a connected graph with one internal edge, then contract the internal edge. The cyclic order on the legs of each spider induces a cyclic ordering of the legs of the spider which results from mating. Mating is extended linearly to all spiders, i.e. to spiders whose "body" consists of a linear combination of cyclic orderings.

2.2.3 The Lie operad

In the Lie operad, $\mathcal{O}[m]$ is the vector space spanned by all rooted planar binary trees with m numbered leaves, modulo the subspace spanned by all anti-symmetry and IHX relators. Specifying a planar embedding of a tree is equivalent to giving a cyclic ordering of the edges adjacent to each interior vertex. The anti-symmetry relation AS says that switching the cyclic order at any vertex reverses the sign. The IHX (Jacobi) relation is well-known (see Figure 5).

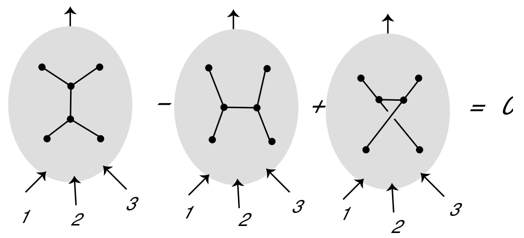

Figure 5: The IHX relator

The symmetric group Σ_m acts by permuting the numbering of the leaves. The composition rule in the operad attaches the root of one tree to a leaf of another to form an interior edge of a new planar tree, then suitably renumbers the remaining leaves.

The space of $\mathcal{O}$-spiders, $\mathcal{OS}$, is spanned by planar binary trees with m numbered leaves, modulo AS and IHX, but with no particular leaf designated as the root. Mating is accomplished by gluing two such trees together at a leaf to form a single planar binary tree, then renumbering the remaining leaves suitably. Note that mating does not involve an edge collapse, as it did in the commutative and associative cases.

2.3 Graph homology of a cyclic operad $\mathcal{O}$

In this section we construct a functor

$$\{\text{Cyclic operads}\} \to \{\text{Chain complexes}\},$$

where the chain complexes are spanned by oriented graphs with an element of $\mathcal{OS}$ attached to each vertex. We begin with a subsection discussing the appropriate notion of orientation on graphs. This subsection includes results which will be needed later when working with specific operads; the reader interested only in the basic construction can stop at the definition of orientation on first reading.

2.3.1 Oriented graphs

By a *graph* we mean a finite 1-dimensional CW complex. The set of edges of a graph X is denoted $E(X)$, the set of vertices $V(X)$ and the set of half-edges $H(X)$. Let $H(e)$ denote the set of (two) half-edges contained in an edge e. There is an involution $x \mapsto \bar{x}$ on $H(X)$, swapping the elements of $H(e)$ for each $e \in E(X)$.

For an n-dimensional vector space V, set $\det(V) := \wedge^n V$. An orientation on V can be thought of as a unit vector in $det(V)$. For a set Z, we denote by $\mathbb{R}Z$ the real vector space with basis Z.

Definition 2 An *orientation* on a graph X is a unit vector in

$$\det \mathbb{R}V(X) \otimes \bigotimes_{e \in E(X)} \det \mathbb{R}H(e).$$

In other words, *an orientation on X is determined by ordering the vertices of X and orienting each edge of X.* Two orientations are the same if they are obtained from one another by an even number of edge-orientation switches and vertex-label swaps.

Our definition is different from the definition given in Kontsevich's papers [12],[13] but, as we show below, it is equivalent for connected graphs. (Note that Kontsevich defines his graph homology using only connected graphs.) We follow ideas of Dylan Thurston [21], and begin by recording a basic observation:

Lemma 1 *Let $0 \to A \to B \to C \to D \to 0$ be an exact sequence of finite-dimensional vector spaces. Then there is a canonical isomorphism*

$$\det(A) \otimes \det(C) \to \det(B) \otimes \det(D).$$

Proof For any short exact sequence $0 \to U \xrightarrow{f} V \to W \to 0$ of finite-dimensional vector spaces, with $s\colon W \to V$ a splitting, the map

$$\det(U) \otimes \det(W) \to \det(V)$$

given by $u \otimes w \mapsto f(u) \wedge s(w)$ is an isomorphism, and is independent of the choice of s. The lemma now follows by splitting $0 \to A \to B \to C \to D \to 0$ into two short exact sequences. $\square$

Kontsevich defines an orientation on a graph to be an orientation of the vector space $H_1(X;\mathbb{R}) \oplus \mathbb{R}E(X)$. The following proposition shows that this is equivalent to our definition for connected graphs.

Proposition 1 *Let X be a connected graph. Then there is a canonical isomorphism*

$$\det(\mathbb{R}V(X)) \otimes \bigotimes_e \det(\mathbb{R}H(e)) \cong \det(H_1(X;\mathbb{R})) \otimes \det(\mathbb{R}E(X)).$$

Proof For any graph X, we have an exact sequence $0 \to H_1(X;\mathbb{R}) \to C_1(X) \to C_0(E) \to H_0(X;\mathbb{R}) \to 0$, so Lemma 1 gives a canonical isomorphism

$$\det(H_1(X;\mathbb{R})) \otimes \det(C_0(X)) \cong \det(C_1(X)) \otimes \det(H_0(X;\mathbb{R})). \tag{1}$$

$C_0(X)$ has a canonical basis consisting of the vertices of X, so that $C_0(X)$ can be identified with $\mathbb{R}V(X)$:

$$\det(C_0(X)) \cong \det(\mathbb{R}V(X)).$$

In order to give a chain in $C_1(X)$, on the other hand, you need to prescribe orientations on all of the edges, so that

$$\det(C_1(X)) \cong \det(\oplus_e \det \mathbb{R}H(e)) \cong \det(\mathbb{R}E(X)) \otimes \bigotimes_e \det(\mathbb{R}H(e)). \qquad (2)$$

The second isomorphism follows since both expressions are determined by ordering and orienting all edges.

If X is connected, $H_0(X;\mathbb{R}) \cong \mathbb{R}$ has a canonical (ordered!) basis. Combining this observation with isomorphisms (1) and (2) gives

$$\det(H_1(X;\mathbb{R})) \otimes \det(\mathbb{R}V(X)) \cong \det(\mathbb{R}E(X)) \otimes \bigotimes_e \det(\mathbb{R}H(e)). \qquad (3)$$

Now note that $\mathbb{R}E(X)$ and $\mathbb{R}V(X)$ have canonical unordered bases, which can be used to identify $\det(\mathbb{R}E(X))$ and $\det(\mathbb{R}V(X))$ with their duals. Since $V^* \otimes V$ is canonically isomorphic to $\mathbb{R}$, we can use this fact to "cancel" copies of $\det(\mathbb{R}E(X))$ or $\det(\mathbb{R}V(X))$, effectively moving them from one side of a canonical isomorphism to the other. In particular, from equation (3) we get the desired canonical isomorphism

$$\det(H_1(X;\mathbb{R})) \otimes \det(\mathbb{R}E(X)) \cong \det(\mathbb{R}V(X)) \otimes \bigotimes_e \det(\mathbb{R}H(e)). \qquad \square$$

Other equivalent notions of orientation, which we will use for particular types of graphs, are given in the next proposition and corollaries. First we record an easy but useful lemma.

Lemma 2 (Partition Lemma) *Let N be a finite set, and $P = \{P_1, \dots, P_k\}$ a partition of N. Then there is a canonical isomorphism*

$$\det(\mathbb{R}N) \cong \otimes \det(\mathbb{R}P_i) \otimes \det\Big(\bigoplus_{|P_i|\, odd} \mathbb{R}\Big),$$

which is independent of the ordering of the P_i.

Proof For x_i in N, the map "regroups" $x_1 \wedge \ldots \wedge x_{|N|}$ so that all of the x_i which are in P_1 come first, etc. $\square$

Proposition 2 *For X a connected graph, let $H(v)$ denote the set of half-edges adjacent to a vertex v of X. There is a canonical isomorphism*

$$\det(H_1(X;\mathbb{R})) \otimes \det(\mathbb{R}E(X)) \cong \bigotimes_{v \in V(X)} \det(\mathbb{R}H(v)) \otimes \det\Big(\bigoplus_{|H(v)|\, even} \mathbb{R}\Big)$$

Proof By Lemma 2, grouping half-edges according to the edges they form gives an isomorphism

$$\bigotimes_{e\in E(X)} \det(\mathbb{R}H(e)) \cong \det \mathbb{R}H(X);$$

on the other hand, grouping according to the vertices to which they are adjacent gives an isomorphism

$$\det \mathbb{R}H(X) \cong \bigotimes_{v} \det(\mathbb{R}H(v)) \otimes \det(\bigoplus_{|H(v)|odd} \mathbb{R}v).$$

Combining these isomorphisms, and substituting into the canonical isomorphism of Proposition 1, we get

$$\begin{aligned}
&\det(H_1(X;\mathbb{R})) \otimes \det(\mathbb{R}E(X)) \\
&\qquad \cong \det(\mathbb{R}V(X)) \otimes \bigotimes_{v\in V(X)} \det(\mathbb{R}H(v)) \otimes \det(\bigoplus_{|H(v)|odd} \mathbb{R}v) \\
&\qquad \cong \bigotimes_{v\in V(X)} \det(\mathbb{R}H(v)) \otimes \det(\bigoplus_{|H(v)|even} \mathbb{R}v)
\end{aligned}$$

The last isomorphism follows from the fact that

$$\det \mathbb{R}V(X) \cong \det(\bigoplus_{|H(v)|even} \mathbb{R}v) \otimes \det(\bigoplus_{|H(v)|odd} \mathbb{R}v)$$

which, in turn, follows by the partition lemma combined with the observation that a graph cannot have an odd number of vertices of odd valence and an odd number of vertices of even valence. □

Corollary 1 *Let T be a trivalent graph. Then an orientation on T is equivalent to a cyclic ordering of the edges incident to each vertex.*

Proof Since T is trivalent, all vertices have odd valence, so the orientation is determined by ordering the sets $H(v)$, up to cyclic (i.e. even) permutation, at each vertex v. □

This equivalence was mentioned in Kontsevich's papers [12],[13] and is also an ingredient in the isomorphism between commutative graph *cohomology* and the diagram algebras arising in the study of finite type invariants of three-manifolds and knots.

Corollary 2 *Let T be a connected binary tree. An orientation on T is equivalent to an ordering of the edges of T, or to a cyclic ordering of the edges incident to each interior vertex of T.*

Proof Since T is a connected binary tree, $H_1(T)$ is zero, and all of the vertices have odd valence (1 or 3), so that the isomorphism in the statement of of Proposition 2 reduces to:

$$\det(\mathbb{R}E(T)) \cong \bigotimes_v \det(\mathbb{R}H(v)) \tag{5}$$ □

Note that a cyclic ordering of the edges incident to each interior vertex of a tree can be thought of as an embedding of the tree into the plane.

2.3.2 Chain groups

We can now define the chain groups of the graph complex associated to a cyclic operad.

Definition 3 A vertex v of a graph is *$\mathcal{O}$-colored* if the half-edges incident to v are identified with the legs of an $\mathcal{O}$-spider. An *$\mathcal{O}$-graph* is an oriented graph without univalent vertices with an $\mathcal{O}$-spider coloring each vertex.

We represent an $\mathcal{O}$-graph pictorially as in Figure 6.

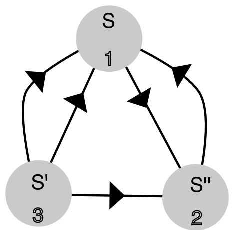

Figure 6: $\mathcal{O}$-graph

Definition 4 The *group of k-chains* $\mathcal{OG}_k$ is a quotient of the vector space spanned by $\mathcal{O}$-graphs with k vertices:

$$\mathcal{OG}_k = \mathbb{R}\{\mathcal{O}\text{-graphs with } k \text{ vertices}\}/\textit{relations}$$

where the relations are of two kinds:

(1) (Orientation) $(\mathbf{X}, or) = -(\mathbf{X}, -or)$

(2) (Vertex linearity) If a vertex v of $\mathbf{X}$ is colored by the element $S_v = aS + bT$, where $a, b \in \mathbb{R}$ and $S, T \in \mathcal{OS}[m]$, then $\mathbf{X} = a\mathbf{X}_S + b\mathbf{X}_T$, where $\mathbf{X}_S, \mathbf{X}_T$ are formed by coloring v by S and T respectively.

Thus $\mathcal{OG}_k$ is spanned by $\mathcal{O}$-graphs with k vertices, each colored by a basis element of $\mathcal{OS}$. We set

$$\mathcal{OG} = \bigoplus_{k \geq 1} \mathcal{OG}_k.$$

We also define the *reduced chain groups* $\overline{\mathcal{OG}}_k$ to be $\mathcal{OG}_k$ modulo the subspace of graphs that have at least one vertex colored by $1_{\mathcal{O}}$. In the commutative, associative and Lie cases, $\overline{\mathcal{OG}}_k$ is spanned by $\mathcal{O}$-graphs without bivalent vertices, since $\mathcal{O}[1]$ is spanned by $1_{\mathcal{O}}$ in these three cases.

2.3.3 Hopf algebra structure

Both $\mathcal{OG}$ and $\overline{\mathcal{OG}}$ have a Hopf algebra structure whose product is given by disjoint union. More precisely, $X \cdot Y$ is defined to be the disjoint union $X \cup Y$ where the orientation is given by shifting the vertex ordering of Y to lie after that of X. The coalgebra structure is defined so that connected graphs are primitive; the comultiplication is then extended multiplicatively to disjoint unions of graphs. The multiplicative unit is the empty graph, and the counit is dual to this unit. The antipode reverses the orientation of a graph.

The primitive parts (i.e. the subspaces of $\mathcal{OG}$ and $\overline{\mathcal{OG}}$ spanned by connected graphs) will be denoted $P\mathcal{OG}$ and $P\overline{\mathcal{OG}}$ respectively.

2.3.4 Boundary map

Let $\mathbf{X}$ be an $\mathcal{O}$-graph, with underlying oriented graph X, and let e be an edge of X. We define a new $\mathcal{O}$-graph $\mathbf{X}_e$ as follows. If e is a loop, then $\mathbf{X}_e$ is zero. If e has distinct endpoints v and w, then the underlying graph X_e is the graph obtained from X by collapsing e. The orientation on X_e is determined by the following rule: choose a representative of the orientation on X so that v is the first vertex, w is the second vertex, and e is oriented from v to w. The orientation on X_e is then induced from that of X: the uncollapsed edges are oriented as they were in X, the new vertex resulting from collapsing e is first in the vertex ordering, and the other vertices retain their ordering. The

$\mathcal{O}$-colorings at all the vertices besides v and w stay the same. Let S_v, S_w be the $\mathcal{O}$-spiders coloring v and w respectively, with legs λ of S_v and μ of S_w identified with the two half-edges of e. Then the $\mathcal{O}$-spider $(S_v,\lambda)\circ(\mu,S_w)$ colors the vertex obtained by collapsing e (Figure 7).

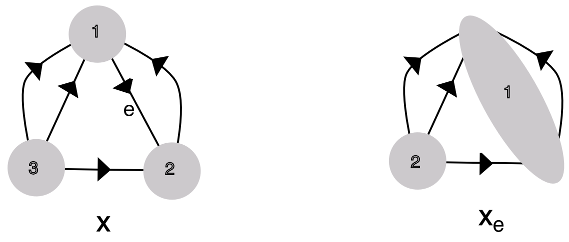

Figure 7: Edge collapse

With this orientation convention, and using the associativity axiom of operads, the map

$$\partial_E(\mathbf{X}) = \sum_{e\in X} \mathbf{X}_e$$

is a boundary operator. This makes $\mathcal{OG}$ into a chain complex, and we have

Definition 5 The *$\mathcal{O}$-graph homology* of the cyclic operad $\mathcal{O}$ is the homology of $\mathcal{OG} = \oplus\mathcal{OG}_k$ with respect to the boundary operator ∂_E.

Note that $\overline{\mathcal{OG}}$, $P\mathcal{OG}$ and $P\overline{\mathcal{OG}}$ are all chain complexes with respect to ∂_E.

We conclude this section with a nice observation, which we won't actually need, and whose proof is left to the reader.

Proposition 3 *The assignment $\mathcal{O}\mapsto\mathcal{OG}$ respects morphisms, and hence is a functor from cyclic operads to chain complexes.*

2.4 The Lie algebra of a cyclic operad and its homology

In this section we associate a sequence of Lie algebras to any cyclic operad. We show that each of these Lie algebras contains a symplectic Lie algebra as a subalgebra, and that under certain finiteness assumptions the Lie algebra homology may be computed using the subcomplex of symplectic invariants in the exterior algebra.

2.4.1 Symplectic Lie algebra of a cyclic operad

For each integer $n \geq 1$ we will define a functor {Cyclic Operads} $\to$ {Symplectic Lie Algebras}, sending $\mathcal{O}$ to $\mathcal{L}\mathcal{O}_n$. We then take a limit as $n \to \infty$ to obtain an infinite-dimensional Lie algebra $\mathcal{L}\mathcal{O}_\infty$.

Fix a $2n$-dimensional real vector space V_n basis $B_n = \{p_1, \dots, p_n, q_1, \dots, q_n\}$ corresponding to the standard symplectic form ω. Given a cyclic operad $\mathcal{O}$, the idea is to form a Lie algebra by putting elements of V_n on the legs of $\mathcal{O}$-spiders, and defining the bracket of two such objects by summing over all possible matings, with coefficients determined by the symplectic form.

