

CUT VERTICES IN COMMUTATIVE GRAPHS

by JAMES CONANT[†]

(Department of Mathematics, University of Tennessee at Knoxville, Knoxville,
TN 37996-1300, USA)

FERENC GERLITS[‡]

(Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences, Reáltanoda utca
13-15, H-1053, Budapest, Hungary)

and KAREN VOGTMANN[§]

(Department of Mathematics, Cornell University, Ithaca, NY 14853-4201, USA)

[Received 2 September 2004]

Abstract

The homology of Kontsevich's commutative graph complex parametrizes finite type invariants of odd-dimensional manifolds. This *graph homology* is also the twisted homology of Outer Space modulo its boundary, so gives a nice point of contact between geometric group theory and quantum topology. In this paper we give two different proofs (one algebraic, one geometric) that the commutative graph complex is quasi-isomorphic to the quotient complex obtained by modding out by graphs with cut vertices. This quotient complex has the advantage of being smaller and hence more practical for computations. In addition, it supports a Lie bialgebra structure coming from a bracket and cobracket we defined in a previous paper. As an application, we compute the rational homology groups of the commutative graph complex up to rank 7.

1. Introduction

Graph homology was introduced by Kontsevich [8, 9], who showed that it computes the homology of a certain infinite-dimensional Lie algebra c_∞ , and also parametrizes invariants of certain odd-dimensional manifolds. The best understood of these invariants are those associated to (rational) homology three-spheres. These are known as 'finite type' invariants, and are analogs of the Goussarov–Vassiliev knot invariants. Alternative constructions of these finite type invariants have been found by Le, Murakami and Ohtsuki [11], Kuperberg and Thurston [10], and Bar-Natan, Garoufalidis, Rozansky and Thurston [1].

Graph homology has a very simple definition. The degree k term of the graph complex \mathcal{G} is spanned (over a field of characteristic 0) by connected, 'oriented' graphs with k vertices, and the boundary operator $\partial : \mathcal{G}_k \rightarrow \mathcal{G}_{k-1}$ is defined on a graph G by adding together all oriented graphs

[†]E-mail: jconant@math.utk.edu

[‡]E-mail: fgerlits@renyi.hu

[§]Corresponding author. E-mail: vogtmann@math.cornell.edu

which can be obtained from G by collapsing a single edge. The notion of orientation is the most subtle part of the whole story (see [5, section 2.3.1] for the many equivalent notions), but suffice it to say that it guarantees that $\partial^2 = 0$. Graph homology is then the homology of this complex.

Graph homology is a special case of a more general construction, where the graph complex is spanned by graphs decorated at each vertex by an element of a cyclic operad (see, for example [5]). Ordinary graph homology corresponds to the commutative operad. Two other important examples, also studied by Kontsevich [8, 9], are obtained by using the associative operad and the Lie operad; the homology of the resulting graph complexes is closely related to the cohomology of mapping class groups of punctured surfaces and the cohomology of outer automorphisms of free groups, respectively.

Each general graph complex \mathcal{G} may be considered as the primitive part of a graded Hopf algebra \mathcal{HG} , where the product on \mathcal{HG} is given by disjoint union. In [4] we introduced a Lie bracket and cobracket on \mathcal{HG} . These do not form a compatible bialgebra structure on \mathcal{HG} , and they do not restrict to \mathcal{G} . However, in [4] we also introduced the subcomplex \mathcal{B} of \mathcal{G} spanned by connected graphs with no separating edges, and showed that the Lie bracket and cobracket do restrict to \mathcal{B} ; furthermore, they are compatible, so give \mathcal{B} a Lie bialgebra structure. In the associative and Lie cases, the subcomplex \mathcal{B} is quasi-isomorphic to \mathcal{G} , but this is not true in the commutative case: \mathcal{B} and \mathcal{G} do not have the same homology.

In this paper we use a different approach in the commutative case to find a smaller chain complex quasi-isomorphic to \mathcal{G} which carries a Lie bialgebra structure. Specifically, we consider the subcomplex \mathcal{C} of \mathcal{G} spanned by graphs with at least one cut vertex, where a cut vertex is defined as a vertex whose deletion disconnects the graph (Fig. 1). In section 2 we use a spectral sequence argument to prove the following.

THEOREM 1.1 *The quotient map of chain complexes $\mathcal{G} \rightarrow \mathcal{G}/\mathcal{C}$ is an isomorphism on homology.*

In section 3 we recall the geometric interpretation of graph homology in terms of Outer Space from [5], and re-prove Theorem 1.1 from this point of view. Specifically, we show that the standard deformation retraction of Outer Space onto the subspace of graphs with no separating edges extends to the Bestvina–Feighn bordification of Outer Space, and we show that the image of the points at infinity consists precisely of the closure of the space of graphs with cut vertices. This deformation retraction induces a homology isomorphism on certain twisted chain complexes, which can be identified with \mathcal{G} and \mathcal{G}/\mathcal{C} .

In section 4 we recall the definition of the Lie bracket and cobracket, and prove the following.

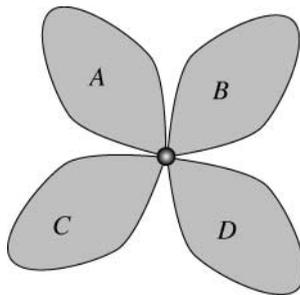


Fig. 1 Schematic for a cut vertex

THEOREM 1.2 *The Lie bracket and cobracket on \mathcal{HG} induce a compatible graded Lie bialgebra structure on \mathcal{G}/\mathcal{C} .*

REMARK Despite the fact that the entire graph complex \mathcal{HG} does not support a Lie bialgebra structure, Wee-Liang Gan [7] has recently shown that it supports a strongly homotopy Lie bialgebra structure, and that this reduces to our Lie bialgebra structure when one mods out by graphs with cut vertices.

Finally, in the last section we exploit the fact that the quotient complex \mathcal{G}/\mathcal{C} is smaller than \mathcal{G} to do some computer-aided calculations of graph homology. Specifically, the elimination of cut vertices reduces the size of the vector spaces involved by about 30 per cent, allowing us to calculate graph homology up to rank 7.

2. Graphs with cut vertices

Let G be an oriented graph with no separating edges. An oriented graph H is said to *retract to G* if G can be obtained from H by collapsing each separating edge of H to a point. Denote by \mathcal{R}^G the subspace of \mathcal{G} spanned by all graphs H which retract to G . Notice that according to our definitions, \mathcal{R}^G is one-dimensional (with basis G) unless G has a cut vertex. Define a boundary operator $\partial_s : \mathcal{R}_k^G \rightarrow \mathcal{R}_{k-1}^G$ by

$$\partial_s(G) = \sum_{e \text{ separating}} G_e,$$

where G_e denotes the graph obtained from G by collapsing the edge e .

