Available online at www.sciencedirect.com
ScienceDirect
Advances in Mathematics 225 (2010) 2117–2161
ADVANCES IN Mathematics
www.elsevier.com/locate/aim

Finite type invariants and fatgraphs

Jørgen Ellegaard Andersen [a], Alex James Bene [b,*], Jean-Baptiste Meilhan [d], R.C. Penner [c,a]

[a] *Department of Mathematics, Center for the Topology and Quantization of Moduli Spaces, Aarhus University, DK-8000 Aarhus C, Denmark*
[b] *Departments of Mathematics, University of Southern California, Los Angeles, CA 90089, USA*
[c] *Departments of Mathematics and Physics/Astronomy, University of Southern California, Los Angeles, CA 90089, USA*
[d] *Institut Fourier, Université Grenoble 1, 38402 St Martin d'Hères, France*

Received 2 August 2009; accepted 4 April 2010

Available online 5 August 2010

Communicated by David Ben-Zvi

Abstract

We define an invariant $\nabla_G(M)$ of pairs M, G, where M is a 3-manifold obtained by surgery on some framed link in the cylinder $\Sigma \times I$, Σ is a connected surface with at least one boundary component, and G is a fatgraph spine of Σ. In effect, ∇_G is the composition with the ι_n maps of Le–Murakami–Ohtsuki of the link invariant of Andersen–Mattes–Reshetikhin computed relative to choices determined by the fatgraph G; this provides a basic connection between 2d geometry and 3d quantum topology. For each fixed G, this invariant is shown to be universal for homology cylinders, i.e., ∇_G establishes an isomorphism from an appropriate vector space $\overline{\mathcal{H}}$ of homology cylinders to a certain algebra of Jacobi diagrams. Via composition $\nabla_{G'} \circ \nabla_G^{-1}$ for any pair of fatgraph spines G, G' of Σ, we derive a representation of the Ptolemy groupoid, i.e., the combinatorial model for the fundamental path groupoid of Teichmüller space, as a group of automorphisms of this algebra. The space $\overline{\mathcal{H}}$ comes equipped with a geometrically natural product induced by stacking cylinders on top of one another and furthermore supports related operations which arise by gluing a homology handlebody to one end of a cylinder or to another homology handlebody. We compute how ∇_G interacts with all three operations explicitly in terms of natural products on Jacobi diagrams and certain diagrammatic constants. Our main result gives an explicit extension of the LMO invariant of 3-manifolds to the Ptolemy groupoid in terms of these operations, and this groupoid extension nearly fits the paradigm of a TQFT. We finally re-derive the Morita–Penner cocycle representing the first Johnson homomorphism using a variant/generalization of ∇_G.

* Corresponding author.
E-mail addresses: andersen@imf.au.dk (J.E. Andersen), bene@usc.edu (A.J. Bene), jean-baptiste.meilhan@ujf-grenoble.fr (J.-B. Meilhan), rpenner@math.usc.edu (R.C. Penner).

0001-8708/$ – see front matter
doi:10.1016/j.aim.2010.04.008

Keywords: Finite type invariants; Fatgraphs; Mapping class groups; Teichmüller spaces; Homology cylinders

1. Introduction

In [21], Le, Murakami and Ohtsuki constructed an invariant $Z^{LMO}(M)$ of a closed oriented 3-manifold M from the Kontsevich integral Z (see Section 2.2) of a framed link with k components, where Z takes values in the space $\mathcal{A}(\circlearrowleft^k)$ of Jacobi diagrams with core $\circlearrowleft^k$, a collection of k oriented circles (see Section 2.1.1). The Kontsevich integral Z is universal among rational-valued Vassiliev invariants, i.e., any other factors through it. The LMO invariant $Z^{LMO}(M) \in \mathcal{A}(\emptyset)$ takes values in Jacobi diagrams with empty core and arises as a suitably normalized post-composition of Z with mappings

$$\iota_n : \mathcal{A}\big(\circlearrowleft^k\big) \to \mathcal{A}(\emptyset),$$

which are of key importance for LMO and effectively "replace circles by sums of trees" (see Section 3.1.1). The LMO invariant is universal among rational-valued finite type invariants of integral and of rational homology spheres.

In [2], Mattes, Reshetikhin and the first-named author defined a universal Vassiliev invariant of links (see Section 2.5 for a partial review) in the product manifold $\Sigma \times I$, where $\Sigma = \Sigma_{g,n}$ is a fixed oriented surface of genus $g \geqslant 0$ with $n \geqslant 1$ boundary components and I is the closed unit interval, which generalizes the Kontsevich integral. Actually, the determination of this AMR invariant depends on a certain decomposition of the surface Σ into polygons.

In [30–32], the last-named author described an ideal cell decomposition of the decorated Teichmüller space of a bordered surface in terms of marked fatgraphs G embedded in Σ (see Section 2.3 for the definitions) and introduced the Ptolemy groupoid $\mathfrak{Pt}(\Sigma)$ and its canonical presentation in terms of Whitehead moves (see Section 5.2 for both the moves and the presentation). A key point is that the natural quotient of $\mathfrak{Pt}(\Sigma)$ contains the mapping class group $MC(\Sigma)$ of Σ as the stabilizer of any object. A more speculative point (discussed further in Section 6) is that the Whitehead moves which generate $\mathfrak{Pt}(\Sigma)$ may themselves be interpreted as triangulated cobordisms of triangulated surfaces.

In fact, the specification of a marked fatgraph G in Σ suffices to determine a polygonal decomposition (see Section 2.4 for this construction) as required for the definition of the AMR invariant. This is a basic connection between decorated Teichmüller theory and finite type invariants which we exploit here.