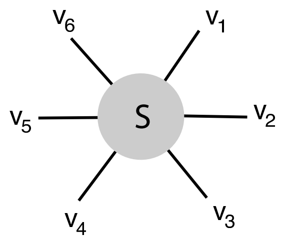

Figure 8: Symplecto-spider

Formally, we "put elements of V_n on the legs of an $\mathcal{O}$-spider" via a coinvariant construction like the one used to define spiders, i.e. we set

$$\mathcal{L}\mathcal{O}_n = \bigoplus_{m=2}^{\infty} (\mathcal{O}\mathcal{S}[m] \otimes V_n^{\otimes m})_{\Sigma_m},$$

where the symmetric group Σ_m acts simultaneously on $\mathcal{O}\mathcal{S}[m]$ and on $V_n^{\otimes m}$. We will refer to elements of $\mathcal{L}\mathcal{O}_n$ of the form $[S \otimes v_1 \otimes \ldots \otimes v_m]$, where S is an $\mathcal{O}$-spider, as *symplecto-spiders* (Figure 8).

The bracket of two symplecto-spiders is defined as follows. Let $\mathbb{S}_1 = [S_1 \otimes v_1 \otimes \ldots \otimes v_m]$ and $\mathbb{S}_2 = [S_2 \otimes w_1 \otimes \ldots \otimes w_l]$ be two symplecto-spiders. Let λ be a leg of S_1 and μ a leg of S_2, with associated elements $v_\lambda, w_\mu \in V_n$. Recalling that ω is the symplectic form on V_n, define $(\mathbb{S}_1, \lambda)!(\mu, \mathbb{S}_2)$ to be $\omega(v_\lambda, w_\mu)$ times the symplecto-spider obtained by mating S_1 and S_2 using λ and μ, erasing the elements v_λ and w_μ, and retaining the elements of V_n on the remaining legs (see Figure 9).

Now define the bracket by setting

$$[\mathbb{S}_1, \mathbb{S}_2] = \sum_{\lambda \in \mathbb{S}_1, \mu \in \mathbb{S}_2} (\mathbb{S}_1, \lambda)!(\mu, \mathbb{S}_2);$$

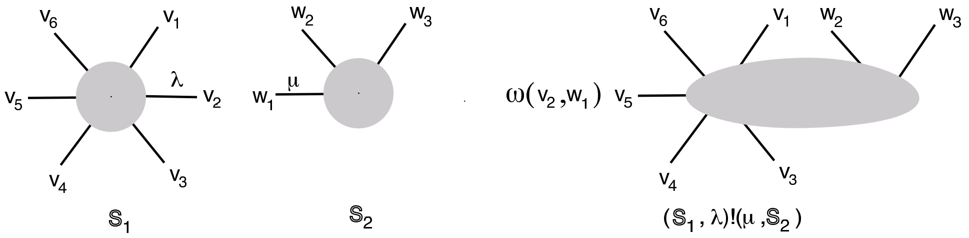

Figure 9: $(\mathbb{S}_1, \lambda)!(\mu, \mathbb{S}_2)$

and extending linearly to all of $\mathcal{LO}_n$.

Proposition 4 *The bracket is antisymmetric and satisfies the Jacobi identity, for any cyclic operad.*

Remark 1 A related construction appears in [11], Theorem 1.7.3, where, for any operad, a bracket on what might be called the space of rooted spiders is defined. One sums over all ways of plugging the root (output) of one spider into an input of another, and then subtracts the results of doing this in the other order. In the cyclic case there is no specified root, so one would have to sum over all choices of root; but then subtracting off the other order would be the same and give you a trivial operation. The needed axiom in this case is that of an *anticyclic* operad ([7, 17]), which ensures that the order of "plugging in" determines a sign. When $\mathcal{O}$ is a cyclic operad and V is a symplectic vector space, the collection $\mathcal{O}[m] \otimes V^{m+1}$ is an anticyclic operad. Thus the bracket defined here is a generalization to the anticyclic case of the one defined by Kapranov and Manin.

Remark 2 This Lie algebra structure on $\mathcal{LO}_n$ is quite natural. Let $T^{\geq 2}V_n$ denote the tensor algebra in degrees ≥ 2. $T^{\geq 2}V_n$ has a Lie bracket induced by the symplectic form, and we can give $\mathcal{OS}$ an abelian Lie algebra structure. Give the tensor product of associative algebras $\mathcal{OS} \otimes T^{\geq 2}V_n$ the natural bracket which is a derivation in each variable, and which extends the brackets on each tensor factor. Then the natural map $\mathcal{OS} \otimes T^{\geq 2}V_n \to \mathcal{LO}_n$ is a Lie algebra homomorphism.

In the commutative case the Lie algebra $\mathcal{LO}_n$ may be identified with the Lie algebra of polynomials with no constant or linear term in the variables $p_1 \ldots p_n, q_1, \ldots, q_n$, under the standard Poisson bracket of functions.

In the Lie case the Lie algebra $\mathcal{L}\mathcal{O}_n$ is isomorphic to the Lie algebra $D_*(V_n)\otimes\mathbb{R}$ that has arisen in the study of the mapping class group (see [14], [18]). Here $D_k(V_n)$ is defined to be the kernel of the map $V_n\otimes L_k(V_n)\to L_{k+1}(V_n)$ sending $v\otimes x\mapsto [v,x]$, where $L_*(V_n)$ refers to the free Lie algebra on V_n.

We record the functoriality of our construction without proof.

Proposition 5 *For each $n\geq 1$, the assignment $\mathcal{O}\mapsto\mathcal{L}\mathcal{O}_n$ respects morphisms, and hence is a functor from cyclic operads to Lie algebras.*

Note that bracketing a symplecto-spider with 2 legs and one with m legs results in a sum of (at most two) symplecto-spiders, each with m legs. In particular, the subspace of $\mathcal{L}\mathcal{O}_n$ spanned by symplecto-spiders with two legs forms a Lie subalgebra of $\mathcal{L}\mathcal{O}_n$. If we consider only (two-legged) symplecto-spiders with vertex colored by the identity $1_{\mathcal{O}}$, we obtain an even smaller subalgebra, denoted $\mathcal{L}\mathcal{O}_n^0$. The next proposition identifies $\mathcal{L}\mathcal{O}_n^0$ with the symplectic Lie algebra $\mathfrak{sp}(2n)$.

Proposition 6 *Let S_0 denote the (unique) $\mathcal{O}$-spider colored by the identity element $1_{\mathcal{O}}$. The subspace $\mathcal{L}\mathcal{O}_n^0$ of $\mathcal{L}\mathcal{O}_n$ spanned by symplecto-spiders of the form $[S_0\otimes v\otimes w]$ is a Lie subalgebra isomorphic to $\mathfrak{sp}(2n)$.*

Proof The map $\mathcal{L}\mathcal{O}_n^0\to S^2V$ sending $[S\otimes v\otimes w]$ to vw is easily checked to be a Lie algebra isomorphism, where the bracket on S^2V is the Poisson bracket.

Recall that $\mathfrak{sp}(2n)$ is the set of $2n\times 2n$ matrices A satisfying $AJ+JA^T=0$, where $J=\begin{pmatrix}0 & I\\ -I & 0\end{pmatrix}$. The symmetric algebra S^2V can be identified with the subspace of $V\otimes V$ spanned by elements of the form $v\otimes w+w\otimes v$. Consider the composition of isomorphisms

$$V\otimes V\to V^*\otimes V\to Hom(V,V)$$

where the first map is induced by the isomorphism $V\to V^*$ given by $v\mapsto \omega(v,\_)$. Tracing through these isomorphisms, we see that

$$p_iq_j\leftrightarrow p_i\otimes q_j+q_j\otimes p_i\mapsto\begin{pmatrix}-E_{ji} & 0\\ 0 & E_{ij}\end{pmatrix}$$
$$p_ip_j\leftrightarrow p_i\otimes p_j+p_j\otimes p_i\mapsto\begin{pmatrix}0 & 0\\ E_{ij}+E_{ji} & 0\end{pmatrix}$$
$$q_iq_j\leftrightarrow q_i\otimes q_j+q_j\otimes q_i\mapsto\begin{pmatrix}0 & -E_{ij}-E_{ji}\\ 0 & 0\end{pmatrix}$$

where E_{ij} is the $n\times n$ matrix with (i,j)-entry equal to 1 and zeroes elsewhere. It is now straightforward to check that this gives a Lie algebra isomorphism $S^2V \to \mathfrak{sp}(2n)$. $\square$

The subalgebra $\mathcal{LO}_n^0 \cong \mathfrak{sp}(2n)$ acts on $\mathcal{LO}_n$ via the bracket (i.e. the adjoint action). Using the remark following Proposition 4, we see that the $\mathfrak{sp}(2n)$ action on $(\mathcal{OS}[m]\otimes V_n^{\otimes m})_{\Sigma_m}$ is given by $\xi\cdot[S\otimes v_1\otimes\cdots\otimes v_m] = \sum_{i=1}^m [S\otimes v_1\otimes\cdots\otimes(\xi\cdot v_i)\otimes\cdots\otimes v_m]$, where V_n has the standard $\mathfrak{sp}(2n)$-module structure.

The natural inclusion $V_n\to V_{n+1}$ induces an inclusion $\mathcal{LO}_n\to\mathcal{LO}_{n+1}$ of Lie algebras, which is compatible with the inclusion $\mathfrak{sp}(2n)\to\mathfrak{sp}(2(n+1))$.

Definition 6 The infinite dimensional symplectic Lie algebra $\mathcal{LO}_\infty$ is the direct limit

$$\mathcal{LO}_\infty = \lim_{n\to\infty}\mathcal{LO}_n.$$

2.4.2 Lie algebra homology

The Lie algebra homology of $\mathcal{LO}_n$ is computed from the exterior algebra $\wedge\mathcal{LO}_n$ using the standard Lie boundary operator $\partial_n\colon \wedge^k\ \mathcal{LO}_n\to\wedge^{k-1}\mathcal{LO}_n$ defined by

$$\partial_n(\mathbb{S}_1\wedge\ldots\wedge\mathbb{S}_k) = \sum_{i<j}(-1)^{i+j+1}[\mathbb{S}_i,\mathbb{S}_j]\wedge\mathbb{S}_1\wedge\ldots\widehat{\mathbb{S}_i}\wedge\ldots\wedge\widehat{\mathbb{S}_j}\wedge\ldots\wedge\mathbb{S}_k.$$

The map $\wedge\mathcal{LO}_n\to\wedge\mathcal{LO}_{n+1}$ induced by the natural inclusion is a chain map, so that

$$H_k(\mathcal{LO}_\infty;\mathbb{R}) = \lim_{n\to\infty}H_k(\mathcal{LO}_n;\mathbb{R}).$$

Proposition 7 *$H_k(\mathcal{LO}_\infty;\mathbb{R})$ has the structure of a Hopf algebra.*

Proof To define the product $H_*(\mathcal{LO}_\infty)\otimes H_*(\mathcal{LO}_\infty)\to H_*(\mathcal{LO}_\infty)$, consider maps $E\colon B_\infty\to B_\infty$ sending $p_i\mapsto p_{2i}$ (resp. $q_i\mapsto q_{2i}$), and $O\colon B_\infty\to B_\infty$ sending $p_i\mapsto p_{2i-1}$ (resp $q_i\mapsto q_{2i-1}$). These induce maps E and O on $\mathcal{LO}_\infty$. The product on $H_*(\mathcal{LO}_\infty)$ is induced by the map

$$\mathcal{LO}_\infty\oplus\mathcal{LO}_\infty\to\mathcal{LO}_\infty$$

which sends $x\oplus y$ to $E(x)+O(y)$.

The coproduct $H_*(\mathcal{LO}_\infty) \to H_*(\mathcal{LO}_\infty) \otimes H_*(\mathcal{LO}_\infty)$ is induced by the diagonal map $\mathcal{LO}_\infty \to \mathcal{LO}_\infty \oplus \mathcal{LO}_\infty$. More explicitly, the coproduct is induced by the map $\wedge\mathcal{LO}_\infty \to \wedge\mathcal{LO}_\infty \otimes \wedge\mathcal{LO}_\infty$ sending

$$\mathbb{S}_1 \wedge \mathbb{S}_2 \wedge \cdots \wedge \mathbb{S}_k \mapsto \sum_{[k]=I\cup J} \epsilon(I,J)\mathbb{S}_I \otimes \mathbb{S}_J,$$

where the sum is over unordered partitions of $[k] = \{1,\ldots,k\}$, where $\mathbb{S}_I = \mathbb{S}_{i_1} \wedge \ldots \wedge \mathbb{S}_{i_{|I|}}$ if I consists of $i_1 < i_2 < \ldots < \ldots i_{|I|}$, and where $\epsilon(I,J)$ is a sign determined by the equation $\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k = \epsilon(I,J)\mathbb{S}_I \wedge \mathbb{S}_J$.

The unit is $1 \in \mathbb{R} \cong \wedge^0\mathcal{LO}_\infty$, and the counit is dual to this. □

2.4.3 The subcomplex of $\mathfrak{sp}(2n)$-invariants

In the remainder of this paper, we assume that the vector spaces $\mathcal{O}[m]$ are finite-dimensional. In this section we show that in this case the homology of $\mathcal{LO}_n$ is computed by the subcomplex $(\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)}$ of $\mathfrak{sp}(2n)$ invariants (where an $\mathfrak{sp}(2n)$ "invariant" is an element which is killed by every element of $\mathfrak{sp}(2n)$.)

In general, the exterior algebra $\wedge\mathcal{LO}_n$ breaks up into a direct sum of pieces $\Lambda_{k,m}$, spanned by wedges of k symplecto-spiders with a total of m legs:

$$\begin{aligned}\wedge\mathcal{LO}_n &= \bigoplus_{k,m} \Lambda_{k,m}\\ &= \bigoplus_{k,m}\Big(\bigoplus_{m_1+\ldots+m_k=mm+} (\mathcal{OS}[m_1]\otimes V^{\otimes m_1})_{S_{m_1}} \wedge \ldots \wedge (\mathcal{OS}[m_k]\otimes V^{\otimes m_k})_{S_{m_k}}\Big).\end{aligned}$$

If the vector spaces $\mathcal{O}[m]$ are finite-dimensional, then these pieces $\Lambda_{k,m}$ are all finite-dimensional, as are the following subspaces:

Definition 7 The *(k,m)-cycles* $Z_{k,m}$, the *(k,m)-boundaries* $Z_{k,m}$ and the *(k,m)-homology* $H_{k,m}(\mathcal{LO}_n;\mathbb{R})$ are defined by

$$Z_{k,m} = \Lambda_{k,m}\cap ker(\partial_n), \quad B_{k,m} = \Lambda_{k,m}\cap im(\partial_n), \quad H_{k,m}(\mathcal{LO}_n;\mathbb{R}) = Z_{k,m}/B_{k,m}.$$

With this definition, we have

$$H_k(\mathcal{LO}_n;\mathbb{R}) = \bigoplus_m H_{k,m}(\mathcal{LO}_n;\mathbb{R}).$$

Proposition 8 *The invariants $(\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)}$ form a subcomplex of $\wedge\mathcal{LO}_n$. If $\mathcal{O}[m]$ is finite dimensional for all m, then the inclusion $(\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)} \to \wedge\mathcal{LO}_n$ is an isomorphism on homology.*

Proof The first statement follows since ∂_n is an $\mathfrak{sp}(2n)$-module morphism.

The proof of the second statement depends on the following remark: every finite dimensional $\mathfrak{sp}(2n)$-module, E, decomposes as a direct sum of $\mathfrak{sp}(2n)$-modules in the following way:

$$E = E^{\mathfrak{sp}(2n)} \oplus \mathfrak{sp}(2n) \cdot E.$$

(Proof: $\mathfrak{sp}(2n)$ is well-known to be reductive, which means that for every finite dimensional module E, every submodule $E' \subset E$ has a complementary submodule E'' with $E = E' \oplus E''$. Thus $E = E^{\mathfrak{sp}(2n)} \oplus E''$. Now E'' is a direct sum of simple modules on which $\mathfrak{sp}(2n)$ acts nontrivially. Therefore, $\mathfrak{sp}(2n) \cdot (E^{\mathfrak{sp}(2n)} \oplus E'') = \mathfrak{sp}(2n) \cdot E'' = E''$, and the proof is complete.)