LEMMA 2.1 *Let G be a connected graph with at least one cut vertex, but no separating edges. Then $(\mathcal{R}^G, \partial_s)$ is an acyclic complex.*

Proof. First we prove the lemma under the assumption that G has no automorphisms. Fix a cut vertex v of G , and let c_1, \dots, c_l be the connected components of $G \setminus \{v\}$. Let \mathcal{R}^v denote the subcomplex of \mathcal{R}^G spanned by graphs whose separating edges form a tree which collapses to v . We can consider \mathcal{R}^v to be graded by the number of edges in this tree. In order to prove that \mathcal{R}^v is acyclic, we will identify it with the augmented chain complex of a contractible simplicial complex X^v .

A low-degree example of the complex X^v , in the case where v cuts the graph into three components, is shown in Fig. 2. In this case there is only one way to replace v by a trivalent tree, and the complex X^v is a single 2-simplex. Face maps of this simplex correspond to collapsing edges of the tree. Recall that generators of \mathcal{G} are *oriented* graphs. One notion of orientation for a commutative graph is an orientation of the vector space $\mathbb{R}^{\{\text{edges}\}} \oplus H_1(\text{Graph}; \mathbb{R})$. Thus the orientation on G is given by ordering of the edges of G and orienting the first homology. To compatibly orient a generator of \mathcal{R}^v we need only to order the additional edges, that is, the edges of the tree collapsing to v . Notice that such an ordering induces an orientation of the 2-simplex in our example.

To describe X^v in general, it is easier to translate from the language of trees to that of *compatible partitions* of $\{c_i\}$. If H is a graph in \mathcal{R}^v and T is the tree in H which collapses to v , then each edge e of T partitions the components c_i into two disjoint sets (more precisely, it partitions the preimages in H of the c_i). Different edges correspond to different partitions, which are

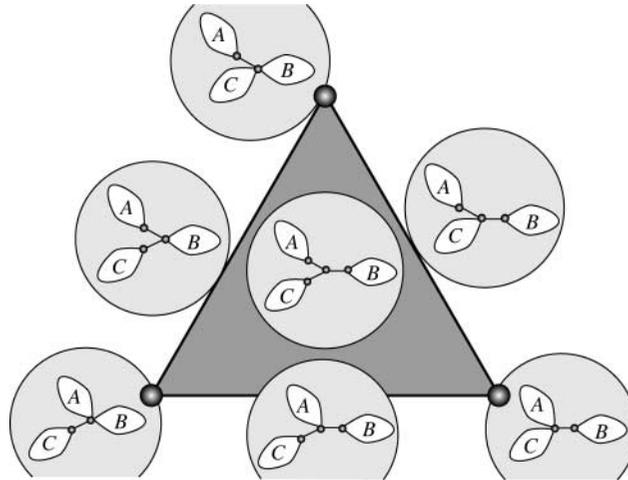


Fig. 2 The complex X^v , when $G - v$ has three components

compatible in the following sense: If $P_1 = X_1 \cup Y_1$ and $P_2 = X_2 \cup Y_2$, then

$$X_1 \subset X_2, X_1 \subset Y_2, Y_1 \subset X_2, \text{ or } Y_1 \subset Y_2.$$

Conversely, any set of pairwise compatible partitions determines a pair (H, T) which collapses to (G, v) . **Figure 3** shows a tree that blows up a vertex and two compatible partitions corresponding to two edges of the tree.

Note that \mathcal{R}_0^v is one-dimensional, spanned by G . Thus \mathcal{R}^v is the augmented chain complex of the simplicial complex X^v whose vertices are partitions of the set $\{c_1, \dots, c_l\}$ into two subsets. A set of $k + 1$ partitions forms a k -simplex of X^v if partitions in the set are pairwise compatible. Let P be the partition which separates c_1 from all other components. Given any set S of pairwise compatible partitions, $S \cup \{P\}$ is also a set of pairwise compatible partitions. Thus X^v is a cone on the vertex P ; in particular, it is contractible, and \mathcal{R}^v is acyclic.

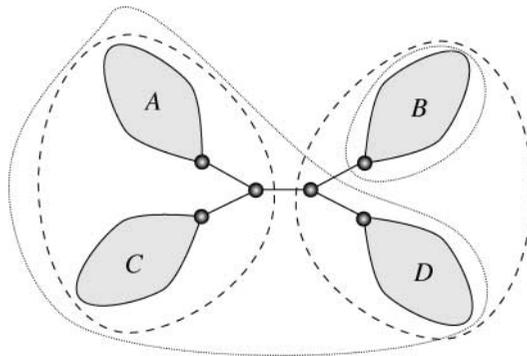


Fig. 3 Two compatible partitions of the set $\{A, B, C, D\}$ corresponding to two edges of a tree

If we grade the entire complex $(\mathcal{R}^G, \partial_s)$ by the number of edges in the forest formed by all separating edges, then it is the tensor product of the complexes $(\mathcal{R}^v, \partial_s)$ for cut vertices v of G . Thus \mathcal{R}^G is acyclic.

Now we return to the general case when the graph has a non-trivial automorphism group. Let $\hat{\mathcal{R}}^G$ be the chain complex of graphs obtained from \mathcal{R}^G by distinguishing all edges and vertices in each graph. (This kills off automorphisms.) Note that $Aut(G)$ acts on $\hat{\mathcal{R}}^G$, and that $\hat{\mathcal{R}}^G/Aut(G) \cong \mathcal{R}^G$. Since \mathcal{R}^G is acyclic, and since finite groups have no rational homology, the chain complexes $\hat{\mathcal{R}}^G$ and \mathcal{R}^G are rationally quasi-isomorphic. Now we are back in the case when graphs have no automorphisms, and we are done.