Indeed, we define an invariant ∇_G (see Section 3.1 for the definition and Theorem 3.1 for its invariance) taking values in the space $\mathcal{A}_h$ of h-labeled Jacobi diagrams without strut components (see Section 2.1.1 for the definitions), where $h = 2g + n - 1$ is the rank of the first homology group of Σ if $\Sigma = \Sigma_{g,n}$ has genus g and n boundary components. Specifically, our invariant is defined for any "cobordism" M, i.e., $\nabla_G(M)$ is defined for any 3-manifold $M = (\Sigma \times I)_L$ arising from Dehn surgery on a framed link $L \subset \Sigma \times I$ and for any marked fatgraph G in Σ. In fact, the fatgraph G determines not only the polygonal decomposition necessary for an AMR invariant but also other choices which are required for our new invariant (see Section 2.6 for these other choices called systems of "latches" and "linking pairs").

The invariant ∇_G is defined in analogy to Z^{LMO} in the sense that it arises as a suitably normalized post-composition of the AMR invariant determined by G with ι_n, so the AMR invariant (actually, a weakened forgetful version of it) plays for us the role of the Kontsevich integral in LMO. We show (see Theorem 3.2) that ∇_G is universal for so-called "homology cylinders", which arise for surgeries along a particular class of links called claspers (see Section 3.2 for the definitions of homology cylinders and claspers).

Since ∇_G is universal for homology cylinders, it induces an isomorphism

$$\nabla_G : \overline{\mathcal{H}}_\Sigma \to \mathcal{A}_h,$$

where $\overline{\mathcal{H}}_\Sigma$ is a quotient of the vector space freely generated by homology cylinders over Σ (see Section 3.2.1 for the precise definition of $\overline{\mathcal{H}}_\Sigma$). It is this manifestation of universality that has useful consequences for the Ptolemy groupoid $\mathfrak{Pt}(\Sigma)$ since given two marked fatgraphs G and G' in Σ, there is the composition

$$\nabla_{G'} \circ \nabla_G^{-1} : \mathcal{A}_h \to \mathcal{A}_h.$$

For essentially formal reasons, this turns out to give a representation

$$\xi : \mathfrak{Pt}(\Sigma_{g,n}) \to \mathrm{Aut}(\mathcal{A}_h)$$

of the Ptolemy groupoid in the algebra automorphism group of $\mathcal{A}_h$.

There are several well-known and geometrically natural operations on $\overline{\mathcal{H}}$. Firstly, there is the "stacking" induced by gluing homology cylinders top-to-bottom. Secondly, given a homology cylinder over the once-bordered surface $\Sigma_{g,1}$ and a genus g homology handlebody (see Section 4.1.1 for the definition), we can take their "shelling product" by identifying the boundary of the latter with the bottom of the former. Thirdly and finally, we can glue two homology handlebodies along their boundaries to get a closed 3-manifold in the spirit of Heegaard decompositions which is called the "pairing" between the homology handlebodies (see Section 4.1.2 for details on all three operations).

In Section 4.3 we explicitly define three algebraic maps

$$\bullet : \mathcal{A}_{2g} \times \mathcal{A}_{2g} \to \mathcal{A}_{2g}, \qquad \star : \mathcal{A}_{2g} \times \mathcal{A}_g \to \mathcal{A}_g, \quad \text{and} \quad \langle , \rangle : \mathcal{A}_g \times \mathcal{A}_g \to \mathcal{A}(\emptyset),$$

which respectively correspond (under conjugation with a normalized version of ∇_G explained in Section 3.3) to the stacking product, the shelling product, and the pairing (as proved in Theorem 4.4). Furthermore, these operations are computed in terms of a basic "concatenation product" $\odot$ (see Section 4.2) with three particular tangles T_g, R_g, S_g (see Section 4.3 and Fig. 13 for the definitions of these tangles) respectively corresponding to the three operations; this gives a purely diagrammatic interpretation and scheme of computation for each operation.

Our penultimate result relies on a groupoid representation

$$\rho : \mathfrak{Pt}(\Sigma_{g,1}) \to \mathcal{A}_{2g},$$

defined by combining our invariant with a representation from [1], to extend the LMO invariant of integral homology spheres to the Ptolemy groupoid in the following sense. Let f be an element of the Torelli group of $\Sigma_{g,1}$ and let

$$G \xrightarrow{W_1} G_1 \xrightarrow{W_2} \cdots \xrightarrow{W_k} G_k = f(G)$$

be a sequence of Whitehead moves representing f in the sense of decorated Teichmüller theory (see Section 5.2). Our result then states that the LMO invariant of the integral homology 3-sphere S^3_f obtained by the Heegaard construction via f is given by

$$Z^{LMO}\big(S^3_f\big) = \big\langle v_0, \big(\rho(W_1) \bullet \rho(W_2) \bullet \cdots \bullet \rho(W_k)\big) \star v_0\big\rangle \in \mathcal{A}(\emptyset),$$

where v_0 is an explicit diagrammatic constant (see Theorem 5.4 for the precise statement) and the operations are fully determined diagrammatically as discussed before. This formalism shows the sense in which the LMO invariant extends to the Ptolemy groupoid as a kind of weakened version of TQFT; whereas the Ptolemy groupoid has not made contact with the LMO invariant previously, similar TQFT phenomena and remarks are reported in [27,9].

Finally (in Section 5.6), we use our invariant (actually, a variation/generalization $\nabla_G^{I_g}$ of ∇_G, which takes values in Jacobi diagrams with core $2g$ intervals and depends upon a "general system of latches" I_g), in order to associate to the Whitehead move $G \xrightarrow{W} G'$ the quotient $\mathcal{J}(W) = \nabla_G^{I_g}(\Sigma_{g,1} \times I)/\nabla_{G'}^{I_g}(\Sigma_{g,1} \times I)$ and define a representation

$$\mathcal{J}^{\mathsf{Y}} : \mathfrak{Pt}(\Sigma_{g,1}) \to \Lambda^3 H_1(\Sigma_{g,1}; \mathbb{Q}),$$

which coincides with that defined by Morita and Penner [26] to give a canonical cocycle extension of the first Johnson homomorphism [16]; see [7] for analogous cocycles extending all of the higher Johnson homomorphisms. It thus seems reasonable to expect that higher-order calculations should provide a corresponding formula for the second Johnson homomorphism and, in light of [25], also for the Casson invariant. In fact, one motivation for the present work was to investigate whether the known extensions to the Ptolemy groupoid of the Johnson homomorphisms [26,1,7] might be special cases of a more general extension of Z^{LMO}, cf. [12,14,23].