By hypothesis $\mathcal{O}[m]$ is finite-dimensional, so that $\Lambda_{k,m}$ is a finite dimensional $\mathfrak{sp}(2n)$-module. Since $\mathfrak{sp}(2n)$ is simple, this means that $\Lambda_{k,m}$ decomposes as a direct sum $\Lambda_{k,m} = \Lambda_{k,m}^{\mathfrak{sp}(2n)} \oplus \mathfrak{sp}(2n) \cdot \Lambda_{k,m}$. Since the boundary is an $\mathfrak{sp}(2n)$ module morphism, the space of (k,m)-cycles $Z_{k,m}$ and the space of (k,m)-boundaries $B_{k,m}$ are both $\mathfrak{sp}(2n)$ modules. Therefore, they decompose as $Z_{k,m}^{\mathfrak{sp}(2n)} \oplus \mathfrak{sp}(2n) \cdot Z_{k,m}$ and $B_{k,m}^{\mathfrak{sp}(2n)} \oplus \mathfrak{sp}(2n) \cdot B_{k,m}$ respectively. Thus the homology

$$\begin{aligned} H_k(\mathcal{L}\mathcal{O}_n) &= \bigoplus_m H_{k,m}(\mathcal{L}\mathcal{O}_n) \\ &= \bigoplus_m Z_{k,m}/B_{k,m} \\ &= \bigoplus_m \frac{Z_{k,m}^{\mathfrak{sp}(2n)}}{B_{k,m}^{\mathfrak{sp}(2n)}} \oplus \frac{\mathfrak{sp}(2n) \cdot Z_{k,m}}{\mathfrak{sp}(2n) \cdot B_{k,m}} \end{aligned}$$

Hence it suffices to show that $\mathfrak{sp}(2n) \cdot Z_{k,m} = \mathfrak{sp}(2n) \cdot B_{k,m}$. For every $\xi \in \mathfrak{sp}(2n)$ define $i_\xi^k : \wedge^{k-1}\mathcal{L}\mathcal{O}_n \to \wedge^k \mathcal{L}\mathcal{O}_n$ by $a \mapsto \xi \wedge a$. Then one easily checks that for $a \in \wedge^k \mathcal{L}\mathcal{O}_n$, $\xi \cdot a = (\partial_n i_\xi^{k+1} + i_\xi^k \partial_n)(x)$. Thus, if $\xi \in \mathfrak{sp}(2n)$ and $z \in Z_{k,m}$, then $\xi \cdot z = \partial_n i_\xi^{k+1}(z)$. Hence $\mathfrak{sp}(2n) \cdot Z_{k,m} \subset B_{k,m}$, which implies $[\mathfrak{sp}(2n), \mathfrak{sp}(2n)] \cdot Z_{k,m} \subset \mathfrak{sp}(2n) \cdot B_{k,m}$. Since $\mathfrak{sp}(2n)$ is simple, $[\mathfrak{sp}(2n), \mathfrak{sp}(2n)] = \mathfrak{sp}(2n)$ and the proof is complete. □

2.5 Relation between graph homology and Lie algebra homology

Now we describe the construction at the heart of Kontsevich's proof, namely the identification of $\mathfrak{sp}(2n)$ invariants with oriented graphs. There are three principle players that need to be introduced:

$$\phi_n\colon \mathcal{OG} \to \wedge\mathcal{LO}_n$$
$$\psi_n\colon\ \wedge\,\mathcal{LO}_n \to \mathcal{OG}$$
$$M_n\colon \mathcal{OG} \to \mathcal{OG}$$

2.5.1 The map ϕ_n

We need to produce wedges of symplecto-spiders from an $\mathcal{O}$-graph $\mathbf{X}$. Let X be the underlying oriented graph, and fix an ordering of the vertices and directions on the edges representing the orientation. A *state* of $\mathbf{X}$ is an assignment of an element of $B_n = \{p_1, \ldots, p_n, q_1 \ldots, q_n\}$ to each half-edge of X and a sign ± 1 to each edge, subject to the following constraints:

- if one half-edge of e is labeled p_i, the other half-edge must be labeled q_i, and vice versa;
- the sign on an edge is positive if the initial half-edge is labeled p_i, and negative if the initial half-edge is labeled q_i.

Given a state s, The *sign* of s, denoted $\sigma(s)$, is the product of the signs on all edges, and we define an element $\mathbf{X}\{s\}$ of $\wedge\mathcal{LO}_n$ as follows. Cut each edge at its midpoint, thereby separating $\mathbf{X}$ into a disjoint union of symplecto-spiders $\mathbb{S}_1, \ldots, \mathbb{S}_k$ (where the subscript comes from the vertex ordering on X), and set $\mathbf{X}\{s\} = \mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k$ (See Figure 10).

Now define $\phi_n(\mathbf{X})$ by summing over all possible states of $\mathbf{X}$:

$$\phi_n(\mathbf{X}) = \sum_s \sigma(s)\mathbf{X}\{s\}.$$

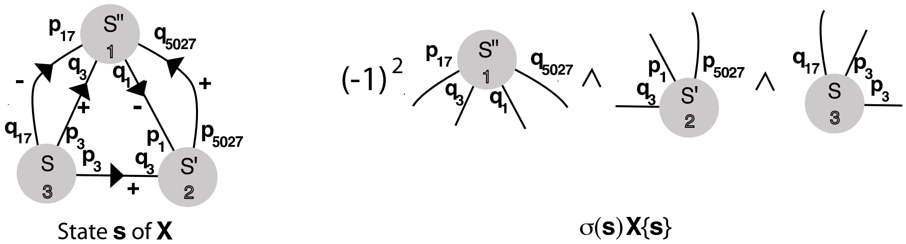

Figure 10: State of $\mathbf{X}$ and corresponding term of $\phi_n(\mathbf{X})$

2.5.2 The map ψ_n

We define ψ_n, too, as a state sum. This time we need to produce $\mathcal{O}$-graphs from a wedge of symplecto-spiders $\mathbb{S}_1 \wedge \mathbb{S}_2 \wedge \ldots \wedge \mathbb{S}_k$. In this case a state π is a pairing of the legs of the spiders. We obtain a new $\mathcal{O}$-graph $(\mathbb{S}_1 \wedge \cdots \wedge \mathbb{S}_k)^\pi$ by gluing the spider legs together according to π, and orienting each edge arbitrarily. Each edge of $(\mathbb{S}_1 \wedge \cdots \wedge \mathbb{S}_k)^\pi$ carries an element $v_1 \in V_n$ on its initial half-edge and $v_2 \in V_n$ on its terminal half-edge. We define the *weight* of this edge to be $\omega(v_1, v_2)$, and denote the product of the weights of all edges by $w(\pi)$ (see Figure 11). With this definition, the product $w(\pi)(\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^\pi$ is independent

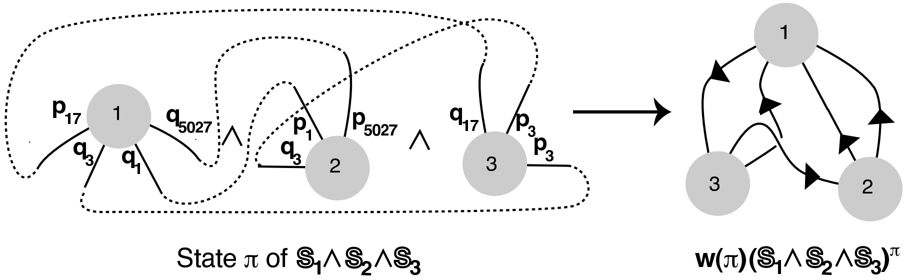

Figure 11: A pairing π and the resulting $\mathcal{O}$-graph

of the choice of edge-orientations, and we define

$$\psi_n(\mathbb{S}_1 \wedge \mathbb{S}_2 \wedge \ldots \wedge \mathbb{S}_k) = \sum_\pi w(\pi)(\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^\pi,$$

where the sum is over all possible pairings π. Note that the vertex-linearity axiom in the definition of $\mathcal{OG}_k$ is required here to make ψ_n linear.

2.5.3 The map M_n

As in the definition of ψ_n, a pairing π of the half-edges $H(X)$ of a graph X determines a new graph X^π, obtained by cutting all edges of X, then re-gluing the half-edges according to π. The *standard pairing* σ pairs x with $\bar{x}$ for all half-edges x of X; X^σ is of course just X.

If X is oriented, there is an induced orientation on X^π, given as follows. A pairing can be represented by a chord diagram on a set of vertices labeled by the half-edges of X. The union of the chord diagrams for π and for the standard pairing σ forms a one-dimensional closed manifold $C(\pi)$, a union of circles. Choose a representative for the orientation of X so that the chords from the initial half-edge of each edge e to the terminal half-edge of e are oriented coherently in each of these circles. Now each edge of X^π inherits a natural orientation from each pair of half-edges determined by π. The ordering of the vertices of X^π is inherited from X. If X is the underlying graph of an $\mathcal{O}$-graph $\mathbf{X}$, we let $\mathbf{X}^\pi$ be the induced $\mathcal{O}$-graph based on X^π (see Figure 12).

Define $c(\pi)$ to be the number of components of $C(\pi)$. and define the map $M_n\colon \mathcal{OG} \to \mathcal{OG}$ by

$$M_n(\mathbf{X}) = \sum_\pi (2n)^{c(\pi)} \mathbf{X}^\pi,$$

where the sum is over all possible pairings π of $H(X)$.

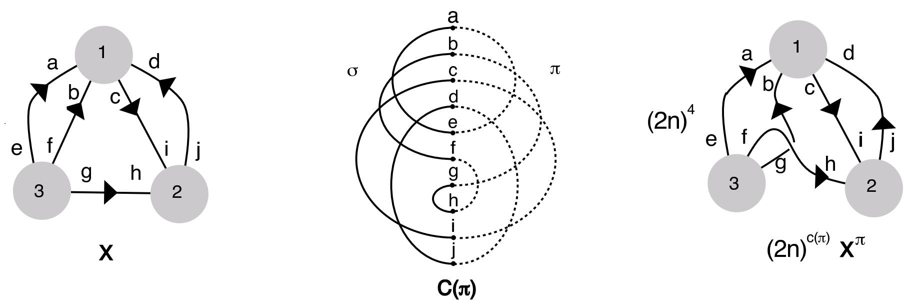

Figure 12: A term in $M_n(\mathbf{X})$

The map M_n decomposes as a direct sum as follows. Write

$$\mathcal{OG} = \bigoplus_{k,m} \mathcal{OG}_{k,m}$$

where $\mathcal{OG}_{k,m}$ is spanned by $\mathcal{O}$-graphs with k vertices and m half-edges (i.e $m/2$ edges). The subspaces $\mathcal{OG}_{k,m}$ are invariant under M_n, and we denote by $M_n^{k,m}$ the restriction of M_n to $\mathcal{OG}_{k,m}$.

Proposition 9 *If $\mathcal{O}[m]$ is finite-dimensional, then for large enough n, the restriction*

$$M_n^{k,m}\colon \mathcal{OG}_{k,m} \to \mathcal{OG}_{k,m}$$

of M_n is an isomorphism.

Proof If $\mathcal{O}[m]$ is finite-dimensional, then so is $\mathcal{OG}_{k,m}$, and we can think of the restriction $M_n^{k,m}$ of M_n to $\mathcal{OG}_{k,m}$, as a matrix. The matrix entries are polynomials in n. The maximum of $c(\pi)$ occurs when π is the standard pairing

σ, and is equal to m. Thus the diagonal entries are of the form $(2n)^m +$ *(lower degree terms)*, whereas the off diagonal terms are all of lower degree. Therefore, for large enough n, the matrix is invertible, i.e. the map $M_n^{k,m}$ is an isomorphism. □

The three maps are related by the following Proposition.

Proposition 10 $\psi_n \circ \phi_n = M_n$

Proof Applying ϕ_n to an $\mathcal{O}$-graph $\mathbf{X}$ means that we are assigning elements of B_n to the endpoints of the chord diagram for the standard pairing σ of $H(X)$, as on the left of Figure 13. Each chord must connect a pair $\{p_i, q_i\}$, for some

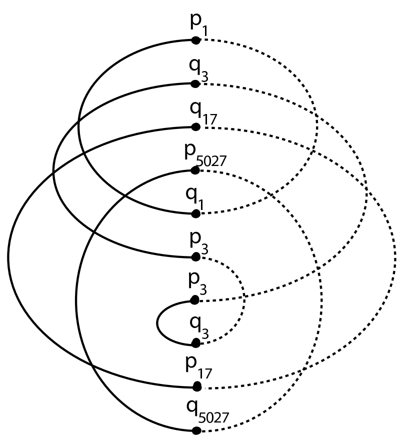

Figure 13: Proof of Proposition 10

i. To then apply ψ_n, we consider all possible pairings π of $H(X)$, and reglue to get $\mathcal{O}$-graphs $\mathbf{X}^\pi$. The weight $w(\pi)$ will only be non-zero if every chord for π also connects a pair $\{p_i, q_i\}$ for some i, as on the right of Figure 13. Thus the label of a vertex in $C(\pi)$ determines the labels of all the vertices of each connected component of $C(\pi)$: they must alternate $p_i, q_i, p_i, q_i, \ldots$ as you travel around a circuit. There are $2n = |B_n|$ choices of label for each component, so there are $(2n)^{c(\pi)}$ possible pairings π with non-zero weight. Keeping track of the orientations, we see that each of these terms has weight 1, so that the composition $\psi_n \circ \phi_n$ is exactly equal to M_n. □

Recall that $\wedge\mathcal{LO}_n = \bigoplus \Lambda_{k,m}$, where $\Lambda_{k,m}$ is spanned by wedges of k symplecto-spiders with a total of m legs. Let $\psi_{k,m}$ denote the restriction of ψ to $\Lambda_{k,m}$, and notice that $\psi_{k,m}(\Lambda_{k,m}) \subset \mathcal{OG}_{k,m}$. Similarly denote the restriction of ϕ_n to $\Lambda_{k,m}$ by $\phi_{k,m}$ and notice that $\phi_{k,m}(\Lambda_{k,m}) \subset \mathcal{OG}_{k,m}$.

Corollary 3 *For n sufficiently large with respect to fixed k and m,*

i) *the map $\psi_{k,m}\colon \Lambda_{k,m} \to \mathcal{OG}_{k,m}$ is onto, and*

ii) *the map $\phi_{k,m}\colon \mathcal{OG}_{k,m} \to \Lambda_{k,m}$ is injective.*

2.5.4 Graphs and invariants

The following proposition shows that the map ϕ_n gives instructions for constructing an $\mathfrak{sp}(2n)$-invariant from an $\mathcal{O}$-graph, and that all $\mathfrak{sp}(2n)$-invariants are constructed in this way.

Proposition 11 $im(\phi_n) = (\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)}$

Proof In order to determine the $\mathfrak{sp}(2n)$-invariants in $\wedge\mathcal{LO}_n$, we first lift to the tensor algebra $T(\mathcal{LO}_n)$. The quotient map $p\colon T(\mathcal{LO}_n) \to \wedge\mathcal{LO}_n$ sending $\mathbb{S}_1 \otimes \ldots \otimes \mathbb{S}_k \mapsto \mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k$ is an $\mathfrak{sp}(2n)$-module map, so restricts to $p\colon (T(\mathcal{LO}_n))^{\mathfrak{sp}(2n)} \to (\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)}$. To see that p is surjective, note that composition

$$\wedge\mathcal{LO}_n \xrightarrow{i} T\mathcal{LO}_n \xrightarrow{p} \wedge\mathcal{LO}_n,$$

where the map i is defined by $i(\mathbb{S}_1\wedge\ldots\wedge\mathbb{S}_k) = \frac{1}{k!}\sum_{\sigma\in\Sigma_m} sgn(\sigma)\mathbb{S}_{\sigma(1)}\otimes\ldots\ldots\otimes\mathbb{S}_{\sigma(k)}$, is the identity. Since i is also an $\mathfrak{sp}(2n)$-module homomorphism, we get restrictions

$$(\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)} \xrightarrow{i} (T(\mathcal{LO}_n))^{\mathfrak{sp}(2n)} \xrightarrow{p} (\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)}$$

whose composition is the identity. In particular the restriction of p is onto.

We next lift even further. Recall that $\mathcal{LO}_n = \oplus_{m=2}^{\infty}(\mathcal{OS}[m] \otimes V^{\otimes m})_{\Sigma_m}$, where $V = V_n$. Define $\widehat{\mathcal{LO}}_n$ to be $\oplus_{m=2}^{\infty}(\mathcal{OS}[m] \otimes V^{\otimes m})$. The map $q\colon T(\widehat{\mathcal{LO}}_n) \to T(\mathcal{LO}_n)$ induced by the quotient maps $(\mathcal{OS}[m] \otimes V^{\otimes m}) \to (\mathcal{OS}[m] \otimes V^{\otimes m})_{\Sigma_m}$ is an $\mathfrak{sp}(2n)$ module map, so restricts to $q\colon T(\widehat{\mathcal{LO}}_n)^{\mathfrak{sp}(2n)} \to T(\mathcal{LO}_n)^{\mathfrak{sp}(2n)}$. To see that this is surjective, consider the composition

$$T(\mathcal{LO}_n) \xrightarrow{j} T(\widehat{\mathcal{LO}}_n) \xrightarrow{q} T(\mathcal{LO}_n),$$

where j is induced by the maps

$$(\mathcal{OS}[m] \otimes V^{\otimes m})_{\Sigma_m} \to \mathcal{OS}[m] \otimes V^{\otimes m}$$

sending x to $\frac{1}{m!}\sum_{\sigma\in\Sigma_m}\sigma\cdot x$. These maps are well defined (they send x and $\sigma\cdot x$ to the same element) and the composition $q\circ j$ is the identity. In addition, j is an $\mathfrak{sp}(2n)$-module homomorphism, so the restriction

$$(T(\mathcal{LO}_n))^{\mathfrak{sp}(2n)} \xrightarrow{j} (T(\widehat{\mathcal{LO}_n}))^{\mathfrak{sp}(2n)} \xrightarrow{q} (T(\mathcal{LO}_n))^{\mathfrak{sp}(2n)}$$

is the identity; in particular, the restriction of q is onto.

We now compute the invariants in $T(\widehat{\mathcal{LO}_n})$. We have

$$T(\widehat{\mathcal{LO}_n}) = \bigoplus_{k\geq 2, m\geq 1} \widehat{T}_{k,m},$$

where

$$\widehat{T}_{k,m} = \bigoplus_{m_1+\cdots+m_k=m} (\mathcal{OS}[m_1]\otimes V^{\otimes m_1})\otimes\cdots\otimes(\mathcal{OS}[m_k]\otimes V^{\otimes m_k})$$

The action of $\mathfrak{sp}(2n)$ on each $\mathcal{OS}[k]\otimes V^{\otimes k}$ affects only the $V^{\otimes k}$ factor, so that

$$(\widehat{T}_{k,m})^{\mathfrak{sp}(2n)} = \bigoplus_{m_1+\cdots+m_k=m} \Big((V^{\otimes m_1}\otimes\cdots\otimes V^{\otimes m_k})^{\mathfrak{sp}(2n)}\otimes(\mathcal{OS}[m_1]\otimes\cdots\otimes\mathcal{OS}[m_k])\Big)$$

Computation of the invariants $(V^{\otimes N})^{\mathfrak{sp}(2n)}$ is a classical result of Weyl [24]. A convenient way of parameterizing these is by oriented chord diagrams. Consider a set of N vertices representing the tensor factors of $V^{\otimes N}$. An oriented chord diagram is a choice of directed edges that pair up the vertices; such a diagram gives rise to an invariant as follows. As in the definition of ϕ_n, define a *state* to be an assignment of pairs $\{p_i,q_i\}$ to the vertices joined by each chord. The invariant is then the sum, over all possible signed states, of the associated elements of $V^{\otimes N}$. For example, when $N=4$, consider the chord diagram $1\to 3, 4\to 2$. One state for this diagram assigns p_1 to the first vertex, q_1 to the third vertex, q_6 to the fourth vertex and p_6 to the second vertex. The corresponding term of $V^{\otimes 4}$ is $-p_1\otimes p_6\otimes q_1\otimes q_6$. The chord $4\to 2$ contributes a minus sign, since the p_6 occurs at the head instead of the tail of the chord.