THEOREM 2.2 *The subcomplex \mathcal{C} of \mathcal{G} spanned by graphs with at least one cut vertex is acyclic.*

Proof. Define $\mathcal{C}_{s,k}$ to be the subspace of \mathcal{C}_k spanned by graphs with exactly s separating edges. Then

$$\mathcal{C}_k = \bigoplus_{s \geq 0} \mathcal{C}_{s,k}.$$

The boundary operator $\partial : \mathcal{C}_k \rightarrow \mathcal{C}_{k-1}$ is the sum of two boundary operators ∂_s and ∂_{ns} , where ∂_s collapses only separating edges, and ∂_{ns} collapses only non-separating edges. These two boundary operators make \mathcal{C} into the total complex of a double complex E^0 , where $E^0_{p,q} = \mathcal{C}_{p,p+q}$, the vertical arrows are given by ∂_s and the horizontal arrows by ∂_{ns} .

$$\begin{array}{ccccc} & & \downarrow & & \downarrow \\ \leftarrow & \mathcal{C}_{p,p+q-1} & \xleftarrow{\partial_{ns}} & \mathcal{C}_{p,p+q} & \leftarrow \\ & \downarrow \partial_s & & \downarrow \partial_s & \\ \leftarrow & \mathcal{C}_{p-1,p+q-2} & \xleftarrow{\partial_{ns}} & \mathcal{C}_{p-1,p+q-1} & \leftarrow \\ & \downarrow & & \downarrow & \end{array}$$

Note that a graph with p separating edges has at least $p + 1$ vertices, so that $E^0_{p,q} = \mathcal{C}_{p,p+q} = 0$ for $q < 1$, and the double complex is a first quadrant double complex.

We consider the spectral sequence associated to the vertical filtration of this double complex. This spectral sequence converges to the homology of the total complex \mathcal{C} . The $E^1_{p,q}$ term is equal to $H_p(E^0_{*,q}, \partial_s)$, that is, the p th homology of the q th column.

For each q , the column $E^0_{*,q}$ breaks up into direct sum of chain complexes E^G_* , one for each graph G with q vertices (at least one of which is a cut vertex) and no separating edges. A graph in $\mathcal{C}_{p,p+q}$ is in E^G_* if G is the result of collapsing all of its separating edges, that is, $E^G_* = \mathcal{R}^G$. By Lemma 2.1, \mathcal{R}^G has no homology, so that $E^1_{p,q} = 0$ for all p and q , and the complex \mathcal{C} is acyclic.

Theorem 1.1 now follows immediately by the long exact homology sequence of the pair $(\mathcal{G}, \mathcal{C})$.

3. Geometric interpretation, in terms of Outer Space

In this section we will sketch a geometric proof of the main theorem. This proof relies on the identification of the graph homology chain complex with a twisted relative chain complex for Outer Space, as described in [5], and also on a generalization of the *Borel–Serre bordification*

of Outer Space defined by Bestvina and Feighn [2]. A similar generalization is mentioned as a remark in their paper, but details of proofs are not worked out.

Recall that Outer Space X_n is a topological space which parametrizes finite marked metric graphs with (free) fundamental group of rank n (see [12]). Outer Space can be decomposed as a union of open simplices, and there are several ways to add a boundary to this space. The simplest is to formally add the union of all missing faces to obtain a simplicial complex \bar{X}_n , called the *simplicial closure of Outer Space*. The bordification is more subtle; it is a blown-up version of \bar{X}_n , which we will denote by \hat{X}_n . The interiors of \bar{X}_n and \hat{X}_n are both homeomorphic to X_n , and the action of $Out(F_n)$ extends to the boundaries $\partial\bar{X}_n$ and $\partial\hat{X}_n$. There is a natural quotient map $q: \hat{X}_n \rightarrow \bar{X}_n$, which is a homeomorphism on the interiors and in general has contractible point inverses.

In [5] we showed that the subcomplex $\mathcal{G}^{(n)}$ of the graph complex \mathcal{G} spanned by graphs with fundamental group of rank n can be identified with the relative chains on $(\bar{X}_n, \partial\bar{X}_n)$, twisted by the non-trivial *determinant* action of $Out(F_n)$ on \mathbb{R} . Blowing up the boundary does not change this picture; $\mathcal{G}^{(n)}$ is also identified with the relative chains on $(\hat{X}_n, \partial\hat{X}_n)$, twisted by the same non-trivial action of $Out(F_n)$ on \mathbb{R} .

In this section we define an equivariant deformation retraction $\hat{X}_n \rightarrow \hat{Y}_n$, where Y_n is the subspace of X_n consisting of graphs with no separating edges, and \hat{Y}_n denotes the closure of Y_n in \hat{X}_n . The image of $\partial\hat{X}_n$ under this retraction, denoted by \hat{Z}_n , is the union of $\partial\hat{Y}_n$ and the set of graphs with a cut vertex but no separating edges. The deformation retraction induces an isomorphism

$$C_*(\hat{X}_n, \partial\hat{X}_n) \otimes_{Out(F_n)} \mathbb{R} \cong C_*(\hat{Y}_n, \hat{Z}_n) \otimes_{Out(F_n)} \mathbb{R}.$$

Tracing through the identification of $C_*(\hat{X}_n, \partial\hat{X}_n) \otimes_{Out(F_n)} \mathbb{R}$ with $\mathcal{G}^{(n)}$, we see that the chains

$$C_*(\hat{Y}_n, \hat{Z}_n) \otimes_{Out(F_n)} \mathbb{R}$$

are identified with $\mathcal{G}_{ns}^{(n)}/\mathcal{C}_{ns}^{(n)}$, where the subscript *ns* denotes the subcomplex spanned by graphs with no separating edges. Since all graphs with separating edges also have cut vertices, this is naturally isomorphic to $\mathcal{G}^{(n)}/\mathcal{C}^{(n)}$. This completes the sketch of the proof of Theorem 1.1, modulo the definition of the bordification and the retraction. The remainder of the section is devoted to just that.

It has long been known that Y_n is an equivariant deformation retract of X_n , but the deformation retraction, which uniformly shrinks all separating edges while uniformly expanding all other edges, does not extend to \bar{X}_n . One can see this even for $n = 2$ by considering the 2-simplex corresponding to the ‘barbell’ graph (see Fig. 4). The deformation retraction sends each horizontal slice linearly on to the bottom edge of the triangle, so that the deformation cannot be extended continuously to the top vertex of the closed triangle.

This difficulty can be resolved by blowing up the vertex of the triangle to a line, which records the (constant) ratio of the lengths of the two loops of the barbell graph along a geodesic in X_2 coming into the vertex (see Fig. 4). This is the idea of the Bestvina–Feighn bordification \hat{X}_n . Similar ideas are also used in the Borel–Serre compactification of a homogenous space and the Fulton–MacPherson and Axelrod–Singer compactifications of configuration spaces of points in a manifold.

To describe \hat{X}_n in general, we need the notion of a *core graph*, which is defined to be a (not necessarily connected) graph with no separating edges and no vertices of valence 0 or 1. Every graph has a unique maximal core subgraph, called its *core*.