We have learned here that Z^{LMO} indeed extends to the Ptolemy groupoid, and in particular have derived an explicit purely diagrammatic extension of Z^{LMO} which is "nearly a TQFT", but whose formulas are not particularly simple or natural largely owing to their dependence upon certain combinatorial algorithms from [1].

On the other hand by a related construction (in Section 5.6.1), we have in the context of finite type invariants derived an elegant and natural Ptolemy groupoid representation which may give a simpler extension of Z^{LMO}. We expect that there is a precursor for this in the early days of development of [26,1,7], where explicit unpleasant formulas were ultimately replaced by simpler and more conceptual ones; see Section 6 for a further discussion.

Standard Notation. We shall fix a compact connected and oriented surface $\Sigma = \Sigma_{g,n}$ of genus $g \geqslant 0$ with $n \geqslant 1$ boundary components, fix a basepoint $p \in \partial\Sigma$ and let $h := 2g + n - 1$ denote the rank of $H_1(\Sigma; \mathbb{Z})$. We shall often write $1_\Sigma = \Sigma \times I$, where $I = [0, 1]$.

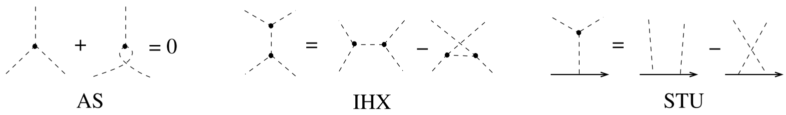

Fig. 1. The relations AS, IHX and STU.

2. Definitions

2.1. Jacobi diagrams

We first recall the spaces of diagrams in which the Kontsevich, LMO and our new invariants take values.

2.1.1. Definitions

A *Jacobi diagram* is a finite graph with only univalent and trivalent vertices, or a so-called "uni-trivalent" graph, such that each trivalent vertex is equipped with a cyclic ordering of its three incident half-edges. In other words, a Jacobi diagram is exactly a uni-trivalent "fatgraph" as discussed separately in Section 2.3. The *Jacobi degree* or simply *J-degree* of a Jacobi diagram is half its number of vertices.

Let $S = \{s_1, \ldots, s_m\}$ be some finite linearly ordered set and let X be a 1-manifold, where we tacitly assume that X is compact and oriented and that its components come equipped with a linear ordering. A Jacobi diagram G *lies on* (X, S) if the set of univalent vertices of G partitions into two disjoint sets, where elements of one of these sets are labeled by elements of S, and elements of the other are disjointly embedded in X; X is called the *core* of the Jacobi diagram. As usual [3,29] for figures, we use bold lines to depict the 1-manifold X and dashed ones to depict the Jacobi diagram (though fatgraphs will sometimes also be depicted with bold lines), and we take the cyclic ordering at a vertex given by the counter-clockwise orientation in the plane of the figure, which is used to determine the "blackboard framing".

Let $\mathcal{A}(X, S)$ denote the $\mathbb{Q}$-vector space generated by Jacobi diagrams on (X, S), subject to the AS, IHX and STU relations depicted in Fig. 1. Consider the respective vector subspaces $\mathcal{A}_k(X, S)$ and $\mathcal{A}_{\leqslant k}(X, S)$ generated by Jacobi diagrams lying on (X, S) of J-degree k and $\leqslant k$, with respective projections of $x \in \mathcal{A}(X, S)$ denoted x_k and $x_{\leqslant k}$. Abusing notation slightly, let $\mathcal{A}(X, S)$ furthermore denote the J-degree completion of $\mathcal{A}(X, S)$ with its analogous projections to $\mathcal{A}_k(X, S)$ and $\mathcal{A}_{\leqslant k}(X, S)$. The empty diagram in $\mathcal{A}(X, S)$ is often denoted simply 1.

We shall primarily be interested in certain specializations of this vector space:

- When $S = \emptyset$, we write simply $\mathcal{A}(X) = \mathcal{A}(X, \emptyset)$. If X is the disjoint union of m copies of S^1, respectively, m copies of the unit interval, then $\mathcal{A}(X)$ is also respectively denoted by $\mathcal{A}(\circlearrowleft^m)$ and $\mathcal{A}(\uparrow^m)$. There is an obvious surjective "closing map"

$$\pi : \mathcal{A}(\uparrow^m) \to \mathcal{A}(\circlearrowleft^m)$$

which identifies to a distinct point the boundary of each component of X.
- When $X = \emptyset$, we write simply $\mathcal{B}(S) = \mathcal{A}(\emptyset, S)$, called the vector space of *S-colored Jacobi diagrams*, and when $S = \{1, \ldots, m\}$, we write $\mathcal{B}(m) = \mathcal{B}(S)$, called the space of *m-colored Jacobi diagrams*.

In fact, $\mathcal{A}(X,S)$ has the structure of a Hopf algebra provided $X=\emptyset$, $\circlearrowleft$, or $\uparrow^m$, cf. the next section.

When $S=\{s_1,\dots,s_m\}$ is a linearly ordered set with cardinality m, there is a standard [3] graded isomorphism

$$\chi_S:\mathcal{B}(S)\to\mathcal{A}(\uparrow^m),$$

called the Poincaré–Birkhoff–Witt isomorphism, which maps a diagram to the average of all possible combinatorially distinct ways of attaching its s_i-colored vertices to the ith interval, for $i=1,\dots,m$. When $S=\{1,\dots,m\}$, we simply write $\chi=\chi_S$; more generally, given a 1-manifold X with a submanifold $X'\subset X$ which is isomorphic to and identified with $\uparrow^m$, we have the isomorphism

$$\chi_{X',S}:\mathcal{A}(X-X',S)\to\mathcal{A}(X),$$

which arises by applying χ_S only to the S-labeled vertices.