Having identified the invariants in $T(\widehat{\mathcal{LO}_n})$, we follow the surjective maps

$$(T(\widehat{\mathcal{LO}_n}))^{\mathfrak{sp}(2n)} \to (T(\mathcal{LO}_n))^{\mathfrak{sp}(2n)} \to (\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)}$$

to identify the invariants in $\wedge\mathcal{LO}_n$. Let $c\otimes(S_1\otimes\ldots\otimes S_k)\in(V_n^{\otimes m_1}\otimes\cdots\otimes V_n^{\otimes m_k})^{\mathfrak{sp}(2n)}\otimes(\mathcal{OS}[m_1]\otimes\cdots\otimes\mathcal{OS}[m_k])$ where c is a chord diagram invariant. Mapping to

$$(\mathcal{OS}[m_1]\otimes V_n^{\otimes m_1})_{\Sigma_{m_1}}\otimes\cdots\otimes(\mathcal{OS}[m_k]\otimes V_n^{\otimes m_k})_{\Sigma_{m_k}}$$

has the effect of putting the vectors p_i, q_i in the terms of c on the legs of the spiders $S_1, \ldots, S_k$. The chord diagram c gives instructions for gluing the spider legs to form an $\mathcal{O}$-graph $\mathbf{X}$, and each state of c induces a state of $\mathbf{X}$. Summing over all states of c, we obtain the sum over all states of $\mathbf{X}$. Mapping then to $\wedge\mathcal{LO}_n$ explains the antisymmetry of the vertex ordering of $\mathbf{X}$.

The invariant just described is precisely $\phi_n(\mathbf{X})$. Since these invariants span the space of all invariants, the proposition is proved. □

Corollary 4 *For large enough n with respect to fixed k, m, the restriction*

$$\psi_{k,m} \colon (\Lambda_{k,m})^{\mathfrak{sp}(2n)} \to \mathcal{OG}_{k,m}$$

is an isomorphism.

Proof This is immediate from Propositions 9, 10 and 11. □

2.5.5 Isomorphism with graph homology

The map ϕ_n is *not* a chain map, so does not induce an isomorphism on homology; this was an oversight in Kontsevich's original papers [12],[13]. We show instead that ψ_n is a chain map, and induces an isomorphism on homology after stabilization with respect to n.

Proposition 12 *$\psi_n \colon \Lambda\mathcal{LO}_n \to \mathcal{OG}$ is a chain map.*

Proof We need to show the following diagram commutes.

$$\begin{array}{ccc} \wedge\mathcal{LO}_n & \xrightarrow{\psi_n} & \mathcal{OG} \\ {\scriptstyle\partial_n}\downarrow & & \downarrow{\scriptstyle\partial_E} \\ \wedge\mathcal{LO}_n & \xrightarrow{\psi_n} & \mathcal{OG} \end{array}$$

Let $\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k \in \wedge^k \mathcal{LO}_n$. Then

$$
\begin{aligned}
&\psi_n \partial_n (\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k) \\
&\quad = \psi_n \left(\sum_{i<j} (-1)^{i+j+1} [\mathbb{S}_i, \mathbb{S}_j] \wedge \mathbb{S}_1 \wedge \ldots \wedge \hat{\mathbb{S}}_i \wedge \ldots \wedge \hat{\mathbb{S}}_j \wedge \ldots \wedge \mathbb{S}_k \right) \\
&\quad = \psi_n \left(\sum_{i<j} (-1)^{i+j+1} \sum_{\lambda \in \mathbb{S}_i, \mu \in \mathbb{S}_j} (\mathbb{S}_i, \lambda)!(\mu, \mathbb{S}_j) \wedge \mathbb{S}_1 \wedge \ldots \wedge \hat{\mathbb{S}}_i \wedge \ldots \wedge \hat{\mathbb{S}}_j \wedge \ldots \wedge \mathbb{S}_k \right) \\
&\quad = \sum_{i<j} \sum_{\lambda \in \mathbb{S}_i, \mu \in \mathbb{S}_j} \\
&\qquad\quad \sum_{\pi} (-1)^{i+j+1}\, w(\pi) \Big((\mathbb{S}_i, \lambda)!(\mu, \mathbb{S}_j) \wedge \mathbb{S}_1 \wedge \ldots \wedge \hat{\mathbb{S}}_i \wedge \ldots \wedge \hat{\mathbb{S}}_j \wedge \ldots \wedge \mathbb{S}_k \Big)^{\pi} \qquad (*)
\end{aligned}
$$

where π runs over all pairings of legs of the $\mathbb{S}_i$ other than λ and μ.

On the other hand

$$
\begin{aligned}
\partial_E \psi_n (\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k) &= \partial_E \sum_{\tau} w(\tau) (\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^{\tau} \\
&= \sum_{\tau} w(\tau) \sum_{e \in (\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^{\tau}} (\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^{\tau}_{e} \qquad (1)
\end{aligned}
$$

where τ runs over all pairings of the legs of the $\mathbb{S}_i$. Two legs λ and μ form an edge e of $(S_1 \wedge \ldots \wedge S_k)^{\tau}$ if and only if they are paired by τ, i.e. $\{\lambda, \mu\} \in \tau$. If λ and μ are in the same S_i for some i, then the edge e is a loop, and $(S_1 \wedge \ldots \wedge S_k)^{\tau}_{e} = 0$. Therefore we can rewrite (1) as

$$
\begin{aligned}
&= \sum_{\tau} w(\tau) \sum_{\{\lambda, \mu\} \in \tau} (\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^{\tau}_{\lambda \cup \mu} \\
&= \sum_{\tau} \sum_{i<j} \sum_{\lambda \in \mathbb{S}_i, \mu \in \mathbb{S}_j, \{\lambda, \mu\} \in \tau} w(\tau) (\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^{\tau}_{\lambda \cup \mu}. \qquad (2)
\end{aligned}
$$

Now set $\tau' = \tau - \{\lambda, \mu\}$, i.e. τ' is the pairing on the legs other than λ and μ induced by τ. Then $w(\tau) = w(\tau')\omega(v_\lambda, v_\mu)$, and

$$
\begin{aligned}
&w(\tau)(\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)^{\tau}_{\lambda \cup \mu} \\
&\qquad = (-1)^{i+j+1} w(\tau') \omega(v_\lambda, v_\mu) (\mathbb{S}_i \wedge \mathbb{S}_j \wedge \mathbb{S}_1 \wedge \ldots \wedge \hat{\mathbb{S}}_i \wedge \ldots \wedge \hat{\mathbb{S}}_j \wedge \ldots \wedge \mathbb{S}_k)^{\tau}_{\lambda \cup \mu} \\
&\qquad = (-1)^{i+j+1} w(\tau') ((\mathbb{S}_i, \lambda)!(\mu, \mathbb{S}_j) \wedge \mathbb{S}_1 \wedge \ldots \wedge \hat{\mathbb{S}}_i \wedge \ldots \wedge \hat{\mathbb{S}}_j \wedge \ldots \wedge \mathbb{S}_k)^{\tau'}.
\end{aligned}
$$

Substituting into the sum (2) gives

$$
= \sum_{\tau} \sum_{i<j} \sum_{\lambda \in \mathbb{S}_i, \mu \in \mathbb{S}_j, \{\lambda, \mu\} \in \tau} (-1)^{i+j+1}\, w(\tau') ((\mathbb{S}_i, \lambda)!(\mu, \mathbb{S}_j) \wedge \mathbb{S}_1 \wedge \ldots \wedge \hat{\mathbb{S}}_i \wedge \ldots \wedge \hat{\mathbb{S}}_j \wedge \ldots \wedge \mathbb{S}_k)^{\tau'}
$$

Changing the order of summation now recovers the formula (*) for $\psi_n\partial_n(\mathbb{S}_1 \wedge \ldots \wedge \mathbb{S}_k)$. □

We now want to stabilize by letting $n \to \infty$. We need

Proposition 13 *The system of maps ψ_n gives rise to a map $\psi_\infty \colon \wedge\mathcal{LO}_\infty \to \mathcal{OG}$, i.e. the following diagram commutes:*

$$\begin{array}{ccc} \wedge\mathcal{LO}_n & \xrightarrow{\psi_n} & \mathcal{OG} \\ \downarrow & & \downarrow = \\ \wedge\mathcal{LO}_{n+1} & \xrightarrow{\psi_{n+1}} & \mathcal{OG} \end{array}$$

Proof Write $V_{n+1} = V_n \oplus V_1$, where V_1 is spanned by p_{n+1}, q_{n+1}. Commutativity of the diagram follows since the gluing instructions given by ψ_{n+1} are the same as those of ψ_n if the symplecto-spiders happen to have legs labeled only by elements of V_n. □

We are now ready to prove the main theorem, identifying the homology of the infinite-dimensional Lie algebra with $\mathcal{O}$-graph homology.

Theorem 1 *The map $\psi_\infty : \wedge\mathcal{LO}_\infty \to \mathcal{OG}$ induces a Hopf algebra isomorphism on homology $H_*(\mathcal{LO}_\infty;\mathbb{R}) \cong H_*(\mathcal{OG})$.*

Proof We have

$$H_*(\wedge\mathcal{LO}_\infty;\mathbb{R}) = \lim_{n\to\infty} H_*(\wedge\mathcal{LO}_n;\mathbb{R}) = \bigoplus_{k,m} \lim_{n\to\infty} H_{k,m}(\wedge\mathcal{LO}_n;\mathbb{R}).$$

By Proposition 8, $H_{k,m}(\wedge\mathcal{LO}_n;\mathbb{R}) \cong H_{k,m}((\wedge\mathcal{LO}_n)^{\mathfrak{sp}(2n)};\mathbb{R})$. By Corollary 4, there is an N such that $(\Lambda_{k,m})^{\mathfrak{sp}(2n)}, (\Lambda_{k+1,m+2})^{\mathfrak{sp}(2n)}$, and $(\Lambda_{k-1,m-2})^{\mathfrak{sp}(2n)}$ are isomorphic via ψ_n to

$$\mathcal{OG}_{k,m}, \mathcal{OG}_{k+1,m-2}, \mathcal{OG}_{k-1,m+2}$$

respectively for $n > N$. The fact that these isomorphisms respect the boundary map (Proposition 12) implies that for $n > N$ $(\psi_n)_* \colon H_{k,m}(\wedge\mathcal{LO}_n;\mathbb{R}) \to H_{k,m}(\mathcal{OG})$ is an isomorphism. Since $H_{k,m}(\mathcal{OG})$ is independent of n, we have

$$\bigoplus_{k,m} \lim_{n\to\infty} H_{k,m}(\wedge\mathcal{LO}_n;\mathbb{R}) \cong \bigoplus_{k,m} H_{k,m}(\mathcal{OG}) = H_*(\mathcal{OG}),$$

showing that ψ_∞ induces an isomorphism on homology.

To see that ψ_∞ is a Hopf algebra isomorphism, it suffices to show that the map on the chain level is a Hopf algebra morphism. We first show that it is compatible with the product structures, i.e. the following diagram commutes.

$$\begin{array}{ccc} \wedge(\mathcal{LO}_\infty)\otimes\wedge(\mathcal{LO}_\infty) & \longrightarrow & \wedge\mathcal{LO}_\infty \\ \Big\downarrow{\scriptstyle \psi_\infty\otimes\psi_\infty} & & \Big\downarrow{\scriptstyle \psi_\infty} \\ \mathcal{OG}\otimes\mathcal{OG} & \longrightarrow & \mathcal{OG} \end{array}$$

Let $x\otimes y\in\wedge(\mathcal{LO}_\infty)\otimes\wedge(\mathcal{LO}_\infty)$, where x and y are each wedges of symplecto-spiders with legs labeled by $B_\infty=\cup B_n$ (i.e. the legs are labeled by the basis vectors $\{p_i,q_i\}_{i=1}^\infty$). The top map sends $x\otimes y$ to $E(x)\wedge O(y)$. The labels of $E(x)$ are disjoint from those of $O(y)$, hence gluing up the legs with ψ_∞ will act separately on $E(x)$ and $O(y)$ yielding the disjoint union $\psi_\infty(E(x))\cup\psi_\infty(O(y))$. This says exactly that the diagram commutes.

Now we show it is a coalgebra morphism on the chain level:

$$\begin{array}{ccc} \wedge\mathcal{LO}_\infty & \longrightarrow & \wedge\mathcal{LO}_\infty\otimes\wedge\mathcal{LO}_\infty \\ \Big\downarrow{\scriptstyle \psi_\infty} & & \Big\downarrow{\scriptstyle \psi_\infty\otimes\psi_\infty} \\ \mathcal{OG} & \longrightarrow & \mathcal{OG}\otimes\mathcal{OG} \end{array}$$

Consider a wedge of B_∞-labeled symplecto-spiders $\mathbb{S}_1\wedge\ldots\wedge\mathbb{S}_k$. Both routes around the rectangle involve partitioning the spiders and gluing up the legs. Let π be a gluing, i.e. a pairing of the legs of the symplecto-spiders. Let $P=(I,J)$ be a partition of the spiders into two sets. In order for P to give a summand of $\Delta(\mathbb{S}_1\wedge\ldots\wedge\mathbb{S}_k)^\pi$, π should only pair legs of symplecto-spiders in the same part of the partition. In this case we say that π and P are *consistent*. On the other hand, a typical term of $\Delta(\mathbb{S}_1\wedge\ldots\wedge\mathbb{S}_k)$ is $\epsilon(I,J)\mathbb{S}_I\otimes\mathbb{S}_J$. In order for π to give a term of $(\psi\otimes\psi)(\epsilon(I,J)\mathbb{S}_I\otimes\mathbb{S}_J)$, π and P should be consistent. Therefore, going both ways around the rectangle, the summands are parameterized by consistent pairs (P,π).

It remains only to check the coefficients. Given a consistent pair (P,π), we can subdivide π into two pairings π_I and π_J, which pair the legs in the I and J symplecto-spiders respectively. Now the (P,π) term in $\psi_\infty\Delta(\mathbb{S}_1\wedge\ldots\wedge\mathbb{S}_k)$ is $\epsilon(I,J)(\mathbb{S}_I\otimes\mathbb{S}_J)^\pi=\epsilon(I,J)\mathbb{S}_I^{\pi_I}\otimes\mathbb{S}_J^{\pi_J}$. The (P,π) term in $\Delta\psi_\infty(\mathbb{S}_1\wedge\ldots\wedge\mathbb{S}_k)=\epsilon(I,J)\Delta\psi_\infty(\mathbb{S}_I\wedge\mathbb{S}_J)$ is the P term in $\epsilon(I,J)(\mathbb{S}_I\wedge\mathbb{S}_J)^\pi=\epsilon(I,J)\mathbb{S}_I^{\pi_I}\wedge\mathbb{S}_J^{\pi_J}$, which is just $\epsilon(I,J)\mathbb{S}_I^{\pi_I}\otimes\mathbb{S}_J^{\pi_J}$.

We illustrate the commutativity of the coalgebra diagram in Figure 14. In the upper right is a possible gluing, π, of the legs, with numbering of the vertices

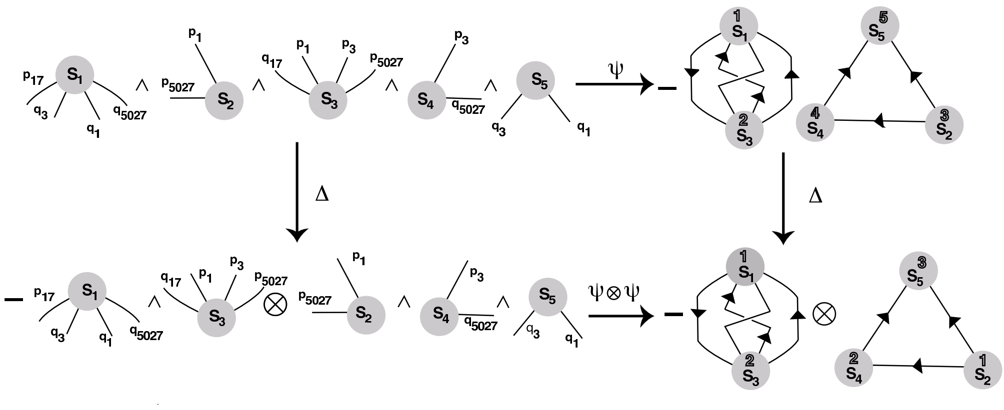

Figure 14: Commutativity of the coalgebra diagram

depicted by the outlined numbers. We have switched the numbers 2 and 3, resulting in a minus sign. In the lower left is a possible partition, P, of the symplecto-spiders, which is consistent with the gluing. In the lower right, we have the result of applying π to the lower left, which is the same as applying the partition P to the upper right. $\square$

Recall that the primitive elements of $\mathcal{OG}$ are the connected $\mathcal{O}$-graphs, spanning the subcomplex $P\mathcal{OG}$. If we denote the primitive elements of $H_*(\mathcal{LO}_\infty;\mathbb{R})$ by $Prim\,(H_*(\mathcal{LO}_\infty;\mathbb{R}))$, we have

Corollary 5 $Prim\,(H_*(\mathcal{LO}_\infty;\mathbb{R})) \cong H_*(P\mathcal{OG})$.