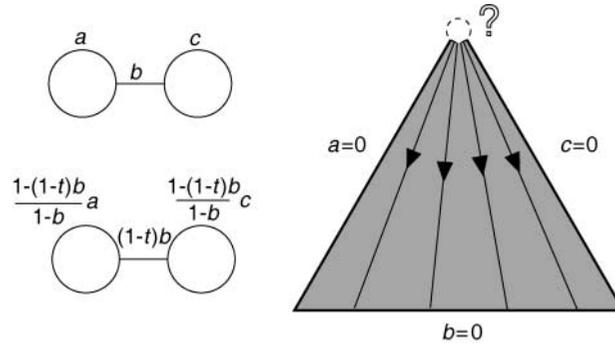


Fig. 4 Why the simplicial closure does not work

A point in X_n is a marked, non-degenerate metric graph of rank n with total volume 1, where *non-degenerate* means no edge is assigned the length 0. In the bordification \hat{X}_n , we allow edges of a core subgraph to have length 0, but in this case there is a secondary metric, also of volume 1, given on the core subgraph. The secondary metric may also be zero on a smaller core subgraph, in which case there is a third metric of volume 1 on that core subgraph, etc.

In general, a point of \hat{X}_n consists of a marked metric graph Γ_0 and a properly nested (possibly empty) sequence

$$\Gamma_1 \supset \Gamma_2 \supset \dots \supset \Gamma_k$$

of core subgraphs of Γ_0 . Each Γ_i is equipped with a metric of volume 1; the metric on Γ_0 is the *primary metric*, the metric on Γ_1 the *secondary metric*, etc. Each Γ_i is the subgraph of Γ_{i-1} spanned by all edges of length 0, and the chain is *non-degenerate* in the sense that every edge of Γ_0 has non-zero length in exactly one Γ_i . The space \hat{X}_n is stratified as a union of open cells; the dimension of the cell containing $x = (\Gamma_0 \supset \Gamma_1 \supset \dots \supset \Gamma_k)$ is $e(\Gamma_0) - k - 1$, where $e(\Gamma_0)$ is the number of edges of Γ_0 .

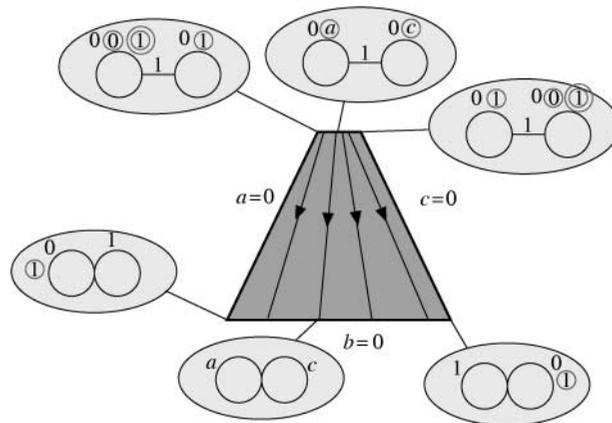


Fig. 5 The deformation retraction on a cell of the bordification

This construction is illustrated for $n = 2$ in Fig. 5. In this figure, the number of circles surrounding an edge length corresponds to the hierarchy of metrics. A sequence of graphs on which the volume of a core subgraph is shrinking to zero will approach a point on the boundary which depends on the relative lengths of edges in the core subgraph. If the metric is shrinking uniformly on the core subgraph, then the limit is the graph whose primary metric vanishes on the core subgraph, and where the secondary metric on the core subgraph is a rescaled version of the metric restricted to the shrinking core subgraph. If parts of the core subgraph are shrinking at a faster rate than others, the sequence will land in a face of higher codimension.

Bestvina and Feighn prove the following theorem.

THEOREM 3.1 (Bestvina–Feighn) *\hat{Y}_n is contractible, and the $Out(F_n)$ action on the interior extends to the whole space.*

The following theorem shows that \hat{X}_n is also contractible, and identifies the image of the boundary $\partial\hat{X}_n$.

THEOREM 3.2 *\hat{X}_n equivariantly deformation retracts onto \hat{Y}_n . Under this retraction, the image \hat{Z}_n of $\partial\hat{X}_n$ is the union of $\partial\hat{Y}_n$ with the set of graphs in the interior Y_n which have a cut vertex.*

Proof. Let $x = (\Gamma_0 \supset \dots \supset \Gamma_k)$ be a point in \hat{X}_n . We define an equivariant deformation retraction $\phi(x, t)$ as follows.

If Γ_1 is not the maximal core of Γ_0 , then the deformation retraction ϕ changes the metric on Γ_0 by uniformly shrinking all separating edges and rescaling the primary metric on the rest of the graph by a global factor to retain total volume 1. The metrics on Γ_k for $k \geq 1$ are not affected. If, on the other hand, Γ_1 is the maximal core of Γ_0 , then the deformation retraction shrinks the separating edges in Γ_0 (that is, it shrinks $\Gamma_0 \setminus \Gamma_1$) while simultaneously blowing up the initially degenerate Γ_1 by an appropriate factor in the primary metric. In other words, Γ_1 immediately disappears from the filtration. The metrics on Γ_i for $i \geq 2$ are not affected.

We now give an explicit formula for $\phi(x, t)$. Let m_i denote the metric on Γ_i , let S denote the set of separating edges in Γ_0 , and let $(\Gamma_i)_S$ be the image of Γ_i under the map which collapses each edge in S to a point. The formula depends on whether $m_0(S) = 1$ (so that Γ_1 is the entire core of Γ_0) or $m_0(S) < 1$.

If $m_0(S) < 1$, then for $0 < t < 1$ we have $\phi(x, t) = (\Gamma_0 \supset \Gamma_1 \supset \dots \supset \Gamma_k)$, where the new metric n_i on Γ_i is given by

$$n_i(e) = \begin{cases} \left(1 + t \left(\frac{m_0(S)}{1 - m_0(S)}\right)\right) m_0(e), & i = 0, e \in \Gamma_0 \setminus S, \\ (1 - t) m_0(e), & i = 0, e \in S, \\ m_i(e), & i > 0, e \in \Gamma_i. \end{cases}$$

For $t = 1$, $\phi(x, t) = ((\Gamma_0)_S \supset (\Gamma_1)_S \supset \dots \supset (\Gamma_k)_S)$, where $(\Gamma_i)_S$ is the image of Γ_i under the map which collapses each edge in S to a point. The length of each edge $e \in (\Gamma_0)_S$ is $(1 - m_0(S))^{-1} m_0(e)$, and the length of $e \in (\Gamma_i)_S$ is $(m_i)e$.