The *internal degree* or *i-degree* of a Jacobi diagram is its number of trivalent vertices. We call a connected Jacobi diagram of i-degree zero a *strut*, and we denote by $\mathcal{B}^{\mathsf{Y}}(m)$ the vector space generated by m-colored Jacobi diagrams without strut components modulo the AS and the IHX relations. As these two relations (unlike STU) are homogeneous with respect to the internal degree, $\mathcal{B}^{\mathsf{Y}}(m)$ is graded by the i-degree. The i-degree completion is also denoted $\mathcal{B}^{\mathsf{Y}}(m)$ and is canonically isomorphic to the J-degree completion.

In the rest of this paper, we shall use the simplified notation $\mathcal{A}_m=\mathcal{B}^{\mathsf{Y}}(m)$.

2.1.2. Operations on Jacobi diagrams

There are several basic operations [3] on Jacobi diagrams as follows:

First of all, disjoint union of 1-manifolds X_1 and X_2 gives a tensor product

$$\otimes:\mathcal{A}(X_1)\times\mathcal{A}(X_2)\to\mathcal{A}(X_1\sqcup X_2),$$

where the linear ordering on the components of $X_1\sqcup X_2$ is the lexicographic one with components of X_1 preceding those of X_2. Secondly, if $V_i\subseteq\partial X_i$ for $i=1,2$, and V_1 is identified with the reversal of V_2 as linearly ordered sets of points to form a new 1-manifold X from X_1 and X_2, then the *stacking product*

$$\cdot:\mathcal{A}(X_1)\times\mathcal{A}(X_2)\to\mathcal{A}(X)$$

arises by gluing together pairs of identified points and combining Jacobi diagrams in the natural way.

Suppose that $Y\subseteq X$ is a connected component of a 1-manifold X. $Y^{(n)}$ denotes the union of n ordered parallel copies of Y. The *comultiplication map* $\Delta_Y:\mathcal{A}(X)\to\mathcal{A}(Y^{(2)}\cup X-Y)$ is defined as follows. Given a diagram $D\in\mathcal{A}(X)$ with c univalent vertices on Y, replace Y by $Y^{(2)}$ and take the sum of all 2^c possible ways of distributing these c vertices to the components of $Y^{(2)}$. More generally, we can recursively define maps

$$\Delta_Y^{(n)}:\mathcal{A}(X)\to\mathcal{A}\big(Y^{(n)}\cup X-Y\big)$$

by $\Delta_Y^{(n)}=\Delta_{Y_1^{(n-1)}}\circ\Delta_Y^{(n-1)}$, where $Y_1^{(n-1)}$ denotes the first copy of Y in $Y^{(n-1)}$.

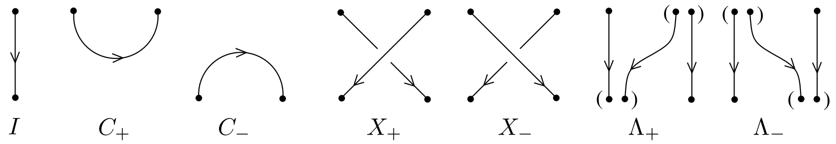

Fig. 2. The elementary q-tangles I, $C_\pm$, $X_\pm$ and $\Lambda_\pm$.

If $Y \subseteq X$ is a union of components, let $\overline{Y}$ denote the result of reversing the orientation on Y. The *antipode map*

$$S_Y : \mathcal{A}(X) \to \mathcal{A}(\overline{Y} \cup X - Y)$$

is defined by $S_Y(D) = (-1)^c \overline{D}$, where the diagram $D \in \mathcal{A}(X)$ contains c univalent vertices attached to Y, and $\overline{D}$ arises from D by reversing the orientation of Y.

Finally, let X be a 1-manifold with distinguished component C. Let $D \in \mathcal{A}(X)$ and $D' \in \mathcal{A}(S^1)$ be two Jacobi diagrams. Remove a small arc from each of C and S^1 which does not contain any vertex of D and D', and connect the boundary points of the removed arcs such that the orientations are compatible and such that the result is an element of $\mathcal{A}(X)$. The result is called the *connected sum* with D' on the component C of D. The fact that it does not depend on the choice of the removed arc follows from the STU relation, see [3].

2.2. The Kontsevich integral of framed tangles

Let M be a compact connected oriented 3-manifold whose boundary is endowed with an identification to the boundary of the standard cube $C := [0,1]^3$, and let X be a 1-manifold possibly with boundary. A *tangle with core X in M* is a proper embedding of X in M such that all boundary points of X lie on the segments $[0,1] \times \frac{1}{2}$ in the upper and lower squares $[0,1]^2 \times \{1\}$ and $[0,1]^2 \times \{0\}$ of C. We shall identify such an embedding with the (isotopy class relative to the boundary) of its image. A *framed tangle* is a tangle together with a non-vanishing normal vector field. A *q-tangle* is a framed tangle enhanced with a “bracketing”, i.e., a consistent collection of parentheses on each of the naturally linearly ordered sets of boundary points in the segments $[0,1] \times \{\frac{1}{2}\} \times \{\varepsilon\}$; $\varepsilon = 0, 1$.

We define two operations on q-tangles in C as follows. The tensor product $T \otimes T'$ of two q-tangles T and T' is obtained by horizontal juxtaposition and natural bracketing, with T to the left of T' (and reparameterization of the ambient cube). If the upper end of T coincides with the lower end of T', i.e., they coincide as bracketed sets of dots, then the composition $T \cdot T'$ is obtained by stacking T' on top of T (and reparameterizing the ambient cube).

A fundamental fact [29] is that any q-tangle in C can be (non-uniquely) decomposed as a composition of tensor products of (oriented) copies of the elementary q-tangles I, $X_\pm$, $C_\pm$ and $\Lambda_\pm$ of Fig. 2 together with those obtained by orientation-reversal on certain components.