2.5.6 The primitive irreducible subcomplex

In this section we consider the subcomplex $P\overline{\mathcal{OG}}$ of $P\mathcal{OG}$ spanned by $\mathcal{O}$-graphs with no vertices colored by $1_\mathcal{O}$. This subcomplex is important for calculations, as well as being neccesary to draw the connections to $Out(F_r)$ and moduli space later on. In this section we prove

Proposition 14 *The homology of $P\mathcal{OG}$ splits as a direct sum of vector spaces*

$$H_*(P\mathcal{OG}) \cong H_*(\mathfrak{sp}(\infty)) \oplus H_*(P\overline{\mathcal{OG}})$$

We first consider the subcomplex $\mathcal{B}$ of $P\mathcal{OG}$ spanned by graphs which *only* have vertices colored by $1_\mathcal{O}$ (and is hence spanned by k-gons). Then $P\mathcal{OG} \cong$

$\mathcal{B} \oplus \mathcal{POG}/\mathcal{B}$, and Proposition 14 will follow by showing that $H_*(\mathcal{POG}/\mathcal{B}) \cong H_*(P\overline{\mathcal{OG}})$ (Lemma 4) and that $H_*(\mathcal{B}) \cong H_*(\mathfrak{sp}(\infty))$ (Lemma 5).

The proof of Lemma 4 depends on understanding what happens when trivial bivalent vertices (i.e. vertices colored with $1_{\mathcal{O}}$) are added to an edge of an $\mathcal{O}$ graph. For $k \geq 0$, let E_k be the oriented linear graph with k internal bivalent vertices. Define a chain complex $\mathcal{E}_*$ with one generator in each dimension, corresponding to E_k, and boundary operator ∂_E which sums over all edge collapses.

Lemma 3 *$H_0(\mathcal{E}_*) = \mathbb{R}$, and $H_i(\mathcal{E}_i) = 0$ for $i > 0$.*

Proof Collapsing any edge of E_k results in a copy of E_{k-1}. If we orient E_k by numbering its edges (see Proposition 2), then collapsing the jth edge gives $(-1)^j E_{k-1}$. Thus the boundary operator $\partial\colon \mathcal{E}_j \to \mathcal{E}_{j-1}$ is 0 for j odd and an isomorphism for j even, giving the result. □

Lemma 4 *$H_*(\mathcal{POG}/\mathcal{B}) \cong H_*(P\overline{\mathcal{OG}})$.*

Proof The quotient complex $\mathcal{T} := \mathcal{POG}/\mathcal{B}$ is spanned by $\mathcal{O}$-graphs with at least one vertex which is not colored by $1_{\mathcal{O}}$. We filter $\mathcal{T}$ by the number of vertices not colored by $1_{\mathcal{O}}$. In the associated spectral sequence we have

$$E^0_{p,q} = \mathcal{F}_p\mathcal{T}_{p+q}/\mathcal{F}_{p-1}\mathcal{T}_{p+q}.$$

That is $E^0_{p,q}$ is spanned by graphs with q vertices colored by $1_{\mathcal{O}}$ and p other vertices. Hence along the p axis of the spectral sequence we have $E^0_{p,0} = P\overline{\mathcal{OG}}_p$. The differential $d_0 : E^0_{p,q} \to E^0_{p,q-1}$ is induced by the boundary operator, and is therefore given by the map which only contracts edges incident to $1_{\mathcal{O}}$-colored vertices. Passing to the E^1 term, we claim that only the $E_{p,0}$ terms persist, and that the maps $d_1 : E^1_{p,0} \to E^1_{p-1,0}$ are just the graph boundary operator, which will conclude the proof.

We break up each vertical complex

$$\cdots E^0_{p,q} \to E^0_{p,q-1} \to E^0_{p,q-2} \to \cdots E^0_{p,0}$$

into a direct sum of chain complexes E^X_p, where the sum is over isomorphism classes of graphs with p vertices. An $\mathcal{O}$-graph is in E^X_p if removing the (bivalent) $1_{\mathcal{O}}$-colored vertices from the underlying graph results in X. Let Ξ be the graph X with all vertices and edges distinguished, so that $Aut(\Xi) = \{id\}$. Then $E^X_q = E^{\Xi}_q/Aut(X)$. Since we are in characteristic zero, we can ignore the $Aut(X)$ action, i.e. $H_q(E_*^X;\mathbb{R}) \cong H_q(E_*^\Xi;\mathbb{R})$. Thus we must show $H_q(E_*^\Xi)$ vanishes for $q > 0$.

Since the boundary operator only affects vertices colored by $1_\mathcal{O}$, E_*^Ξ breaks up as a direct sum of chain complexes $E_*^\Xi(S_1, S_2, \ldots, S_p)$, where the S_i are choices of $\mathcal{O}$-spider for each vertex of Ξ, ranging over a basis for the quotient of vector spaces $\mathcal{OS}/ < 1_\mathcal{O} >$. Each of these chain complexes is isomorphic to $\bigotimes_{e\in E(\Xi)} \mathcal{E}_*$, where $\mathcal{E}_*$ is the chain complex of Lemma 3. By this Lemma and the Künneth formula the homology of $\bigotimes_{e\in E(\Xi)} \mathcal{E}_*$ is zero except in degree 0. □

Lemma 5 $H_*(\mathcal{B}) = Prim(H_*(\mathfrak{sp}(2\infty)))$

Proof Direct computation shows that $H_k(\mathcal{B})$ is nonzero iff $k \equiv 3 \mod 4$, in which case it is one dimensional. This is well known to be the primitive homology of $\mathfrak{sp}(2\infty)$, but in fact this follows if we consider the operad which is spanned by the identity in degree 1 and is zero in higher degrees. Then $\mathcal{LO}_n$ consists of symplecto-spiders which have 2 legs and whose internal vertex is labeled by $1_\mathcal{O}$. This is isomorphic to $\mathfrak{sp}(2n)$ by Proposition 6. Thus Theorem 1 implies that the primitive homology of $\mathcal{LO}_\infty = \mathfrak{sp}(2\infty)$ is given by connected $1_\mathcal{O}$-colored graphs. This is precisely the complex $\mathcal{B}$. □

2.5.7 Brackets, cobrackets, compatibility

As we remarked earlier, the map ϕ_n does not establish an isomorphism of homologies. It is not even a chain map. However, $\mathcal{OG}$ has a different boundary operator for which ϕ_n *is* a chain map. The difference between these two boundary operators was studied in [4] in the commutative case, and in [3] for general cyclic operads.

Definition 8 Let x and y be half-edges of a graph X, and let π_{xy} denote the pairing of the half-edges $H(X)$ which differs from the standard pairing only at $x, y, \bar{x}, \bar{y}$, where it pairs $\{x, y\}$ and $\{\bar{x}, \bar{y}\}$. If $\mathbf{X}$ is an $\mathcal{O}$-graph with underlying graph X, the *H-boundary* of $\mathbf{X}$ is

$$\partial_H(\mathbf{X}) = \sum_{x,y,x\neq\bar{y}} (\mathbf{X}^{\pi_{xy}})_{x\cup y}.$$

The following proposition shows that the Lie boundary map ∂_n corresponds in fact to $2n\partial_E + \partial_H$.

Proposition 15 *The following diagram commutes.*

$$\begin{tikzcd} \mathcal{OG} \arrow[r, "\phi_n"] \arrow[d, "2n\partial_E+\partial_H"'] & \wedge\mathcal{LO}_n \arrow[d, "\partial_n"] \\ \mathcal{OG} \arrow[r, "\phi_n"] & \wedge\mathcal{LO}_n \end{tikzcd}$$

The notation ∂_H is justified by the following Lemma, which shows that ∂_H is a boundary operator in its own right.

Lemma 6 *a)* $\partial_H^2 = 0$
b) $\partial_H\partial_E + \partial_E\partial_H = 0$

Proof ϕ_n is injective for large n, and $\partial_n^2 = 0$, implying $(2n\partial_E + \partial_H)^2 = (2n)^2\partial_E^2 + 2n(\partial_H\partial_E + \partial_E\partial_H) + \partial_H^2 = 0$ for infinitely many values of n. Hence each coefficient of this polynomial in n must be zero. □

With respect to the disjoint union product on $\mathcal{OG}$, ∂_E is a derivation, but ∂_H is not. The deviation from being a derivation defines a bracket:

Definition 9 For $\mathbf{X}, \mathbf{Y} \in \mathcal{OG}$, the $\mathcal{O}$*-bracket* $[\mathbf{X}, \mathbf{Y}]$ is given by

$$[\mathbf{X}, \mathbf{Y}] = \partial_H(\mathbf{X} \cdot \mathbf{Y}) - \partial_H(\mathbf{X}) \cdot \mathbf{Y} + (-1)^{|\mathbf{X}|}\mathbf{X} \cdot \partial_H(\mathbf{Y}) = \sum_{x \in \mathbf{X}, y \in \mathbf{Y}} (\mathbf{X} \cdot \mathbf{Y})_{xy},$$

where $|\mathbf{X}|$ is the number of vertices of $\mathbf{X}$.

This bracket is a straightforward generalization of the bracket we defined in [4] for the commutative case, and coincides with the operation ϕ_2 of [3]. In [3], the following proposition is proven.

Proposition 16 *The $\mathcal{O}$-bracket is graded symmetric and satisfies the graded Jacobi identities.*

To obtain a graded *anti*-symmetric Lie bracket, we simply shift the grading (see [4]).

The proof of the following proposition can be found in [3].

Proposition 17 *The $\mathcal{O}$-bracket vanishes canonically on $\mathcal{O}$-graph homology.*

Corollary 6 *$\mathcal{O}$-graph homology $H_*(\mathcal{OG})$, equipped with the differential ∂_H, is a differential graded algebra (DGA).*

Proof By Lemma 6, ∂_H descends to the homology level. By Proposition 17, it is a derivation on the homology level. □

The boundary operator ∂_E is also a co-derivation, but the boundary operator ∂_H is not. Hence we can define

$$\theta(\mathbf{X}) = \Delta\partial_H(\mathbf{X}) - \partial_H\Delta(\mathbf{X}),$$

where ∂_H is extended to $\mathcal{OG}\otimes\mathcal{OG}$ as a graded derivation. The map θ was considered in the commutative case in [4], and coincides with the map θ_2 of [3] when the graph is connected.

Let us describe θ in more detail. A pair x, y of half-edges of a graph X is called *separating* if x and $\bar{x}$ (equivalently y and $\bar{y}$) lie in different components of $X^{\pi_{x,y}}$. (Recall π_{xy} is the pairing which pairs x with y and $\bar{x}$ with $\bar{y}$ and is the same as the standard pairing everywhere else.) Given an $\mathcal{O}$-graph $\mathbf{X}$ based on X, and a separating pair x, y of half-edges in X, write $\mathbf{X}^{\pi_{xy}}$ as a product of connected $\mathcal{O}$-graphs:

$$\mathbf{X}^{\pi_{xy}} = \mathbf{X}'\cdot\mathbf{X}''\cdot\mathbf{X}_1\cdot\mathbf{X}_2\ldots\cdot\mathbf{X}_m,$$

with $x, y\in\mathbf{X}'$ and $\bar{x},\bar{y}\in\mathbf{X}''$. Suppose (I, J) is an ordered partition of $[m]$. Define the sign $\epsilon_1(I, J)$ to satisfy the equation

$$\mathbf{X}'\cdot\mathbf{X}''\cdot\mathbf{X}_1\cdot\mathbf{X}_2\ldots\cdot\mathbf{X}_m = \epsilon_1(I,J)\mathbf{X}_I\cdot\mathbf{X}'\cdot\mathbf{X}_J\cdot\mathbf{X}''$$

and the sign $\epsilon_2(I, J)$ to satisfy the equation

$$\mathbf{X}'\cdot\mathbf{X}''\cdot\mathbf{X}_1\cdot\mathbf{X}_2\ldots\cdot\mathbf{X}_m = \epsilon_2(I,J)\mathbf{X}_I\cdot\mathbf{X}''\cdot\mathbf{X}_J\cdot\mathbf{X}'.$$

Now we define $\theta(\mathbf{X})$

$$= \sum_{x,y\text{ separating}}\sum_{I,J}\epsilon_1(I,J)(\mathbf{X}_I\cdot\mathbf{X}'_{x\cup y})\otimes(\mathbf{X}_J\cdot\mathbf{X}'')+\epsilon_2(I,J)(\mathbf{X}_I\cdot\mathbf{X}'')\otimes(\mathbf{X}_J\cdot\mathbf{X}'_{x\cup y}).$$

It is straightforward to verify that this agrees with the previous definition of θ.

The following proposition is proven in [3].

Proposition 18 *The map $\theta\colon\mathcal{OG}\to\mathcal{OG}\otimes\mathcal{OG}$ restricts to a map $P\mathcal{OG}\to P\mathcal{OG}\otimes P\mathcal{OG}$, and is a (graded) cobracket on $P\mathcal{OG}$.*

Conjecturally, θ is a cobracket on the entire complex. In [4], we proved this for the commutative case, and the proof extends to the associative case. However, it does not extend to the Lie case.

The following was noted without proof in [3] in the connected case.

Proposition 19 *The map θ vanishes canonically on the homology level.*

Proof Define a map $T\colon \mathcal{OG} \to \mathcal{OG} \otimes \mathcal{OG}$ as

$$T(\mathbf{X}) = \sum_{x,y \text{ separating}} \sum_{I,J} \epsilon_1(I,J)(\mathbf{X}_I \cdot \mathbf{X}') \otimes (\mathbf{X}_J \cdot \mathbf{X}''),$$

where $\mathbf{X}$ is written as a product of connected graphs as above. For notational ease, write $T(\mathbf{X}) = \sum_{x,y \text{ separating}} T_{xy}(\mathbf{X})$. We claim that $\theta = \partial_E T - T\partial_E$, i.e. $\partial_E T = T\partial_E + \theta$. To see this, consider the term $\partial_E T_{xy}(\mathbf{X})$ of $\partial_E T(\mathbf{X})$, which contracts all edges of $T_{xy}(\mathbf{X})$, one at a time. If the edge e is not $x \cup y$ or $\bar{x} \cup \bar{y}$, then we have $(T_{xy}(\mathbf{X}))_e = T_{xy}(\mathbf{X}_e)$. In this way we pick up all terms of $T\partial_E(\mathbf{X})$. If the edge is $x \cup y$ or $\bar{x} \cup \bar{y}$, then we pick up the two terms in $\theta(\mathbf{X})$.

Now suppose that $\mathbf{X}$ is a cycle. Then the above claim implies that $\theta(\mathbf{X}) = \partial_E T(\mathbf{X})$, and is hence trivial on the homology level. □

Corollary 7 *$\mathcal{O}$-graph homology $H_*(\mathcal{OG})$ is a differential graded Hopf algebra, with differential ∂_H.*

Corollary 8 *The map ∂_H is a graded differential on $H_*(P\mathcal{OG})$ and $H_*(P\overline{\mathcal{OG}})$.*

Proof Since ∂_H is a derivation on the homology level, it takes primitives to primitives, and so induces a differential on $PrimH_*(\mathcal{OG}) = H_*(P\mathcal{OG})$. This is *not* true on the chain level, as there exist connected graphs with disconnected terms of ∂_H.

Also any graph with a vertex colored by $1_{\mathcal{O}}$ will get mapped by ∂_H to a sum of graphs each with a vertex colored by $1_{\mathcal{O}}$ implying that ∂_H is a differential on $H_*(P\overline{\mathcal{OG}})$. If the pair of half-edges being contracted in ∂_H is not adjacent to a $1_{\mathcal{O}}$-colored vertex, then the $1_{\mathcal{O}}$ vertex survives as claimed. Otherwise, consider a $1_{\mathcal{O}}$ colored vertex. It has two adjacent half-edges h_1, h_2. Given another half-edge k, we claim that the h_1, k term cancels with the $h_2, \bar{k}$ term. In other words all the summands coming from contracting edges adjacent to $1_{\mathcal{O}}$ colored vertices cancel! Clearly these two terms give isomorphic graphs, and checking the orientation one sees that they have opposite sign. □

The following proposition is proven in [4] for the commutative case, and as remarked in [3], the proof carries over to the general cyclic operad case. It explains our interest in graphs with no separating edges.

Proposition 20 *Bracket and cobracket form a compatible Lie bi-algebra structure on the subcomplex of connected, 1-particle irreducible graphs (graphs with no separating edges).*

3 The Lie case

3.1 The forested graph complex and the Lie graph complex

We now specialize to the Lie operad. Recall that $\mathcal{O}[m]$ is the vector space spanned by all rooted vertex-oriented binary trees with m numbered leaves, modulo the subspace spanned by all anti-symmetry and IHX relators. A *vertex orientation* is a choice of cyclic order at each trivalent vertex of a binary tree, and is equivalent to specifying a planar embedding up to isotopy. The composition rule in the operad attaches the root of one tree to a leaf of another, eliminates the resulting bivalent vertex, and then suitably renumbers the remaining leaves.

The space of $\mathcal{O}$-spiders, $\mathcal{OS}$, is spanned by vertex-oriented binary trees with m numbered leaves, but with no particular leaf designated as the root. Mating is defined by gluing two such trees together at a leaf, eliminating the resulting bivalent vertex, then renumbering the leaves suitably.