If $m_0(S) = 1$, then for $0 < t < 1$ we have $\phi(x, t) = \Gamma_0 \supset \Gamma_2 \supset \dots \supset \Gamma_k$, where the new metric n_i on Γ_i is given by

$$n_i(e) = \begin{cases} tm_1(e), & i = 0, e \in \Gamma_0 \setminus S, \\ (1 - t)m_0(e), & i = 0, e \in S, \\ m_i(e), & i > 1, e \in \Gamma_n. \end{cases}$$

If $t = 1$, then $\phi(x, t) = ((\Gamma_1)_S \supset (\Gamma_2)_S \supset \dots \supset (\Gamma_k)_S)$, where the length of each edge e of (Γ_i) is equal to $m_i(e)$ for all i .

The deformation retraction restricted to a cell in the case $n = 2$ is pictured in Fig. 5. The top line corresponds to graphs with a degenerate core, and the flow pushes them into strata of \hat{X}_n of one higher dimension. Everywhere else, the flow stays within strata until $t = 1$, when the dimension of the stratum may decrease.

The fact that points in $\partial \hat{X}_n$ land in \hat{Z}_n is clear. Now we attack the question of continuity. For this, it will be convenient to fix a metric on each closed cell of \hat{X}_n . Every top-dimensional cell is associated to a marked trivalent graph, Γ . Call such a closed cell $\hat{\Sigma}_\Gamma$. Let C be a core subgraph of Γ . For every point of $\hat{\Sigma}_\Gamma$, C has a *level*, which is the unique i such that the metric m_i is defined on C and is not identically zero on C . (If we are looking at a point on $\partial \hat{\Sigma}_\Gamma$ where a subforest has been contracted, then the level is defined for the image of C under this contraction.)

Let x, x' be two points in $\hat{\Sigma}_\Gamma$; let $C \subset \Gamma$ be a core subgraph and let l, l' be the levels of C in x and x' . Then define

$$d_C(x, x') = \sum_{e \in E(C)} \left| \frac{m_l(e)}{m_l(C)} - \frac{m_{l'}(e)}{m_{l'}(C)} \right|.$$

Then the metric on $\hat{\Sigma}_\Gamma$ is defined to be

$$d(x, x') = \sum_{C \subset \Gamma} d_C(x, x'),$$

where the sum is over all core subgraphs of Γ including Γ itself. That this metric generates the appropriate topology follows from [2, Lemma 2.3].

To show that ϕ is continuous it suffices to show that on each closed cell $\hat{\Sigma}_\Gamma$, the functions $\phi(x_0, t)$ are *equicontinuous* as a family indexed by x_0 and that each function $\phi(x, t_0)$ is continuous as a function of x . Recall that *equicontinuous* means that for every $\varepsilon > 0$ there exists $\delta > 0$ such that $|t - s| < \delta \Rightarrow \forall x_0 (d(\phi(x_0, t), \phi(x_0, s)) < \varepsilon)$.

The continuity of $\phi(x_0, t)$ as a function of t is clear except when x_0 represents $\Gamma \supset \Gamma_1 \supset \dots \supset \Gamma_k$, and Γ_1 is the maximal core of Γ . However, here too ϕ is continuous, since by construction of $\hat{X}_n, x_0 := \lim_{t \rightarrow 0} \phi(x_0, t)$. Note that as functions of t , the formula for how the length of each edge changes is a linear map with coefficients bounded by 1. This ensures that the family of functions is equicontinuous, since

$$d(\phi(x_0, t), \phi(x_0, s)) = \sum_{C \subset \Gamma} d_C(\phi(x_0, t), \phi(x_0, s)) \leq \sum_{C \subset \Gamma} \sum_{e \in E(C)} |t - s| \leq N \cdot |t - s|,$$

where N is a constant independent of x_0 .

Now we wish to show continuity in x . Clearly ϕ is continuous on the interiors of cells, so we need to consider what happens as we approach the boundary. It will be simpler to analyse what happens as we go from a cell Σ to a codimension-one stratum B . Suppose Σ corresponds to the sequence of graphs $\Gamma_0 \supset \Gamma_1 \supset \dots \supset \Gamma_k$. Then the face B comes from one of two processes. Either it corresponds to contracting an edge $e : (\Gamma_0)_e \supset (\Gamma_1)_e \supset \dots \supset (\Gamma_k)_e$, or it corresponds to refining the filtration by inserting a new core graph $C : \Gamma_0 \supset \dots \supset \Gamma_i \supset C \supset \Gamma_{i+1} \supset \dots \supset \Gamma_k$. For every point x on B , there is a canonical path, P_x , into Σ . In the first case, it is defined by expanding the contracted edge in the metric that makes sense, shrinking the other edges in that metric to maintain total volume 1. In the second case, the core C is expanded from length zero in the i th metric, using a scaled version of the metric that had been defined on C . The other edges in Γ_i are scaled down to maintain total volume 1.

We will show that, for every $\varepsilon > 0$ there is δ such that

$$\forall x \in B \forall y \in P_x (d(x, y) < \delta \Rightarrow d(\phi(x, t_0), \phi(y, t_0)) < \varepsilon).$$

This condition will be called *boundary equicontinuity*. This is sufficient to ensure continuity. For example, to show continuity at a point z on a codimension-2 face, let x be a nearby point in the top cell. Let y be the projection onto one of the nearby codimension-1 faces (that is, $x \in P_y$), and z' the projection onto the codimension-2 face (that is, $y \in P_{z'}$). Thus if x is sufficiently close to z , then x, y are close, y, z' are close, and z, z' are close. Then

$$d(\phi(x, t_0), \phi(z, t_0)) \leq d(\phi(x, t_0), \phi(y, t_0)) + d(\phi(y, t_0), \phi(z, t_0)) + d(\phi(z, t_0), \phi(z', t_0)),$$

and by the boundary equicontinuity hypothesis we can make the first two terms uniformly less than $\varepsilon/3$ and by continuity on the interior of cells at z we can bound the last term by $\varepsilon/3$.

So now let us show boundary equicontinuity. Let x be an interior point and x' be nearby on a codimension-1 face, such that $x \in P_{x'}$.

As mentioned above, in one case, $x' = (\Gamma_0)_e \supset \dots \supset (\Gamma_k)_e$, where the metric on edges is unchanged except in the image of the (unique) graph Γ_i of the filtration in which e has non-zero length; in $(\Gamma_i)_S$, edges are scaled by $1/(1 - m_i(e))$. The fact that x is close to x' means that $m_i(e)$ is very small.

In the other case $x' = \Gamma_0 \supset \dots \supset \Gamma_i \supset C \supset \Gamma_{i+1} \supset \dots \supset \Gamma_k$. The metric on C is $1/m_i(C)$ times the restriction of m_i to C . The metric on Γ_i is 0 on C , and $\frac{1}{1 - m_i(C)} m_i$ on edges not in C . The fact that x is close to x' means that $m_i(C)$ is small.