The *framed Kontsevich integral* $Z(T)$ of a q-tangle T with core X in the standard cube C lies in the space $\mathcal{A}(X)$ of Jacobi diagrams [3,29]. Insofar as $Z(T \cdot T') = Z(T) \cdot Z(T')$ and $Z(T \otimes T') = Z(T) \otimes Z(T')$, for any two tangles T, T', it is enough to determine Z on any tangle by specifying its values on the elementary q-tangles of Fig. 2 by the fundamental fact. We set $Z(I) = 1 \in \mathcal{A}(\uparrow)$, and

$$Z(C_\pm) = \sqrt{\nu}, \tag{1}$$

where $\nu \in \mathcal{A}(\circlearrowleft) \simeq \mathcal{A}(\uparrow)$ is the Kontsevich integral of the 0-framed unknot (computed in [5]).

Recall that ν is invariant under the antipode map and that projecting away the non-strut components of $\chi^{-1}(\nu)$ produces zero.

Define

$Z(X_\pm) = \exp\Big(\frac{\pm 1}{2}$ 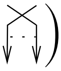 $\Big) = 1 + \sum_{k=1}^{\infty} \frac{(\pm 1)^k}{2^k k!}\Big($  $\Big)^k,$ (2)

where the kth power on the right-hand side denotes the diagram with k parallel dashed chords and set

$$Z(\Lambda_\pm) = \Phi^{\pm 1}, \tag{3}$$

where $\Phi \in \mathcal{A}(\uparrow^3)$ is the choice of an associator (see for example [29, Appendix D]).

While there are many associators that one may choose to define the Kontsevich integral, we shall restrict our choice to an even associator (see [20, Section 3] for a definition), which necessarily satisfies

$$Z\big(r(T)\big) = r\big(Z(T)\big), \tag{4}$$

for any q-tangle T and for any mirror reflection r of its planar projection with respect to any horizontal or vertical line [20]; moreover, if $T_i^{(k)}$ is obtained by taking k parallel copies of the ith component of a q-tangle T, then we have

$$Z\big(T_i^{(k)}\big) = \Delta_i^{(k)}\big(Z(T)\big). \tag{5}$$

2.3. Fatgraphs

A *fatgraph* is a finite graph endowed with a "fattening", i.e., a cyclic ordering on each set of half-edges incident on a common vertex. When depicting a fatgraph in a figure, the fattening is given by the counter-clockwise orientation in the plane of the figure. A fatgraph G determines a corresponding "skinny surface" with boundary in the natural way, where polygons of $2k$ sides corresponding to k-valent vertices of G have alternating bounding arcs identified in pairs as determined by the edges of G. We shall be primarily concerned with the case where such graphs are connected and uni-trivalent with only one univalent vertex, and by a slight abuse of terminology, we shall call such a fatgraph a *bordered fatgraph*. The edge incident on the uni-valent vertex of a bordered fatgraph is called the *tail*.

Suppose that $\mathbf{e}$ is an oriented edge that points towards the vertex v of G. There is a succeeding oriented edge $\mathbf{e}'$ gotten by taking the oriented edge pointing away from v whose initial half edge follows the terminal half edge of $\mathbf{e}$. A sequence of iterated successors gives an ordered collection of oriented edges starting from any oriented edge called a *boundary cycle* of G, which we take to be cyclically ordered and evidently corresponds to a boundary component of the associated skinny surface. We shall call any subsequence of a boundary cycle of G a *sector*. By a *once bordered fatgraph*, we mean a bordered fatgraph with only one boundary cycle, which canonically begins from the tail.

Thus, the oriented edges of any once bordered fatgraph G come in a natural linear ordering, namely, in the order of appearance in the boundary cycle starting from the tail. For a connected bordered fatgraph G, we can also linearly order the oriented edges by defining the *total boundary cycle* as follows. Let the total boundary cycle begin at the tail and continue until it returns again to the tail. If every oriented edge has not yet been traversed, then there is a first oriented edge $\mathbf{e}$ in this sequence such that the oppositely oriented edge $\bar{\mathbf{e}}$ has not yet been traversed by connectivity. We then extend the total boundary cycle by beginning again at $\bar{\mathbf{e}}$ and continuing as before until the boundary cycle containing $\bar{\mathbf{e}}$ has been fully traversed. At this stage, if there is again a first edge $\mathbf{e_1}$ in the total boundary cycle whose opposite orientation $\bar{\mathbf{e}}_\mathbf{1}$ has not yet been traversed, we continue the total boundary cycle at $\bar{\mathbf{e}}_\mathbf{1}$. By iterating this procedure, we eventually traverse every oriented edge of G exactly once. According to our conventions for figures, the total boundary cycle is oriented with G on its left.

Finally, a *marking* of a bordered fatgraph G in a surface $\Sigma = \Sigma_{g,n}$ of genus g with $n > 0$ boundary components with basepoint p in its boundary, is a homotopy class of embeddings $G \hookrightarrow \Sigma$ such that the tail of G maps to a point $q \neq p$ on the same component of the boundary of $\Sigma_{g,n}$ as p and the complement $\Sigma - G$ consists of a disc (corresponding to the boundary component containing p) and $n - 1$ annuli (corresponding to the remaining boundary components). The relative version [32] of decorated Teichmüller theory [30] shows that the natural space of all marked fatgraphs in a fixed bordered surface is identified with a trivial bundle over its Teichmüller space.

2.4. The polygonal decomposition associated to a fatgraph

By a *bigon*, *square* or *hexagon* in a surface Σ with boundary, we mean a (topologically) embedded closed disc $D^2 \hookrightarrow \Sigma$ such that the intersection $D^2 \cap \partial\Sigma$ is the union of one, two or three disjoint closed intervals, respectively, called the *bounding edges*; the closures of the components of the remainder of ∂D^2 are called the *cutting edges*.

Given a marked bordered fatgraph $G \hookrightarrow \Sigma$, its corresponding skinny surface is naturally diffeomorphic to Σ itself thus providing a polygonal decomposition

$$\Sigma = B \cup \left(\bigcup_i S_i\right) \cup \left(\bigcup_j H_j\right),$$

where each trivalent vertex corresponds to a hexagon H_j, each non-tail edge corresponds to a square S_i, and the tail corresponds to a bigon B, such that the intersection of any two of these components consists of a (possibly empty) union of cutting edges. We refer to any $S_i \times I$ or $B \times I$ as a *box* and to the box $B \times I$ associated to the tail of G as the *preferred box*. The faces of the boxes corresponding to cutting edges are called the *cutting faces*. This decomposition of Σ is the *polygonal decomposition associated to the fatgraph G* marking in Σ and is denoted P_G.