Thus a basic $\mathcal{O}$-graph for the Lie operad is a fairly complicated object: an oriented graph with a vertex-oriented trivalent tree attached at each vertex, modulo AS and IHX relations. However, the picture can be considerably simplified since the orientations of the trees at the vertices cancel to a large degree with the orientation of the graph itself. We now describe this simplification.

Let X be a finite graph. A *forest* in X is a sub-graph which contains no cycles. The connected components of a forest are trees, where we allow a tree to consist of a single vertex. An *orientation* of a forest is given by an ordering of its edges (if any); interchanging the order of any two edges reverses the orientation of the forest.

Definition 10 A *forested graph* is a pair (G, Φ), where G is a finite connected trivalent graph and Φ is an oriented forest which contains all vertices of G.

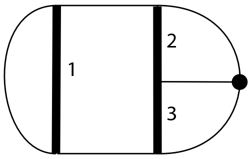

Figure 15: Forested graph

An example of a forested graph with three trees is shown in Figure 15.

We denote by $\widehat{f\mathcal{G}_k}$ the vector space spanned by all forested graphs whose forest contains exactly k trees, modulo the relations $(G,\Phi)=-(G,-\Phi)$.

Let (G,Φ) be a forested graph, and let e be an edge of Φ. Collapsing e produces a new pair (G_e,Φ_e), where G_e has one 4-valent vertex. There are exactly two other forested graphs (which may be isomorphic!), $(G^{'},\Phi')$ and $(G^{''},\Phi'')$, and edges $e'\in\Phi', e''\in\Phi''$ which produce the same pair (G_e,Φ_e) (see Figure 16).

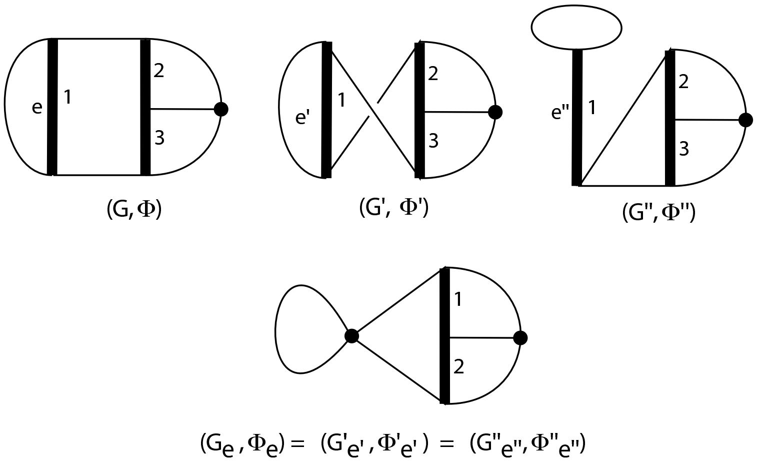

Figure 16: Blow-ups of (G_e,Φ_e)

The vector

$$(G,\Phi)+(G^{'},\Phi')+(G^{''},\Phi'')$$

is called the *basic IHX relator* associated to (G,Φ,e).

Let IHX_k be the subspace of $\widehat{f\mathcal{G}_k}$ spanned by all basic IHX relators, and define $f\mathcal{G}_k$ to be the quotient space $\widehat{f\mathcal{G}_k}/IHX_k$.

Proposition 21 *$f\mathcal{G}_k$ is naturally isomorphic as a vector space to $P\overline{\mathcal{O}\mathcal{G}}_k$, where $\mathcal{O}$ is the Lie operad.*

Proof We define a map $\widehat{f\mathcal{G}_k} \to P\overline{\mathcal{O}\mathcal{G}}_k$ as follows. Starting with a forested graph (G, Φ), collapse each component of Φ to a point to produce a graph X. For each vertex v of X, let T_v be an ϵ-neighborhood of the preimage of v in G, so that T_v is a binary tree whose leaves are identified with the half-edges $H(v)$ adjacent to v (see Figure 17). To specify an $\mathcal{O}$-graph $\mathbf{X} = (X, \{T_v\})$, it remains only to determine the orientations on X and on the trees T_v. In the

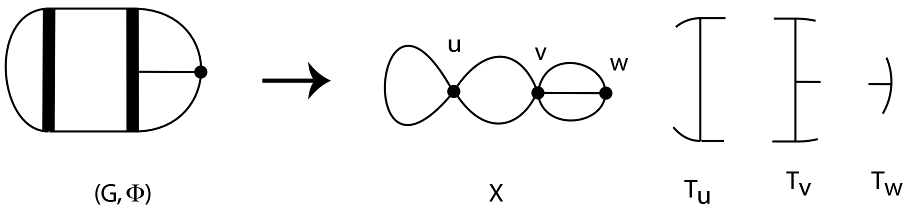

Figure 17: Correspondence between forested graphs and Lie graphs

language of section 2.3.1, this information can be described as a unit vector in

$$\det(\mathbb{R}V(X)) \otimes \bigotimes_{e\in E(X)} \det \mathbb{R}H(e) \otimes \bigotimes_{v\in V(X)} (\otimes_{u\in V(T_v)} \det \mathbb{R}H(u))$$

That is, it is specified by a numbering of the vertices of X, an orientation on each edge of X, and a vertex-orientation on each tree T_v.

Applying Propositions 1 and 2 to the first two tensor factors and applying Corollary 2 to the third tensor factor, this is canonically isomorphic to

$$\det(\bigoplus_{|H(v)| even} \mathbb{R}v) \otimes \bigotimes_{v\in V(X)} \det \mathbb{R}H(v) \otimes \bigotimes_{v\in V(X)} \det(\mathbb{R}E(T_v)) \qquad (1)$$

In other words, an orientation can be specified by an ordering of the even valence vertices of X, an ordering of the half-edges incident to each vertex of X, and an ordering of the edges of each tree.

The forest Φ is the union of the *internal* edges of all of the T_v. Denote the internal edges of T_v by Φ_v, so that $\Phi = \cup_v \Phi_v$. Since $H(v)$ is identified with the leaves of T_v, the Partition lemma implies $\det(\mathbb{R}E(T_v)) \cong \det(\mathbb{R}H(v)) \otimes \det(\mathbb{R}E(\Phi_v))$. (Actually, the Partition lemma tells you to also order the set of odd subsets of $E(T_v)$, but since the number of leaves and the number of interior edges have opposite parity, this would be ordering at most one object and that is not extra information.) Hence, $\det(\mathbb{R}H(v)) \otimes \det(\mathbb{R}E(T_v)) \cong \det(\mathbb{R}E(\Phi_v))$, so that expression (1) becomes

$$\begin{aligned} &\det\Big(\bigoplus_{|H(v)|\,even} \mathbb{R}v\Big) \otimes \bigotimes_{v\in V(X)} \det(\mathbb{R}E(\Phi_v)) \\ &= \det\Big(\bigoplus_{|E(\Phi_v)|\,odd} \mathbb{R}v\Big) \otimes \bigotimes_{v\in V(X)} \det(\mathbb{R}E(\Phi_v)) \\ &\cong \det\Big(\bigoplus_{e\in\Phi} \mathbb{R}e\Big), \end{aligned}$$

where the last isomorphism is again given by the Partition lemma. Thus the orientation data on $\mathbf{X} = (X, \{T_v\})$ is equivalent to an ordering of the edges of Φ.

This defines the map $\widehat{f\mathcal{G}_k} \to P\overline{\mathcal{O}\mathcal{G}}_k$. We claim that a basic IHX relator $(G,\Phi) + (G',\Phi') + (G'',\Phi'')$ maps to $(X,\{T_v\}) - (X,\{T'_v\}) + (X,\{T''_v\}) = (X,\{T_v - T'_v + T''_v\}) = 0$, so the map factors through $f\mathcal{G}_k$. This is obvious except for the signs of the terms. To check these, we need to carefully consider the equivalence of orientations described above. The essential step is the translation from vertex orienting a tree to ordering the tree's edges. We illustrate the argument by doing the case of a tree with 5 edges, and leave the general case to the reader. Suppose we have

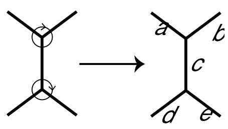

where $\{a,b,c,d,e\} = \{1,2,3,4,5\}$ (not neccesarily in that order). Then, depending on where we place the tree in the plane, we have

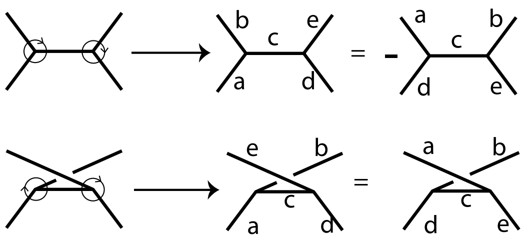

Thus the $I - H + X$ relator of vertex oriented trees gets mapped to an $I + H + X$ relator of trees with ordered edges.

To see that the map $f\mathcal{G}_k \to P\overline{\mathcal{O}\mathcal{G}}_k$ is an isomorphism, define an inverse $P\overline{\mathcal{O}\mathcal{G}}_k \to \widehat{f\mathcal{G}_k} \to f\mathcal{G}_k$ as follows: If $(X,\{T_v\})$ is a basic $\mathcal{O}$-graph, produce a trivalent graph G from this data by replacing each vertex v of X by T_v, using the identification of the leaves of T_v with $H(v)$ as gluing instructions (see Figure 17). The union of the interior edges of the T_v forms a forest Φ in G, and we have seen that the orientation data on $(X, \{T_v\})$ determines an orientation of Φ. $\square$

We have identified the spaces $f\mathcal{G}_k$ with the chain groups $P\overline{\mathcal{OG}}_k$. Under this identification, it can be checked that the boundary map ∂_E is induced by the map on $\widehat{f\mathcal{G}}_k$ given by

$$\partial_E(G,\Phi) = \sum (G, \Phi \cup e),$$

where the sum is over all edges e of $G-\Phi$ such that $\Phi \cup e$ is still a forest. (Note that this only happens if the initial and terminal vertices of e lie in different components of Φ, so that $\Phi \cup e$ has $k-1$ components.) The orientation of $\Phi \cup e$ is determined by ordering the edges of Φ with the labels $1, \ldots, k$ consistent with its orientation, and then labeling the new edge e with $k+1$.

Remark 3 As defined above, ∂_E is a boundary operator on the chain complex $\widehat{f\mathcal{G}}_*$ of forested graphs. There is also a coboundary map $\delta\colon \widehat{f\mathcal{G}}_k \to \widehat{f\mathcal{G}}_{k+1}$, defined by the formula

$$\delta(G,\Phi) = \sum_{e_i \in \Phi} (-1)^i (G, \Phi - e_i),$$

where the orientation on $\Phi - e_i$ is induced from the ordering of the edges of Φ. However, this coboundary does not preserve the IHX-subspaces, so does not induce a coboundary operator on $f\mathcal{G}_*$. Of course, we can get a coboundary operator by dualizing: $\delta = \partial_E^*\colon f\mathcal{G}_k^* \to f\mathcal{G}_{k+1}^*$, but there is no natural basis of $f\mathcal{G}_k$, and hence no natural identification of $f\mathcal{G}_k$ with $f\mathcal{G}_k^*$.

3.1.1 Bracket and cobracket in the forested graph complex

In this section we describe how the bracket and cobracket of section 2.5.7 translate to forested graphs.

Let (G_1, Φ_1) and (G_2, Φ_2) be forested graphs, let x be a half-edge in $G_1 - \Phi_1$ and let y be a half-edge in $G_2 - \Phi_2$. We form a new trivalent graph $(G_1 \cdot G_2)^{\pi_{xy}}$, and forest it by the image of $\Phi_1 \cup \Phi_2$ plus the new edge xy. The orientation is given by shifting the numbering the edges of Φ_2 to lie after Φ_1, and letting

$x \cup y$ come after that. The bracket is then defined by summing over all possible pairs x and y:

$$[(G_1, \Phi_1), (G_2, \Phi_2)] = \sum_{x,y} ((G_1 \cdot G_2)^{\pi_{xy}}, \Phi_1 \cup \Phi_2 \cup x \cup y)$$

The cobracket is defined as follows. Let (G, Φ) be a forested graph, and let x, y be two half-edges not in Φ and not adjacent to the same tree of Φ. such that $G^{\pi_{xy}}$ has two connected components G_1 and G_2, where $x \cup y$ is in G_1. The pair of half-edges $\{x, y\}$ is said to be *separating*. Let $\Phi_i = \Phi \cap G_i$. Add $x \cup y$ to the forest $\Phi_1 \cup \Phi_2$, ordered after everything else. Adjust the representative of the orientation so that the numbering of $\Phi_1 \cup x \cup y$ precedes that of Φ_2. Form the symmetric product $(G_1, \Phi_1 \cup x \cup y) \odot (G_2, \Phi_2)$, where the numbering of Φ_2 is shifted down by the number of edges in $\Phi_1 \cup x \cup y$. The cobracket is given by summing over all such pairs $\{x, y\}$:

$$\theta(G, \Phi) = \sum_{\{x,y\} \text{ separating}} (G_1, \Phi_1 \cup x \cup y) \odot (G_2, \Phi_2)$$

3.2 A filtration of Outer Space, and the associated complex

The group $Out(F_r)$ of outer automorphisms of a free group of rank r acts cocompactly on a contractible simplicial complex K_r, known as the spine of Outer space. For an introduction to Outer space and its spine, see [23]. Point stabilizers of this action are finite, so that the quotient $Q_r = K_r/Out(F_r)$ has the same rational cohomology as $Out(F_r)$. In this section we define a filtration on K_r and use the spectral sequence of this filtration to prove the following theorem:

Theorem 2 *Let $f\mathcal{G}^{(r)}$ denote the subcomplex of $f\mathcal{G}$ spanned by (connected) forested graphs of rank r. Then $H_k(f\mathcal{G}^{(r)}) \cong H^{2r-2-k}(Out(F_r))$.*

Recall that vertices of K_r are pairs (g, X), where X is a connected graph with all vertices at least trivalent, and g is a homotopy equivalence from a standard rose (wedge of circles) R_r to X; thus g gives an identification of $\pi_1(X)$ with the group $F_r = \pi_1(R_r)$. The k-simplices of K_r, for $k \geq 1$, can be identified with chains

$$\emptyset = \Phi_0 \subset \Phi_1 \subset \ldots \subset \Phi_k$$

of forests in a marked graph (g, X); the ith vertex of the k-simplex is obtained from (g, X) by collapsing each edge of Φ_i to a point. $Out(F_r)$ acts by changing the marking g, and the stabilizer of (g, X) is isomorphic to the group of automorphisms of the graph X.

We define a filtration on K_r as follows:

- $\mathcal{F}_0K_r$ consists of all vertices (g,X) with X a (connected) trivalent graph. Note that a trivalent graph X with fundamental group isomorphic to F_r has $2r-2$ vertices.
- $\mathcal{F}_iK_r$ is the sucomplex of K_r spanned by $\mathcal{F}_{i-1}K_r$ together with all vertices (g,X) such that X has $2r-2-i$ vertices.

We have $\mathcal{F}_0K_r \subset \mathcal{F}_1K_r \subset \ldots \subset \mathcal{F}_{2r-3}K_r = K_r$, where the ith subcomplex $\mathcal{F}_iK_r$ has dimension i. Since the action of $Out(F_r)$ simply changes the markings g, it preserves the filtration, and so induces a filtration $\mathcal{F}_0Q_r \subset \mathcal{F}_1Q_r \subset \ldots \subset \mathcal{F}_{2r-3}Q_r = Q_r$ on the quotient Q_r. The cohomology spectral sequence of this filtration has

$$E_1^{p,q} = H^{p+q}(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r)$$

and converges to the cohomology of Q_r. In particular, if we take trivial real coefficients, this spectral sequence converges to the rational cohomology of $Out(F_r)$. All cohomology groups will be assumed to have trivial rational coefficients unless otherwise specified.

Proposition 22 $H^{p+q}(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r) = 0$ *for* $q \neq 0$

Proof Marked graphs in $\mathcal{F}_pK_r$ which are not in $\mathcal{F}_{p-1}K_r$ are those where $\sum_v(|v|-3)$ is exactly p, where $|v|$ is the valence of the vertex v. The link in $\mathcal{F}_{p-1}K_r$ of such a marked graph (g,X) is the join of the blow-up complexes of the vertices of X of valence >3.

A simplicial complex is called *i-spherical* if it is i-dimensional and homotopy equivalent to a wedge of i-spheres. The blow-up complex of a vertex of valence k is $(k-4)$-spherical (see, e.g. [22], Theorem 2.4). Therefore the link of (g,X) in $\mathcal{F}_{p-1}K_r$ is $(p-1)$-spherical. Since graphs in $\mathcal{F}_pK_r - \mathcal{F}_{p-1}K_r$ are not connected by edges, the entire complex $\mathcal{F}_pK_r/\mathcal{F}_{p-1}K_r$ is p-spherical. Therefore $H^{p+q}(\mathcal{F}_pK_r, \mathcal{F}_{p-1}K_r) = 0$ for $q \neq 0$.

Since $Out(F_r)$ preserves the filtration, the homology of $\mathcal{F}_pQ_r/\mathcal{F}_{p-1}Q_r$ is the homology of $\mathcal{F}_pK_r/\mathcal{F}_{p-1}K_r$ modulo the action of $Out(F_r)$. The stabilizer of each graph in $\mathcal{F}_pK_r$ is finite, so the (reduced) homology of the quotient $\mathcal{F}_pQ_r/\mathcal{F}_{p-1}Q_r$ vanishes in dimensions other than p, as was to be shown. □

Proof of Theorem By Proposition 22, the spectral sequence of the filtration $\mathcal{F}_*Q_r$ of Q_r collapses to the cochain complex

$$0 \to H^0(\mathcal{F}_0Q_r) \to H^1(\mathcal{F}_1Q_r, \mathcal{F}_0Q_r) \to H^2(\mathcal{F}_2Q_r, \mathcal{F}_1Q_r) \to \\ \ldots \to H^{2r-3}(\mathcal{F}_{2r-3}Q_r, \mathcal{F}_{2r-4}Q_r) \to 0,$$

which computes the cohomology of $Out(F_r)$.