It is now routine to check that $\phi(x, t)$ is uniformly close to $\phi(x', t)$ in all cases. As an example, we check one of the more complicated cases, when $x' = \Gamma_0 \supset C \supset \Gamma_1 \supset \dots \supset \Gamma_k$, where C is the core of Γ_0 . Let $|e|$ denote the primary metric on x . Then $|S| + |C| = 1$. The primary length of e in x' is 0 if $e \in C$ and $|e|/|S|$ if $e \in S$. The secondary length of $e \in C$ in x' is $|e|/|C|$. We now compute $\phi(x, t)$ and $\phi(x', t)$ using the formulae above.

Note that Γ_1 is *not* the core of Γ_0 , so

$$\phi(x, t) = (\Gamma_0 \supset \Gamma_1 \supset \dots \supset \Gamma_k)$$

for $0 < t < 1$. The primary length of e is

$$|e|_{\phi(x,t)} = \begin{cases} (1-t)|e|, & e \in S, \\ \left(1 + t \frac{|S|}{1-|S|}\right)|e| = ((1-t)|C| + t) \frac{|e|}{|C|}, & e \in C. \end{cases}$$

On the other hand, C is the core of Γ_0 so again

$$\phi(x', t) = (\Gamma_0 \supset \Gamma_1 \supset \dots \supset \Gamma_k)$$

for $0 < t < 1$. Now the primary length of e is

$$|e|_{\phi(x',t)} = \begin{cases} (1-t) \frac{|e|}{|S|}, & e \in S, \\ t \frac{|e|}{|C|}, & e \in C. \end{cases}$$

Now, to show equicontinuity at this boundary, we calculate distances. First, we claim that $d(x, x') = d_{\Gamma_0}(x, x')$. So let $D \neq \Gamma_0$ be a core subgraph of Γ_0 . Then $D \subset C$. If the primary metric of x vanishes on D , then the first non-vanishing metric is the same for both x and x' , and so $d_D(x, x') = 0$. If D does not vanish in the primary metric of x' , D vanishes in the primary metric of x , and is rescaled in the secondary metric: $m'_2|_D = m_1|_D \cdot 1/|C|$. Then

$$d_D(x, x') = \sum_{e \in E(D)} \left| \frac{m_1(e)}{m_1(D)} - \frac{m'_2(e)}{m'_2(D)} \right| = \sum_{e \in E(D)} \left| \frac{m_1(e)}{m_1(D)} - \frac{m_1(e) \cdot |C|^{-1}}{m_1(D) \cdot |C|^{-1}} \right| = 0.$$

So we have

$$\begin{aligned} d(x, x') = d_{\Gamma_0}(x, x') &= \sum_{e \in E(C)} ||e| - |e'|| + \sum_{e \in E(S)} ||e| - |e'|| = \sum_{e \in E(C)} ||e| - 0| + \sum_{e \in E(S)} \left| |e| - \frac{|e|}{|S|} \right| \\ &= |C| + |S||1 - 1/|S|| = 2|C|. \end{aligned}$$

On the other hand

$$\begin{aligned} d(\phi(x, t_0), \phi(x', t_0)) &= d_{\Gamma_0}(\phi(x, t_0), \phi(x', t_0)) = \sum_{e \in E(S)} \left| (1-t_0)|e| - (1-t_0) \frac{|e|}{|S|} \right| \\ &+ \sum_{e \in E(C)} \left| (1-t_0)|C| + t_0 \frac{|e|}{|C|} - t_0 \frac{|e|}{|C|} \right| = (1-t_0)(|S||1 - 1/|S|| + |C|) = 2(1-t_0)|C|. \end{aligned}$$

Thus we can take $\delta = \varepsilon/(1-t_0)$, which is independent of x .

4. Lie bialgebra structure on \mathcal{G}/\mathcal{C}

Recall that \mathcal{HG} denotes the Hopf algebra spanned by all oriented graphs (not necessarily connected). In this section, we will show that the Lie bracket and cobracket on \mathcal{HG} introduced in [4] induce a Lie bracket and cobracket on \mathcal{G}/\mathcal{C} , and that these are compatible on \mathcal{G}/\mathcal{C} .

We first recall the definition of the Lie bracket. Let G be a graph, and let x and y be half-edges of G , terminating at the vertices v and w respectively. Form a new graph as follows: Cut the edges of G containing x and y in half and glue x to y to form a new edge xy , with vertices v and w . If y was not the other half of x (that is, $y \neq \bar{x}$), there are now two ‘dangling’ half-edges \bar{x} and \bar{y} . Glue these to form another new edge $\bar{x}\bar{y}$. Finally, collapse the edge xy to a point. We say the resulting graph, denoted by G_{xy} , is obtained by *contracting the half-edges x and y* .

Recall that the Hopf algebra product $G \cdot H$ is the disjoint union of G and H , with appropriate orientation. The bracket of G and H is defined to be the sum of all graphs obtained by contracting a half-edge of G with a half-edge of H in $G \cdot H$:

$$[G, H] = \sum_{x \in G, y \in H} (G \cdot H)_{xy}.$$

For more information about the bracket, we refer to [4]; there we show for example that there is a second boundary operator on \mathcal{HG} , and the bracket measures how far this boundary operator is from being a derivation.

If x and y belong to separating edges of G and H , then $(G \cdot H)_{xy}$ will not be connected, even if G and H are connected. Thus the bracket on \mathcal{HG} does not restrict to a bracket on \mathcal{G} . It does restrict to a bracket on the subcomplex of \mathcal{G} spanned by graphs with no separating edges, but that subcomplex is not quasi-isomorphic to \mathcal{G} . However, we will show that it does induce a well-defined bracket on \mathcal{G}/\mathcal{C} . The quotient \mathcal{G}/\mathcal{C} has as basis the cosets $G + \mathcal{C}$, where G is connected with no cut vertices. We define the bracket on basis elements by $[G + \mathcal{C}, H + \mathcal{C}] = [G, H] + \mathcal{C}$, where G and H are connected with no cut vertices. To see that this is well defined, we need the following lemma.

LEMMA 4.1 *If $G, H \neq 0$ are connected and have no cut vertices, then each term $(G \cdot H)_{xy}$ of $[G, H]$ is connected and has no cut vertices. If G or H has a cut vertex, then so does each term $(G \cdot H)_{xy}$.*

Proof. Recall that an orientated graph is zero if it has an edge-loop at any vertex. A graph without such loops is connected with no cut vertices if and only if there are at least two disjoint paths between every pair of vertices. Let v be the vertex of G adjacent to x and \bar{v} the vertex adjacent to \bar{x} ; similarly, let w, \bar{w} be the vertices of H adjacent to y, \bar{y} . Choose a path α in G from v to \bar{v} which does not contain x , and β a path in H from w to \bar{w} which does not contain y .