Such a decomposition P_G of Σ into 2-, 4-, and 6-gons, together with a specification of one bounding edge for each hexagon, provides sufficient data to define the AMR invariant of [2], which is discussed in the next section. We call the specified bounding edge of each hexagon (as well as the corresponding sector of G) its *forbidden sector*. One can check that for any choice of

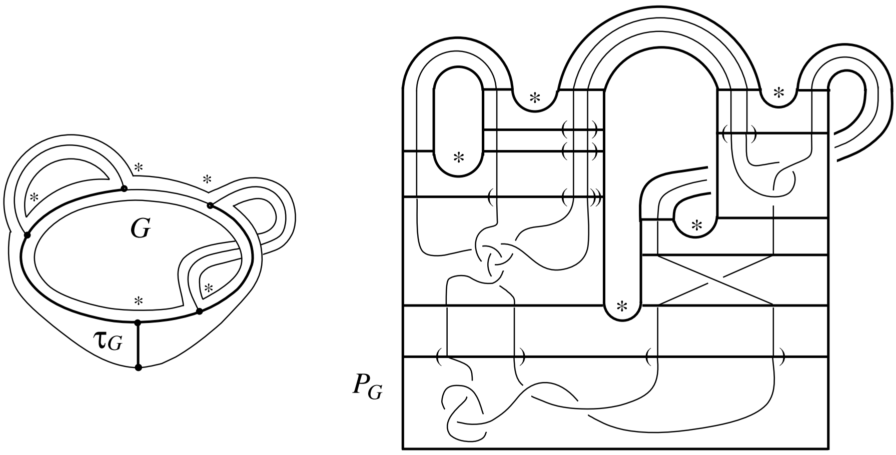

Fig. 3. A fatgraph G marked in the twice-punctured torus $\Sigma_{1,2}$, the maximal tree τ_G, and a knot in admissible position with respect to the polygonal decomposition P_G.

forbidden sectors for P_G, any framed link L in 1_Σ can be isotoped in 1_Σ and endowed with a bracketing of its intersection with the cutting faces so that:

- For each square S_i, the tangle $L_i^S := L \cap (S_i \times I)$ is a q-tangle in the cube $S_i \times I \cong C$ as in Section 2.2. Similarly, the tangle $L^B := L \cap (B \times I)$ is a q-tangle in $B \times I$.
- For each hexagon H_j, the tangle $L_j^H := L \cap (H_j \times I)$ is a "trivial" q-tangle in the sense that: there are no crossings of strands of L_j^H; no strand of L_j^H connects the two edges of ∂H_j adjacent to the forbidden sector; the bracketing of the intersection of L with the cutting face opposite the forbidden sector is the concatenation of the bracketings for the other two cutting faces in the natural way, cf. Fig. 3.
- Pairs of bracketings corresponding to the two sides of a single cutting face must coincide (as follows from their definition).

If a link satisfies these conditions, then we say that it is in *admissible position* with respect to the polygonal decomposition P_G associated to the marked fatgraph G in Σ. An example is given in Fig. 3, where we have labeled each forbidden sector by $*$.

The top and bottom of the various tangles L_i^S, L^B and L_j^H can be easily determined from the fatgraph G as follows. Each box has a natural upward direction induced by the first corresponding oriented edge traversed by the total boundary cycle of G, and the bottom face of each hexagon is the one opposite the forbidden sector. This convention, however, will not be relevant in our construction since we will make use of an even associator.

In fact, a marked fatgraph G in a surface Σ not only determines the required polygonal decomposition P_G of Σ as already discussed, it furthermore determines a collection of forbidden sectors as follows.

By the *greedy algorithm* of [1], there is a canonical maximal tree τ_G in G built by traversing the total boundary cycle of G starting from the tail and "greedily" adding every traversed edge to τ_G provided the resulting graph is simply connected. See Fig. 3. Note that during this process, the corresponding subset of τ_G is always a connected tree, and that the tail and all vertices of G

are included in τ_G. See [1, Section 3] for a detailed exposition of the greedy algorithm as well as its other manifestations and applications.

Given a bordered fatgraph G, its *generators* are the edges in the complement $X_G = G - \tau_G$ of the maximal tree τ_G. Note that there is a natural linear ordering on the set of generators, and each generator comes equipped with an orientation, where the ordering and orientation are determined as the first encountered during the traversal of the total boundary cycle.

By general principles about maximal trees, each vertex v of G is connected to the tail by a unique embedded path in τ_G, and this path contains a unique edge of G incident to v. As each hexagon of the decomposition of Σ corresponds to a vertex of G, we may define the forbidden sector of a hexagon to be the one opposite the edge contained in the path initiating from the corresponding vertex. See Fig. 3.

Lemma 2.1. *For any marked fatgraph G in Σ, the specified forbidden sectors and the corresponding polygonal decomposition P_G have the property that any link L in $\Sigma \times I$ can be isotoped so that it intersects each box except the preferred one in a trivial q-tangle.*

Proof. For the purposes of this proof, we distinguish between the T-boxes, coming from the edges of τ_G, and the G-boxes, coming from the edges of X_G. To begin the isotopy, for any G-box S_x containing a non-trivial tangle L_x^S, isotope L_x^S out of S_x in either direction through the adjacent hexagon and then into an adjacent T-box in a way which avoids producing arcs parallel to the forbidden sector of the hexagon. This results in a link which is trivial in all G-boxes. Next, for any T-box S_t containing a non-trivial tangle L_t^S, similarly isotope L_t^S into a neighboring T-box which is closer to the tail via the path in the maximal tree τ_G. Note that such an isotopy can be performed by our choice of forbidden sectors. Repeated application of this last step results in a link which is trivial in all boxes except the preferred one. □

2.5. The AMR invariant

Andersen, Mattes and Reshetikhin [2] defined a universal Vassiliev invariant of links in $1_\Sigma = \Sigma \times I$, for Σ a surface with boundary, which generalizes the Kontsevich integral; we shall only require a weak version of their more general construction in this paper. These invariants depend on the choice of a polygonal decomposition of the surface Σ together with other essentially combinatorial choices in order to decompose the link into suitable sub-links (as for the Kontsevich integral), and these choices (and more) are provided by a marked fatgraph G in Σ as discussed in the previous section.