To identify the terms $H^p(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r)$ it is helpful to recall the description of K_r as a cubical complex (see [9]). The simplices of K_r organize themselves naturally into cubes (g, X, Φ), one for each forest Φ in a marked graph (g, X) (see Figure 18). An ordering on the edges of Φ determines an orientation of the cube.

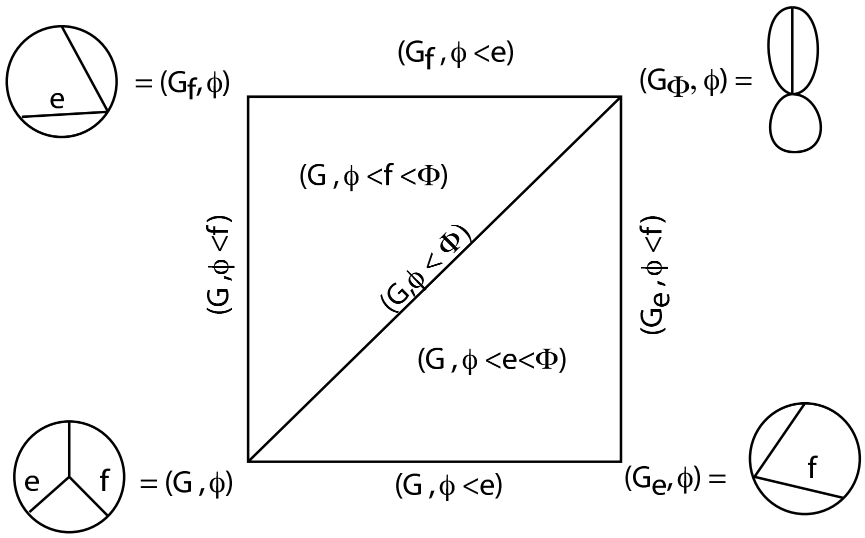

Figure 18: The cube (G, Φ) for $\Phi = \{e, f\}$

The stabilizer of a cube (g, X, Φ) consists of automorphisms of X which preserve Φ (see [20]). The quotient of a cube by its stabilizer is a cone on a rational homology sphere, or on a rational homology disk if there is an automorphism of (X, Φ) which reverses orientation.

Now $\quad H^p(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r)$

$$= C^p(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r)/im(\delta \colon C^{p-1}(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r) \to C^p(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r)).$$

Note that $\mathcal{F}_{p-1}Q_r$ is $(p-1)$-dimensional, and $\mathcal{F}_pQ_r$ has one p-dimensional cube quotient for every pair (G, Φ) such that G is trivalent and Φ has p edges. Therefore

$$C^p(\mathcal{F}_pQ_r, \mathcal{F}_{p-1}Q_r) = C^p(\mathcal{F}_pQ_r) = (\mathbb{R}\{(G, \Phi)\}/AS)^*,$$

where AS is the subspace generated by the antisymmetry relations $(G, \Phi) = -(G, -\Phi)$.

Each p-cube (g, G, Φ) in $\mathcal{F}_p K_r$ has 2^p codimension one faces of the form $(g, G, \Phi - e)$, which are in $\mathcal{F}_{p-1} K_r$. Opposite each is a codimension one face $(g, G, \Phi - e)^{op} = (g_e, G_e, \Phi_e)$. The relative cochain group $C^{p-1}(\mathcal{F}_p K_r, \mathcal{F}_{p-1} K_r)$ consists of functions on the $(p-1)$-cubes of $\mathcal{F}_p K_r$ which vanish on $\mathcal{F}_{p-1} K_r$, i.e. $C^{p-1}(\mathcal{F}_p K_r, \mathcal{F}_{p-1} K_r)$ has as a basis the characteristic functions of the $(p-1)$-cubes (g_e, G_e, Φ_e) of the second type. Each such $(p-1)$-cube is a face of exactly three p-cubes, (g, G, Φ), (g', G', Φ') and (g'', G'', Φ''), with $G_e = G'_{e'} = G''_{e''}$, so that $\delta(g_e, G_e, \Phi_e) = (g, G, \Phi) + (g', G', \Phi') + (g'', G'', \Phi'')$. Passing to the quotient modulo $Out(F_r)$, we get the IHX relator

$$\delta(G_e, \Phi_e) = (G, \Phi) + (G', \Phi') + (G'', \Phi''),$$

showing that in the quotient, $im(\delta)$ is spanned by the (characteristic functions of) the IHX relators $IHX(G, \Phi, e)$.

Thus $H^p(\mathcal{F}_p Q_r, \mathcal{F}_{p-1} Q_r)$

$$= C^p(\mathcal{F}_p Q_r, \mathcal{F}_{p-1} Q_r)/im(\delta) = (\mathbb{R}\{(G, \Phi)\}/AS)^*/IHX^* = f\mathcal{G}_{2r-2-p}$$

The last isomorphism follows since $\mathbb{R}\{(G, \Phi)\}/AS$ and IHX are canonically isomorphic to their duals, even though the quotient is not. This identifies the cochain groups $E_1^{p,0}$ above with the chain groups of the forested graph complex. The d^1 map in the spectral sequence is induced by the coboundary map

$$H^{p-1}(\mathcal{F}_{p-1} Q_r, \mathcal{F}_{p-2} Q_r) \to H^p(\mathcal{F}_p Q_r, \mathcal{F}_{p-1} Q_r)$$

which maps

$$(G, \Phi) \mapsto \sum (G, \Phi \cup e),$$

where the sum is over all edges e in G such that $\Phi \cup e$ is a forest; this is the ∂_E map in the forested graph complex. □

3.3 The subcomplex of connected one-particle irreducible graphs

A graph X is called *one-particle irreducible* if removing an edge does not change the number of connected components, i.e. if X has no separating edges. We have

Proposition 23 *The subcomplex IK_r of K_r spanned by (connected) one-particle irreducible graphs is an equivariant deformation retract of K_r.*

Proof The deformation retraction acts by uniformly shrinking all separating edges to points in all graphs representing vertices of K_r. (see [6]). □

Thus we may compute the rational cohomology of $Out(F_r)$ using IK_r instead of K_r. The filtration $\mathcal{F}_*K_r$ restricts to a filtration on IK_r, and of the quotient IQ_r of IK_r by the action of $Out(F_r)$, giving a spectral sequence

$$E_1^{p,q} = H^{p+q}(\mathcal{F}_p IQ_r, \mathcal{F}_{p-1} IQ_r)$$

The fact that this sequence degenerates to a chain complex computing the homology of $Out(F_r)$ is guaranteed by the following proposition

Proposition 24 $H^{p+q}(\mathcal{F}_p IQ_r, \mathcal{F}_{p-1} IQ_r) = 0$ *for* $q \neq 0$

Proof As in the proof of Proposition 22, what is needed is to prove that the link of a vertex in $\mathcal{F}_p IK_r - \mathcal{F}_{p-1} IK_r$ intersects $\mathcal{F}_{p-1} IK_r$ in a $(p-1)$-spherical subcomplex. This link is called $K_{>(g,G)}$ in [22], and Corollary 3.2 of Theorem 3.1 of that paper proves that it is $(p-1)$-spherical. □

Now let $\mathcal{O}$ be the Lie operad, and consider the complex $PI\overline{\mathcal{O}\mathcal{G}}^{(r)}$ of $P\overline{\mathcal{O}\mathcal{G}}^{(r)}$ which is the quotient of $P\overline{\mathcal{O}\mathcal{G}}^{(r)}$ by graphs with separating edges, which may be internal to vertices. Proposition 20 implies that the bracket and cobracket form a compatible bi-algebra structure on this complex. (One must check that the bracket takes graphs with internal separating edges to graphs with internal separating edges, and check a similar statement for the cobracket.) The proof of Proposition 21 restricts to give an isomorphism of the subspace $If\mathcal{G}_k$ of $f\mathcal{G}_k$ spanned by 1-particle irreducible forested graphs with $PI\overline{\mathcal{O}\mathcal{G}}_k$. The proof of Theorem 2 also restricts, and the combination gives

Theorem 3 $H_k(PI\overline{\mathcal{O}\mathcal{G}}^{(r)}) \cong H_k(If\mathcal{G}^{(r)}) \cong H^{2r-2-k}(Out(F_r))$. □

This theorem, together with Proposition 20, shows there is a chain complex that computes $H^*(Out(F_r);\mathbb{R})$ which carries a highly non-trivial Lie bialgebra structure. This bialgebra structure vanishes at the level of homology, so that $\oplus_r H^*(Out(F_r);\mathbb{R})$ is the primitive part of a differential graded Hopf algebra.

4 The associative case

4.1 The associative graph complex and the forested ribbon graph complex

We now make a few remarks concerning the associative case, i.e when the operad $\mathcal{O}$ is the *associative operad.* In this case $\mathcal{O}[m]$ is $m!$-dimensional, with a basis consisting of of rooted planar trees with one interior vertex and m numbered leaves. An equivalent description, closer to the description of the Lie operad, is that $\mathcal{O}[m]$ is spanned by rooted planar *binary* trees with m numbered leaves, modulo anti-symmetry and associativity (or "IH") relations.

$\mathcal{O}$-spiders in this case are planar spiders with one vertex and m numbered legs. Planarity of the spider can be thought of as a cyclic ordering on the legs, so that an $\mathcal{O}$-graph can be characterized as an oriented graph X together with a cyclic ordering of the edges incident to each vertex. Such objects are known as a "ribbon graphs," and have been studied in conjunction with mapping class groups of punctured surfaces (see, e.g., [19]).

For S an oriented compact surface with boundary, we denote by $\Gamma(S)$ the mapping class group of isotopy classes of homeomorphisms of S which preserve the orientation. Homeomorphisms and isotopies need not be the identity on the boundary, and in particular homeomorphisms may permute the boundary components.

4.1.1 Surface subcomplexes of $P\overline{\mathcal{O}\mathcal{G}}$

Given a (connected) ribbon graph in $P\overline{\mathcal{O}\mathcal{G}}$, one can "fatten" it to a unique compact, connected, oriented surface in such a way that the cyclic orientation at each vertex is induced from the orientation of the surface. Contracting an edge will not change the oriented surface. Thus the connected, reduced associative graph complex breaks up into a direct sum over surfaces S:

$$P\overline{\mathcal{O}\mathcal{G}} = \bigoplus_S \left(P\overline{\mathcal{O}\mathcal{G}}\right)_S .$$

One can think of a ribbon graph as an equivalence class of embeddings of a graph into a surface S which induce an isomorphism on π_1; here two embeddings are considered equivalent if they differ by a homeomorphism of the surface.

4.1.2 Forested ribbon graphs

A *forested ribbon graph*, (G,Φ) is a connected trivalent ribbon graph G and an oriented forest $\Phi \subset G$ that contains all the vertices. Let $\widehat{fr\mathcal{G}}$ denote the vector space spanned by forested ribbon graphs, modulo the anti-symmetry relation $(G,\Phi) = -(G,-\Phi)$. Let $\widehat{fr\mathcal{G}}_S$ denote the subcomplex of $\widehat{fr\mathcal{G}}$ spanned by forested graphs (G,Φ) where G thickens to the surface S.

Recall that an IHX relation comes from blowing up a 4-valent vertex in all possible ways. Similarly, in the case of a ribbon graph, we can define an "IH" relation which is the sum of blowing up a 4-valent vertex in the two possible ways consistent with the ribbon structure. Now define $fr\mathcal{G}$ to be the quotient of $\widehat{fr\mathcal{G}}$ by IH relations, and $fr\mathcal{G}_S$ to be the subcomplex which thickens to S.

Proposition 25 *$fr\mathcal{G}_S$ is naturally isomorphic to $\left(P\overline{\mathcal{OG}}\right)_S$.*

Proof We use the description of the associative operad as a quotient of binary trees by IH relations. Define a map $fr\mathcal{G}_S \to \left(P\overline{\mathcal{OG}}\right)_S$ as in the Lie case by collapsing the forest, and coloring the resulting vertices by ϵ-neighborhoods of preimages. As in the Lie case, the orientation of the forest translates to a vertex-orientation of the tree at each vertex plus an orientation in the usual sense of the underlying graph. Each tree is already canonically vertex-oriented by the ribbon structure, so that we can compare the induced vertex-orientation with the canonical one, incurring a plus or minus sign.

One must check that IH relations get sent to IH relations and that this translation is a chain map, the only issue in both cases being the sign. □

4.2 Ribbon subcomplexes of Outer Space

Fix an oriented surface S with non-empty boundary, and consider the set of all isotopy classes of embeddings of ribbon graphs $(X, \{or_v\}) \to S$, where the embedding respects the cyclic orientations at vertices, and induces an isomorphism on the fundamental group. The set of all such isotopy classes of embeddings forms a simplicial complex L_S, where an edge corresponds to collapsing a forest in X. In fact, if we forget the orientations or_v, this is just a subcomplex of the spine K_r of Outer space. To see, this, we identify $\pi_1(S)$ with the fundamental group of a rose R_r which is a deformation retract of S; then the marking g on X is a homotopy inverse to the embedding followed by the retraction.

This subcomplex of K_r is contractible. This can be seen directly from the proof that Outer space is contractible (for details, see [10]), or via the identification of this complex with a deformation retract of the Teichmüller space of the punctured surface ([19]). The mapping class group $\Gamma(S)$ of S is naturally a subgroup of $Out(F_r)$, namely, $\Gamma(S)$ is the stabilizer, in $Out(F_r)$, of the set of cyclic words represented by small loops around the punctures (see [25]). The action of $Out(F_r)$ on K_r restricts to an action of $\Gamma(S)$ on L_S, so the quotient of L_S by this action, denoted Q_S is a rational $K(\Gamma(S), 1)$. The filtration of K_r by the number of vertices in the graph restricts to a filtration $\mathcal{F}_pL_S$ of L_S which is invariant under the action of $\Gamma(S)$, so the cohomology spectral sequence of the quotient filtration converges to the cohomology of $\Gamma(S)$. We have

$$E_1^{p,q} = H^{p+q}(\mathcal{F}_pQ_S, \mathcal{F}_{p-1}Q_S)$$

As before, we need the analog of Proposition 22:

Proposition 26 $H^{p+q}(\mathcal{F}_pQ_S, \mathcal{F}_{p-1}Q_S) = 0$ *for* $q \neq 0$

Proof In this case, the link of a vertex in $\mathcal{F}_pL_S - \mathcal{F}_{p-1}L_S$ intersects $\mathcal{F}_{p-1}L_S$ in a single sphere, of dimension $(p-1)$. This is basically a consequence of the fact that the "IHX" relation is just an "IH" relation when the graph is restricted to lie on a surface, i.e. when it is necessary to preserve the cyclic orderings at vertices. □

Theorem 4 *Let* S *be a surface with fundamental group* F_r. *Then*

$$H^{2r-2-k}(\Gamma(S);\mathbb{R}) \cong H_k(fr\mathcal{G}_S) \cong H_k((P\overline{\mathcal{OG}})_S)$$

Proof By the above remarks $H_*(\Gamma(S);\mathbb{R})$ is computed by the chain complex

$$\ldots \to H^p(\mathcal{F}_pQ_S, \mathcal{F}_{p-1}Q_S) \to H^{p+1}(\mathcal{F}_{p+1}Q_S, \mathcal{F}_pQ_S) \to \ldots$$

As in the case of $Out(F_r)$, this can be described as a quotient of the vector space spanned by forested ribbon graphs

$$(\mathbb{R}\{(G,\Phi)\}/AS)^*/IH^* \cong fr\mathcal{G}_S.$$ □

The deformation retraction of K_r onto the subcomplex spanned by 1-particle irreducible graphs restricts to a deformation retraction of L_S. This sets up an isomorphism with $IP\overline{\mathcal{OG}}_S$, which is the quotient complex of $P\overline{\mathcal{OG}}_S$ by graphs with separating edges (internal or external). As in the $Out(F_r)$ case, this implies there is a chain complex that computes $H^*(\Gamma(S);\mathbb{R})$ which carries a Lie bialgebra structure. This bialgebra structure vanishes at the level of homology, so that $\oplus_S H^*(\Gamma(S);\mathbb{R})$ is the primitive part of a differential graded Hopf algebra.

5 The commutative case, revisited

5.1 Kontsevich's graph complex

When $\mathcal{O}$ is the *commutative operad*, each vector space $\mathcal{O}[m]$ is one-dimensional, with a canonical basis element consisting of a rooted tree (not planar!) with m numbered leaves. $\mathcal{OS}[m]$ has a basis consisting of $*_m$. Thus an $\mathcal{O}$-graph for the commutative operad is simply an oriented graph, with no additional structure attached to the vertices. In the graph complex associated to the commutative operad, the k-chains $\mathcal{OG}_k$ have as basis oriented graphs with k vertices, and the boundary operator ∂_E acts by collapsing edges. Note that we may assume that a commutative graph has no edges which are loops, since a graph automorphism which reverses the orientation on such an edge also reverses the orientation on the graph, making that graph zero in the chain group. This chain complex of oriented graphs was the one studied in [4]. It was noted there that the subcomplex of one-particle irreducible graphs, on which the bracket and cobracket are compatible, does not have the same homology as the full complex, unlike the associative and Lie cases. We revisit the commutative case here to give a geometric interpretation of the complex and explain the fact that the subcomplex is not quasi-isometric to the full complex. We recently discovered a proof that commutative graph homology without separating edges or separating *vertices* does compute the commutative graph homology. This is explained in [5].