If a and b are two vertices of G , and one of the two disjoint paths between them contains $x\bar{x}$, then we can construct a second disjoint path in $(G \cdot H)_{xy}$ by replacing $x\bar{x}$ by β . Similarly, if a and b are in H , we can replace a path containing $y\bar{y}$ by α . If a is in G and b is in H , then disjoint paths can be constructed as follows. To make the first path, join a to the image of v and the (identical) image of w to b ; The second path is obtained by joining a to \bar{v} , then going across $\bar{x}\bar{y}$, then joining \bar{w} to b .

If a vertex is a cut vertex in G , its image is a cut vertex in $(G \cdot H)_{xy}$.

COROLLARY 4.2 *The bracket induces a well-defined bracket on \mathcal{G}/\mathcal{C} .*

Proof. The only subtlety here is that the bracket of two graphs with separating edges (and hence cut vertices) might not be connected. Let $C_1, C_2 \in \mathcal{C}$, and let \mathcal{HC} be the subspace of \mathcal{HG} spanned by graphs with cut vertices. Then $[G + C_1, H + C_2] = [G, H] + [G, C_2] + [C_1, H] + [C_1, C_2] \in [G, H] + \mathcal{HC}$, by the lemma. We then appeal to the natural isomorphism

$$(G + \mathcal{HC})/\mathcal{HC} \cong \mathcal{G}/\mathcal{C}$$

to identify $[G + C_1, H + C_2]$ with $[G, H]$ in \mathcal{G}/\mathcal{C} .

REMARK 4.3 The lemma shows that the bracket restricts to the subspace of \mathcal{G} spanned by graphs with no cut vertices. However, that subspace is not a subcomplex, as the boundary map does not restrict. This is the reason we are using the quotient complex \mathcal{G}/\mathcal{C} .

The cobracket $\mathcal{HG} \rightarrow \mathcal{HG} \otimes \mathcal{HG}$ is defined as follows. We say a pair $\{x, y\}$ of half-edges of a graph G is *separating* if the number of components of G_{xy} is greater than that of G . If G is connected, define

$$\theta(G) = \sum_{\{x,y\} \text{ separating}} A \otimes B + (-1)^a B \otimes A,$$

where $G_{xy} = A \cdot B$, and a is the number of vertices of A . This gives the coproduct on primitive elements, and extends to all elements in a standard way; see [4]. We have the following.

LEMMA 4.4 *Let $\{x, y\}$ be a separating pair of half-edges in an oriented graph G , with $G_{xy} = A \cdot B$. If G has a cut vertex, then at least one of A or B has a cut vertex. If G is connected, then both A and B are connected.*

Proof. The proof is straightforward.

Thus the cobracket induces a cobracket on $(\mathcal{G} + \mathcal{HC})/\mathcal{HC} \cong \mathcal{G}/\mathcal{C}$ defined on a basis element $G + \mathcal{C}$, where G is a connected graph with no cut vertices, by

$$\theta(G + \mathcal{C}) = \sum_{\{x,y\} \text{ separating}} (A + \mathcal{C}) \otimes (B + \mathcal{C}) + (-1)^a (B + \mathcal{C}) \otimes (A + \mathcal{C}).$$

We now check that the bracket and cobracket are compatible on \mathcal{G}/\mathcal{C} , making \mathcal{G}/\mathcal{C} into a Lie bialgebra.

PROPOSITION 4.5 *The bracket and cobracket satisfy*

$$\theta[G + \mathcal{C}, H + \mathcal{C}] + [\theta(G + \mathcal{C}), H + \mathcal{C}] + (-1)^g [G + \mathcal{C}, \theta(H + \mathcal{C})] = 0.$$

where g is the degree of G .

Proof. Let $G + \mathcal{C}$ and $H + \mathcal{C}$ be basis elements of \mathcal{G}/\mathcal{C} , that is, G and H are connected with no cut vertices. We compute

$$\begin{aligned} & \theta[G + \mathcal{C}, H + \mathcal{C}] + [\theta(G + \mathcal{C}), H + \mathcal{C}] + (-1)^g [G + \mathcal{C}, \theta(H + \mathcal{C})] \\ &= \theta[G, H] + [\theta(G), H] + (-1)^g [G, \theta(H)]. \end{aligned}$$

Because both G and H have no cut vertices, they also have no separating edges, so the last sum is zero by [4, Theorem 1].

5. Calculations

In this section we present our computations of the rational homology of $\mathcal{G}^{(n)}$ for $n \leq 7$, briefly describing the algorithm, but omitting the raw code. Details for a similar algorithm can be found in [6].

The program first enumerates all trivalent graphs with no cut vertices. There is only one such graph with fundamental group of rank 2, the *theta graph*: two vertices connected by three edges. If we have a list of all graphs with fundamental group of rank $n - 1$, we can obtain the list for rank n by applying one of the following two operations to all the graphs in every possible way. The first operation takes two distinct edges of the graph, subdivides them by adding a new vertex at the middle of each, and adds a new edge between the two new vertices. The second operation adds two new vertices in the interior of a single edge and connects the two new vertices by a new edge. It is not hard to see that if G has no cut vertices, then it can be obtained from a lower-rank graph which also has no cut vertices using one of these two operations.

The same graph will be listed several times. To eliminate the duplications, we transform each new graph into a *normal form*; two graphs in normal form are isomorphic if and only if they are identical. The graph is stored as the adjacency matrix a_{ij} for $i < j$; that is, $a_{ij} = k$ if vertex i is connected to vertex j by k edges. The normal form of the graph is the ordering of the vertices which yields the matrix latest in the lexicographic ordering. Permutations of the vertices are listed, and the matrices are compared. The number of permutations needed is reduced by distinguishing three types of vertices: those contained in a multiple edge, those contained in a triangle, and the rest. Only vertices of the same type need to be permuted among themselves.

Next, we enumerate all graphs of valence 3 or higher, with no cut vertices and with fundamental group of rank at most 7 by successively contracting edges of the trivalent graphs. Cut vertices may develop during this process: in this case the graph is discarded. Then we examine each graph to see whether it has any orientation-reversing automorphisms, and if so, discard it. A graph with an orientation-reversing automorphism is zero in the graph complex since such a graph is equal to minus itself and the base field is of characteristic zero.