Let $\mathcal{L}(\Sigma; m)$ be the set of isotopy classes of oriented framed m-component links in the thickened surface 1_Σ. Fixing an even associator $\Phi \in \mathcal{A}(\uparrow^3)$ for the Kontsevich integral Z once and for all and choosing a fatgraph G marking in Σ, we define the *AMR invariant*

$$V_G : \mathcal{L}(\Sigma; m) \to \mathcal{A}(\circlearrowleft^m)$$

as follows. Given a link $L \in \mathcal{L}(\Sigma; m)$ in admissible position, we apply Z to each q-tangle L^B and L_i^S and map each trivial q-tangle L_j^H to the empty Jacobi diagram in $\mathcal{A}(\uparrow^{|L_j^H|})$, where $|L_j^H|$ is the number of connected components of L_j^H. By choosing an even associator, we do not need to distinguish between the top and the bottom of the tangles L^B and L_i^S. We finally compose the resulting Jacobi diagrams as prescribed by the polygonal decomposition P_G associated to G

to produce the desired $V_G(L) \in \mathcal{A}(\circlearrowleft^m)$. The invariant depends on the choice of associator and the fatgraph G. We shall also make use of natural extensions of this invariant to certain framed tangles in 1_Σ with endpoints on $(\partial\Sigma) \times \{\frac{1}{2}\}$.

Our definition of V_G differs from [2] insofar as the original invariant takes values in sums of diagrams on the surface, and we are post-composing with the map that forgets the homotopy data of how these diagrams lie in the surface. This is of course a dramatic loss of information, and we wonder what would be the induced equivalence relation on $\mathcal{L}(\Sigma, m)$ assuming faithfulness of the original invariant [2], which gives not only an isotopy invariant but also a universal Vassiliev invariant of links in 1_Σ. See Section 6 for a further discussion.

2.6. Linking pairs and latches

It is a satisfying point that a marked fatgraph G suffices to conveniently determine the choices required to define the AMR invariant $V_G(L) \in \mathcal{A}(\circlearrowleft^m)$ of an m component link L. The fatgraph furthermore determines several other ingredients required for the definition of our new invariants as we finally describe.

Let M be a closed 3-manifold, possibly with boundary. A *linking pair* in M is a 2-component link K arising from an embedding of a standard torus into M, where the first component of K is the core of the torus, called the "longitude" of the pair, and the second is a small null-homotopic 0-framed meridian of it, called the "meridian" of the pair. We say a link L is *disjoint* from a linking pair K in M if K is a linking pair in $M - L$.

In particular in S^3, any two framed links L and $L' = L \sqcup K$, with K a linking pair and L disjoint from K, are related by Kirby I and Kirby II moves. However, this is no longer the case for 3-manifolds with boundary, and one must introduce a third move, called Kirby III, where a linking pair may be added or removed from a surgery link without changing the resulting 3-manifold. The precise statement of the theorem of [33] is that surgery on two framed links in a 3-manifold with boundary determine homeomorphic 3-manifolds if and only if the two links are related by a finite composition of the three Kirby moves, which are sometimes denoted simply KI–III.

Consider the ordered set of generators $X_G = \{x_1, \dots, x_h\}$ of G. For each $x_i \in X_G$, the two paths from its endpoints to the tail in τ_G combine with x_i to form a closed loop based at the tail. By construction, these based loops comprise a (linearly ordered) set of generators for the fundamental group of Σ. Let l_i denote a simple closed curve, representing the free homotopy classes of the ith loop, framed along $\Sigma \times \{1\}$ and pushed off in the I direction in $1_\Sigma = \Sigma \times I$ to height $1 - i\epsilon$, for some small $\epsilon > 0$ fixed independently of i, and pick a small 0-framed meridian m_i of l_i. This provides a collection of linking pairs

$$K_G := \bigcup_i (l_i \cup m_i) \subset 1_\Sigma,$$

called the *system of linking pairs* for 1_Σ determined by the fatgraph G.

Lemma 2.2. *Let G be a marked fatgraph in Σ and let $L \subset 1_\Sigma$ be a framed link disjoint from K_G. There exists a (non-unique) framed link L_0 in $1_\Sigma - K_G$ contained in the preferred box of P_G, such that $L \cup K_G$ is equivalent under isotopy and Kirby II moves to $L_0 \cup K_G$ in 1_Σ.*

Such a representative for a framed link as in the previous lemma is called a *reduced representative*.

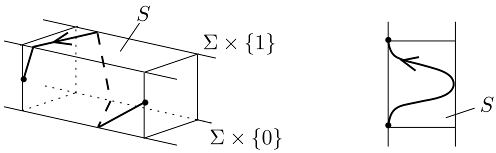

Fig. 4. A latch based at the box S in 1_Σ. The right-hand side is a projection in the I direction of $\Sigma \times I$.