5.2 The simplicial closure of Outer space

Full Outer space U_r, as opposed to the spine K_r, can be described as a union of open simplices, one for each vertex of K_r. Recall that a vertex of K_r is a marked graph (g, X) with $k+1$ edges; points of the corresponding open k-simplex are given by assigning non-zero lengths to the edges of X, subject to the constraint that the sum of the lengths must equal 1. Gluing instructions are given by adjacency relations in K_r: an open face is attached when an edge is allowed to degenerate to length 0, changing the topological type of the graph but not its fundamental group.

This union of open simplices is not a simplicial complex, since each simplex is missing certain faces. Specifically, the missing faces correspond to proper subgraphs which contain cycles; when the edges in such a subgraph are collapsed to points, the fundamental group of the graph changes and the marked graph is no longer in U_r. We can complete U_r to a simplicial complex $\widehat{U}_r$ by formally adding in these missing faces. The action of $Out(F_r)$ extends, and the union A_r of the added simplices forms an invariant subcomplex of $\widehat{U}_r$. We denote by $C_*(\mathbf{X}_r, A_r)$ the corresponding relative simplicial chain complex.

There are two actions of $Out(F_r)$ on $\mathbb{R}$, the trivial action and the action via the composition $Out(F_r) \to Out(\mathbf{Z}^r) = Gl(r, \mathbf{Z}) \to \mathbb{R}$ given by $\alpha \mapsto \bar{\alpha} \mapsto \det(\bar{\alpha})$. We denote by $\tilde{\mathbb{R}}$ the reals with the non-trivial action. We claim that the commutative graph complex can be identified with $C_*(\widehat{U}_r, A_r) \otimes_{Out(F_r)} \tilde{\mathbb{R}}$, so that Kontsevich's graph homology is isomorphic to the equivariant homology $H_*^{Out(F_r)}(\widehat{U}_r, A; \tilde{\mathbb{R}})$ (see [2]).

Kontsevich's graph complex $\mathcal{OG}$ is a direct sum of complexes $\mathcal{OG}^{(r)}$, where $\mathcal{OG}_k^{(r)}$ has as basis connected, oriented graphs with k vertices and Euler characteristic $1 - r$.

Proposition 27 *The equivariant homology $H_p^{Out(F_r)}(\widehat{U}_r, A_r; \tilde{\mathbb{R}})$ is isomorphic to graph homology $H_{p+1-r}(\mathcal{OG}^{(r)})$.*

Proof The relative chains $C_p(\widehat{U}_r, A_r) = C_p(\widehat{U}_r)/C_p(A_r)$ have one basis element for each p-simplex in $\widehat{U}_r$ which is not entirely contained in A_r. Such a p-simplex is given by a marked graph (g, X) with $p+1$ edges. The vertices of this p-simplex correspond to the edges e of X (they are obtained by collapsing all edges other than e), so that ordering the edges of X determines an orientation of the corresponding p-simplex. We write $(g, X, e_0, \ldots, e_p)$ to represent the oriented simplex.

An element $\alpha \in Out(F_r)$ acts on $(g, X, e_0, \ldots, e_p)$ by changing the marking. Specifically, represent α by a homotopy equivalence $f_\alpha \colon R_r \to R_r$ of the standard rose R_r; then $(g, X, e_0, \ldots, e_p) \cdot \alpha = (g \circ f_\alpha, X, e_0, \ldots, e_p)$. This action is transitive on markings: to send $(g, X, e_0, \ldots, e_p)$ to $(g', X.e_0, \ldots, e_p)$, choose any homotopy inverse g^{-1} to g and take $f_\alpha = g^{-1} \circ g'$.

The marking g induces an isomorphism $g_* \colon H_1(R_r) \cong \mathbb{R}^n \to H_1(X)$, thus determining an orientation on $H_1(X)$. Recall that one of the equivalent definitions of an orientation on a connected graph X is a an orientation of $\mathbb{R}E(X) \oplus H_1(X)$ (Proposition 1). Thus we have a map

$$C_p(\widehat{U}_r, A_r) \otimes_{Out(F_r)} \tilde{\mathbb{R}} \to \mathcal{OG}_{p+1-r}$$

sending $(g, X, e_0, \ldots, e_p) \otimes 1 \mapsto (X, or)$, where or is the orientation of $\mathbb{R}E(X) \oplus H_1(X)$ given by the ordering $e_0, \ldots, e_p$ of $E(X)$ together with the orientation of $H_1(X)$ induced by g.

To see that this map is well-defined, note that the action of α changes orientation iff $\det(\bar{\alpha}) = -1$. Since the map is also surjective, and since the vector spaces in question are of the same (finite) dimension, it is an isomorphism. Under this isomorphism, the boundary map

$$\partial \otimes 1 \colon C_p(\widehat{U}_r, A_r) \otimes_{Out(F_r)} \tilde{\mathbb{R}} \to C_{p-1}(\widehat{U}_r, A_r) \otimes_{Out(F_r)} \tilde{\mathbb{R}}$$

is identified with the graph homology boundary map $\partial_E \colon \mathcal{O}\mathcal{G}^{(r)}_{p+1-r} \to \mathcal{O}\mathcal{G}^{(r)}_{p-r}$. □

Note that the twisted coefficients account for the $H_1(X)$ part of a graph's orientation. With trivial coefficients, $H_*^{Out(F_r)}(\widehat{U}_r, A_r; \mathbb{R})$ is isomorphic to a version of graph homology where the orientation is given by ordering the edges only. This homology is considered in a preprint of Bar-Natan and McKay[1].

The twisted equivariant homology groups $H_*^{Out(F_r)}(\widehat{U}_r, A_r; \tilde{\mathbb{R}})$ can be interpreted in terms of certain untwisted relative homology groups as follows. Let $SOut_r$ denote the kernel of the non-trivial action on $\mathbb{R}$, i.e. $SOut_r$ is the preimage of $SL(r,\mathbb{Z})$ under the natural map $Out(F_r) \to GL(r,\mathbb{Z})$.

Theorem 5 *There is a long exact sequence*

$$\cdots \to H_p^{Out(F_r)}(\widehat{U}_r, A_r; \tilde{\mathbb{R}}) \to H_p(\widehat{U}_r/SOut_r, A_r/SOut_r) \to$$
$$\to H_p(\widehat{U}_r/Out(F_r), A_r/Out(F_r)) \to H_{p-1}^{Out(F_r)}(\widehat{U}_r, A_r; \tilde{\mathbb{R}}) \to \cdots.$$

Corollary 9 *Let $\mathcal{O}$ be the commutative operad. Then the shifted $\mathcal{O}$-graph complex $P\overline{\mathcal{O}\mathcal{G}}_{*+1-r}$ fits into a long exact sequence*

$$\cdots \to H_p(P\overline{\mathcal{O}\mathcal{G}}_{*+1-r}) \to H_p(\widehat{U}_r/SOut_r, A_r/SOut_r) \to$$
$$\to H_p(\widehat{U}_r/Out(F_r), A_r/Out(F_r)) \to H_{p-1}(P\overline{\mathcal{O}\mathcal{G}}_{*+1-r}) \to \cdots.$$

In other words, graph homology measures the difference between the homologies of the relative quotients of $(\widehat{U}, A)$ by $SOut$ and by Out.

Proof of Theorem 5 We will suppress the subscripts r to streamline the notation.

Let $\tau \in Out$ be an automorphism with $\det(\bar{\tau}) = -1$ (e.g. $\tau \colon x_1 \leftrightarrow x_2$).

We have $C_p(\widehat{U}/SOut, A/SOut) \cong C_p^+ \oplus C_p^-$, where C_p^+ and C_p^- are the $+1$ and -1 eigenspaces of the involution induced by τ. The short exact sequence

$$1 \to C_p^- \to C_p(\widehat{U}/SOut, A/SOut) \to C_p^+ \to 0$$

gives rise to a long exact homology sequence

$$\cdots \to H_p(C_*^-) \to H_p((\widehat{U}/SOut, A/SOut) \to H_p(C_*^+) \to H_{p-1}(C_*^-) \to \dots$$

Now note that $C_p(\widehat{U}/SOut, A/SOut) \otimes_{\langle\tau\rangle} \mathbb{R} \cong C_p^-$, implying

$$H_p(C_*(\widehat{U}, A) \otimes_{Out} \tilde{\mathbb{R}}) \cong H_p(C_*^-).$$

Furthermore, the map $C_p(\widehat{U}/Out, A/Out) \to C_p(\widehat{U}/SOut, A/SOut)$ given by $\sigma \mapsto \frac{1}{2}(\sigma + \tau\sigma)$ is an isomorphism onto C_p^+. Therefore the above long exact sequence becomes

$$\cdots \to H_p^{Out}(\widehat{U}, A; \tilde{\mathbb{R}}) \to H_p(\widehat{U}/SOut, A/SOut) \to$$
$$\to H_p(\widehat{U}/Out, A/Out) \to H_{p-1}^{Out}(\widehat{U}, A; \tilde{\mathbb{R}}) \to \dots. \qquad \square$$

5.3 Homology of the one-particle irreducible subcomplex versus graph homology

The fact that the subcomplex of one-particle irreducible graphs has the same homology as the full complex for the associative and Lie operads followed because the $\mathcal{O}$ graph homology in these cases can be identified with the homology of a space of graphs, which deformation retracts onto the subspace of one-particle irreducible graphs by uniformly shrinking all separating edges. If the deformation retraction of Outer space U_r extended equivariantly to $(\widehat{U}_r, A_r)$, then our geometric interpretation would show that graph homology, too, could be computed using the subcomplex of one-particle irreducible graphs (see [2], Proposition 7.3). However, there is not even a continuous extension of this deformation retraction to all of $\widehat{U}_r$. The reason is that A_r contains simplices which are obtained by shrinking all *non*-separating edges in a graph to points, leaving a tree. If one then tries to shrink all of the (separating) edges in the tree to points, one is left with a single vertex, which does not correspond to a point in $\widehat{U}_r$.

Let IU_r denote the subspace of U_r corresponding to 1-particle irreducible graphs. Then the homology of the subcomplex of the graph complex spanned by 1-particle irreducible graphs can be identified with the equivariant homology $H_*^{Out(F_r)}(\widehat{IU}_r, A_r \cap \widehat{IU}_r); \tilde{\mathbb{R}})$, as in the proof of Proposition 27, measuring the difference between the relative quotients of $(\widehat{IU}_r, A \cap \widehat{IU}_r)$ modulo $Out(F_r)$ and modulo $SOut(F_r)$. Although $H_*^{Out(F_r)}(\widehat{IU}_r, A_r \cap \widehat{IU}_r; \tilde{\mathbb{R}}) \cong H_*^{Out(F_r)}(\widehat{U}_r, A_r; \tilde{\mathbb{R}})$ for small values of r, this is accidental, as the geometry suggests; for larger values of r the homologies differ. We illustrate the difference between $\widehat{U}$ and $\widehat{IU}$ for $r = 2$, where one can easily draw a picture, below.

For $r = 2$, IU_2 can be identified with the familiar picture of the hyperbolic plane, tiled by ideal triangles (Figure 19). Full Outer space U_2 has one addi-

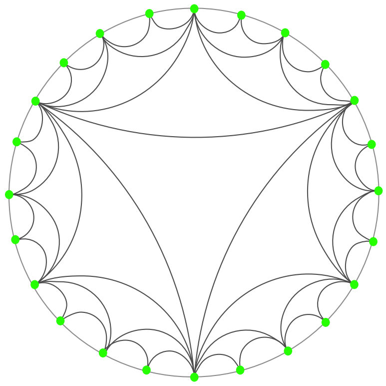

Figure 19: The space $\widehat{IU}_2$

tional open triangle (a "fin") adjacent to each edge in the ideal triangulation. (Figure 20). The simplicial closure $\widehat{IU}_2$ is obtained by adding the vertices of

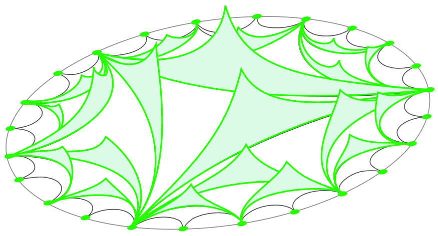

Figure 20: Outer space in rank 2

the ideal triangles; the closure of $\widehat{U}_2$ contains also the additional vertex and two missing edges of each fin.

We conclude by using this example to illustrate Corollary 9. The quotient of $\widehat{U}_2$ by the action of $SOut_2$ is a "pillow" (perhaps more familiar as the quotient of the hyperbolic plane by the action of $SL(2,\mathbb{Z}) \cong SOut_2$) with half of a fin attached (Figure 21); the image of A_2 is a (closed) edge of this fin. Therefore the relative homology $H_*(\widehat{U}_2/SOut_2, A_2/SOut_2)$ has one generator in dimension 2 and is zero elsewhere. In the quotient by the full outer automorphism group $Out(F_2)$, the pillow is flattened to a triangle; again the image of A_2 is an edge of the attached fin, so that the relative homology vanishes. Therefore the only

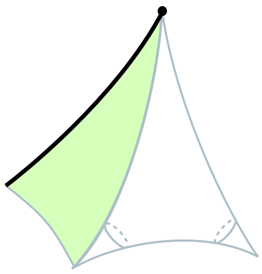

Figure 21: Quotient of $\widehat{U}_2$ modulo $SOut_2$

non-zero terms in the long exact sequence are

$$\cdots 0 \to H_2^{Out(F_2)}(\widehat{U}_2, A_2) \to \mathbb{R} \to 0 \to \cdots .$$

Proposition 27 then tells us that $H_2(\mathcal{OG}^{(2)}) \cong \mathbb{R}$. In terms of oriented graphs, the graph homology generator is the theta graph, with two vertices connected by three edges.

References

[1] **Dror Bar-Natan and Brendan D. McKay**, *Graph cohomology–an overview and some computations*, draft 2001, available at `http://www.ma.huji.ac.il/~drorbn/Misc/`

[2] **Kenneth S. Brown**, *Cohomology of groups.* Graduate Texts in Mathematics, 87. Springer-Verlag, New York, 1982. x+306 pp.

[3] **James Conant**, *Fusion and fission in graph complexes.*, Pacific Journal of Mathematics, Vol. 209, No.2 (2003), 219-230.

[4] **James(Jim) Conant and Karen Vogtmann**, *Infinitesimal operations on chain complexes of graphs.*, Mathematische Annalen, Vol. 327, No. 3 (2003); 545-573.

[5] **James Conant, Ferenc Gerlits and Karen Vogtmann**, *Cut vertices in commutative graphs*, `arXiv:math.QA/0307364`

[6] **Marc Culler and Karen Vogtmann**, *Moduli of graphs and automorphisms of free groups.* Invent. Math. 84 (1986), no. 1, 91–119.

[7] **Ezra Getzler and M. M. Kapranov**, *Cyclic operads and cyclic homology.* Geometry, topology, & physics, 167–201, Conf. Proc. Lecture Notes Geom. Topology, IV, Internat. Press, Cambridge, MA, 1995.

[8] **Ezra Getzler and M.M. Kapranov**, *Modular operads* Compositio Math. 110 (1998), no. 1, 65–126, `arXiv:dg-ga/9408003`

[9] **Allen Hatcher and Karen Vogtmann**, *Rational homology of* $\mathrm{Aut}(F_n)$. Math. Res. Lett. 5 (1998), no. 6,759–780.

[10] **Matthew Horak**, *Mapping class subgroups of outer automorphism groups of free groups*, thesis, Cornell University, 2003, `arXiv:math.GT/0310328`

[11] **M. Kapranov and Y. Manin** *Modules and Morita theory for operads*, Amer. J. Math. 123 (2001) 811–838.

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

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

[14] **Jerome Levine**, *Homology cylinders: an enlargement of the mapping class group.* Algebr. Geom. Topol. 1 (2001), 243–270

[15] **Swapneel Mahajan**, *Symplectic operad geometry and graph homology*, `arXiv:math.QA/0211464`

[16] **Martin Markl**, *Cyclic operads and homology of graph complexes.* The 18th Winter School "Geometry and Physics" (Srn, 1998). Rend. Circ. Mat. Palermo (2) Suppl. No. 59 (1999), `arXiv:math.QA/9801095`

[17] **Martin Markl, Steve Shnider and Jim Stasheff**, *Operads in algebra, topology and physics.* Mathematical Surveys and Monographs, 96. American Mathematical Society, Providence, RI, 2002. x+349 pp

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

[19] **R. C. Penner**, *Perturbative series and the moduli space of Riemann surfaces.* J. Differential Geom. 27 (1988), no. 1, 35–53.

[20] **John Smillie and Karen Vogtmann**, *A generating function for the Euler characteristic of* $\mathrm{Out}(F_n)$. Proceedings of the Northwestern conference on cohomology of groups (Evanston, Ill., 1985). J. Pure Appl. Algebra 44 (1987), no. 1-3, 329–348.

[21] **Dylan Thurston**, Undergraduate thesis, Harvard, `arXiv:math.QA/9901110`

[22] **Karen Vogtmann**, *End invariants of the group of outer automorphisms of a free group.* Topology 34 (1995), no. 3, 533–545

[23] **Karen Vogtmann**, *Automorphisms of free groups and outer space*, Geometriae Dedicata 94 (1): 1-31, October 2002

[24] **Hermann Weyl**, *Classical Groups*, Princeton University Press, Princeton, NJ 1939; second edition, 1946.

[25] **Heiner Zieschang, Elmar Vogt and Hans-Dieter Coldewey**, *Surfaces and planar discontinuous groups.* Translated from the German by John Stillwell. Lecture Notes in Mathematics, 835. Springer, Berlin, 1980. x+334 pp.

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

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

Received: 5 February 2003 Revised: 1 December 2003