Finally, we compute the matrix of the boundary map

$$\partial(G) = \sum_e G_e$$

by transforming each contracted graph G_e into normal form and comparing it with the list of lower-rank graphs. The output of the program is a sparse matrix; its rank was computed by simple Gaussian elimination in the case of the smaller matrices and by the software package SCILAB in the case of the largest ones.

Recall that $\mathcal{G}^{(n)}$ denotes the subcomplex of the graph complex spanned by rank- n graphs, and that $\mathcal{C}^{(n)}$ is the subcomplex spanned by graphs with cut vertices. We obtain the following rank- n quotient complexes $\mathcal{G}^{(n)}/\mathcal{C}^{(n)}$ for values of n less than 8.

$\mathcal{G}^{(3)}/\mathcal{C}^{(3)}$:

$$0 \rightarrow C_4 \xrightarrow{1} C_3 \xrightarrow{0} C_2 \rightarrow 0$$

$\mathcal{G}^{(4)}/\mathcal{C}^{(4)}$:

$$0 \rightarrow C_6 \xrightarrow{3} C_5 \xrightarrow{0} C_4 \xrightarrow{0} C_3 \xrightarrow{1} C_2 \rightarrow 0$$

$\mathcal{G}^{(5)}/\mathcal{C}^{(5)}$:

$$0 \rightarrow C_8 \xrightarrow{12} C_7 \xrightarrow{7} C_6 \xrightarrow{5} C_5 \xrightarrow{7} C_4 \xrightarrow{3} C_3 \xrightarrow{0} C_2 \rightarrow 0$$

$\mathcal{G}^{(6)}/\mathcal{C}^{(6)}$:

$$0 \rightarrow C_{10} \xrightarrow{52} C_9 \xrightarrow{76} C_8 \xrightarrow{101} C_7 \xrightarrow{116} C_6 \xrightarrow{61} C_5 \xrightarrow{11} C_4 \xrightarrow{1} C_3 \xrightarrow{1} C_2 \rightarrow 0$$

$\mathcal{G}^{(7)}/\mathcal{C}^{(7)}$:

$$0 \rightarrow C_{12} \xrightarrow{295} C_{11} \xrightarrow{828} C_{10} \xrightarrow{1560} C_9 \xrightarrow{1969} C_8 \xrightarrow{1393} C_7 \xrightarrow{540} C_6 \xrightarrow{138} C_5 \xrightarrow{35} C_4 \xrightarrow{6} C_3 \xrightarrow{0} C_2 \rightarrow 0.$$

The number printed under the chain group C_i is its dimension, the number printed above the arrow $\partial_i : C_i \rightarrow C_{i-1}$ is the rank of this linear map. Thus we have the following.

THEOREM 5.1 *The rational homology $H_i(\mathcal{G}^{(n)})$ of the commutative graph complex is zero for all $2 \leq n \leq 7$ and all i except for*

$$\begin{aligned} H_2(\mathcal{G}^{(2)}) &\cong \mathbb{Q} \\ H_4(\mathcal{G}^{(3)}) &\cong \mathbb{Q} \\ H_6(\mathcal{G}^{(4)}) &\cong \mathbb{Q} \\ H_8(\mathcal{G}^{(5)}) &\cong \mathbb{Q}^2 \\ H_{10}(\mathcal{G}^{(6)}) &\cong \mathbb{Q}^2 \quad H_7(\mathcal{G}^{(6)}) \cong \mathbb{Q} \\ H_{12}(\mathcal{G}^{(7)}) &\cong \mathbb{Q}^3 \quad H_9(\mathcal{G}^{(7)}) \cong \mathbb{Q}. \end{aligned}$$

For $2 \leq n \leq 5$, the computation takes only a few minutes. For $n = 6$, it took several hours of CPU time, and for $n = 7$, several thousand hours, even though the elimination of graphs with cut vertices reduces the size of the computation by about 30 per cent.

As Kontsevich realized, any metrized Lie algebra produces classes in trivalent graph homology, and so the abundance of top-dimensional homology is perhaps not surprising. Indeed, these trivalent classes correspond to finite type three manifold invariants. On the other hand, the presence of two codimension-3 classes is rather tantalizing.

The source code for the program is available in the source Folder available with the ‘source’ for the first version of this paper on [arXiv.org](https://arxiv.org). Please look at the readme file first. There is also a data Folder available in the same place.

Acknowledgements

The first author was partially supported by NSF grant DMS-0305012. The third author was partially supported by NSF grant DMS-0204185. The computations were done at the NIIFI Supercomputing Centre in Budapest, Hungary. We thank Craig Jensen for pointing out several typographical errors in an earlier draft.

References

1. D. Bar-Natan, S. Garoufalidis, L. Rozansky, and D. Thurston, The Aarhus integral of rational homology 3-spheres I: A highly nontrivial flat connection on S^3 , *Selecta Math. (N.S.)* **8** (2002), 315–319.
2. Mladen Bestvina and Mark Feighn, The topology at infinity of $\text{Out}(F_n)$, *Invent. Math.* **140** (2000), 651–692.
3. James Conant, Fusion and fission in graph complexes, *Pac. J. Math.* **209**, 2 (2003), 219–230.
4. James Conant and Karen Vogtmann, Infinitesimal operations on complexes of graphs, *Math. Ann.* **327** (2003), 545–573.
5. James Conant and Karen Vogtmann, On a theorem of Kontsevich, *Algebr. Geom. Topol.* **3** (2003), 1167–1224.
6. Ferenc Gerlits, *Invariants in chain complexes of graphs*, Ph.D. Thesis, Cornell, 2002.
7. Wee Liang Gan, *On a theorem of Conant–Vogtmann*, preprint 2004, math.QA/0404173.
8. Maxim Kontsevich, *Formal (non)commutative symplectic geometry*, The Gelfand Mathematical Seminars, 1990–1992, Birkhäuser, Cambridge, Boston, MA, 1993, 173–187.
9. Maxim Kontsevich, Feynman diagrams and low-dimensional topology, *First European Congress of Mathematics*, Vol. II, 1994, 97–121, Progr. Math. 120, Birkhäuser, Basel.
10. Greg Kuperberg and Dylan P. Thurston, *Perturbative 3-manifold invariants by cut-and-paste topology*, preprint 1999, UC Davis Math 1999-36, math.GT/9912167.
11. T. Q. T. Le, J. Murakami, and T. Ohtsuki, On a universal perturbative invariant of 3-manifolds, *Topology* **37** (1998), 539–574.
12. Karen Vogtmann, Automorphisms of free groups and Outer Space, *Geometriae Dedicata* **94** (2002), 1–31.