Proof of Lemma 2.2. By an isotopy supported in a neighborhood of the longitudes of K_G, we may arrange that the meridians are all contained in the preferred box. According to Lemma 2.1, we may assume that L is admissible for G and intersects each box except the preferred one in a trivial q-tangle. By Kirby II moves along the meridians, we may arrange that each component of L lies in a different slice of $\Sigma \times I$ than the longitudes. We may furthermore arrange that the link does not meet the box corresponding to any generator X_G of G by sequentially, one generator at a time, performing Kirby II moves along the longitudes of K_G. Each Kirby II move discussed thus far can and furthermore will be performed using bands for the slides that lie within a single box. A final isotopy of the resulting link produces the desired link L_0, and $L \cup K_G$ is equivalent under Kirby II and isotopy to $L_0 \cup K_G$ in 1_Σ by construction. □

One final ingredient, which will serve as the core of the space $\mathcal{A}(\uparrow^h)$ in which our invariant takes its values, is also determined by the generators of the fatgraph G. In each box corresponding to a generator of G, consider an embedded arc in the boundary of 1_Σ as depicted in Fig. 4. Such an arc, called a *latch*, is uniquely determined up to relative homotopy by which side of the box contains its endpoints, and we determine this side as that corresponding to the first oriented edge traversed by the total boundary cycle of G. This collection of latches, one for each generator of G, is called the *system of latches I_G determined* by G, and they occur in a natural orientation and linear order as before. (These standard latches determined by the fatgraph admit a natural generalization given in Section 5.6.1, which is equally well-suited to the construction given in the next section.)

3. The invariant ∇_G

By a *cobordism over Σ*, we mean a 3-manifold $(\Sigma \times I)_L$ obtained by surgery on some framed link L in 1_Σ. In particular, a cobordism over Σ comes equipped with an identification $\partial(\Sigma \times I) \approx \partial(\Sigma \times I)_L$, and two cobordisms are regarded as equivalent if there is a diffeomorphism between them that is equivariant for this identification. Denote by $\mathcal{C}(\Sigma)$ the set of equivalence classes of cobordisms over Σ.

3.1. The invariant ∇_G of cobordisms

Our construction of ∇_G is modeled on the LMO invariant Z^{LMO}, where the role of the Kontsevich integral is now played by the AMR invariant defined in the previous section, and it relies on the LMO maps ι_n, which we next recall and slightly extend.

3.1.1. The map ι_n

This map is a key tool for LMO and for us as well. It “replaces circles by sums of trees” in the rough sense that a core circle component can be erased by suitably summing over all trees spanning the endpoints of a Jacobi diagram in that component.

More precisely, for any pair (m,n) of positive integers, any 1-manifold X without circle components and any linearly ordered $S=\{s_1,\dots,s_m\}$, first define the auxiliary map

$$j_n : \mathcal{A}(X,S) \to \mathcal{A}(X)$$
$$j_n(D) = \begin{cases} O_n(\langle D\rangle) & \text{if } D \text{ has exactly } 2n \text{ vertices labeled with each color,} \\ 0 & \text{otherwise,} \end{cases}$$

where $\langle D\rangle$ is the sum of all possible Jacobi diagrams obtained by pairwise identifying univalent vertices of D having the same color, and O_n serially removes all isolated loops, one at a time and each with a compensatory factor $(-2n)$.

Given $x \in \mathcal{A}(X \sqcup \circlearrowleft^m)$, where X has no circle components and those of $\circlearrowleft^m$ are labeled by $S=\{1,\dots,m\}$ in the natural way, choose an element $y \in \mathcal{A}(X \sqcup \uparrow^m)$ such that $\pi(y)=x$, and consider $\chi^{-1}_{\uparrow^m,S}(y) \in \mathcal{A}(X,S)$. The assignment $\iota_n(x) := (j_n(\chi^{-1}_{\uparrow^m,S}(y)))_{\leqslant n}$ yields a well-defined map

$$\iota_n : \mathcal{A}\big(X \sqcup \circlearrowleft^m\big) \to \mathcal{A}_{\leqslant n}(X).$$

Note that the definition given here is small reformulation of a simplified version [19] of the original [21].

3.1.2. Definition of the invariant $\widehat{\nabla}_G$

Let M be a cobordism over Σ and let G be a marked bordered fatgraph in Σ. G determines the polygonal decomposition P_G of Σ with its forbidden sectors, the system K_G of linking pairs and the system I_G of latches in 1_Σ, as well as the maximal tree τ_G and the system X_G of generators.

Take a representative link $L \subset 1_\Sigma$ for M which is disjoint from I_G and K_G, so that $M=(1_\Sigma)_L=(1_\Sigma)_{L\cup K_G}$. The linking number of two oriented components K_1, K_2 of L in generic position is defined as follows: project K_1, K_2 to $\Sigma \approx \Sigma \times \{0\}$ and sum over all crossings of the projections a sign ± 1 associated to each crossing, where the sign is positive if and only if the projections of the tangent vectors to the over- and under-crossing in this order agree with the given orientation on Σ. For an arbitrary orientation on the link $L\cup K_G$, we denote by $\sigma_+^{L\cup K_G}$, and $\sigma_-^{L\cup K_G}$, the respective number of positive and negative eigenvalues of its linking matrix, which are well-defined independent of choices of orientation on components of $L\cup K_G$.

Denote by V_G the AMR invariant determined by P_G and our choice of even associator. Set

$$\widehat{\nabla}_n^G(L) := \frac{\iota_n(\check{V}_G(L\cup K_G\cup I_G))}{\iota_n(\check{V}_G(U_+))^{\sigma_+^{L\cup K_G}}\,\iota_n(\check{V}_G(U_-))^{\sigma_-^{L\cup K_G}}} \in \mathcal{A}_{\leqslant n}\big(\uparrow^h\big), \tag{6}$$

where $U_\pm$ denotes the ±1-framed unknot in 1_Σ, and $\check{V}_G(\gamma)$ arises from $V_G(\gamma)$ for any framed tangle γ by taking connected sum with ν on each *closed* component, here using that $\{1,\dots,h\}$ is in canonical bijection with X_G.

Theorem 3.1. *For each $n \geqslant 1$, the quantity $\widehat{\nabla}_n^G(L)$ defined in* (6) *does not change under Kirby moves KI–III and does not depend on the orientation of L. Thus, $\widehat{\nabla}_n^G(L)$ is an invariant of the cobordism $(1_\Sigma)_L$.*

Proof. The invariance under Kirby I holds for the usual [29] reason: the change in $\iota_n(\check{V}_G(L \cup K_G \cup I_G))$ under introduction or removal of a ± 1 framed unknot cancels the change in the denominator from $\sigma_\pm^{L\cup K_G}$.

The invariance under KII follows from precisely the same argument as for the LMO invariant, which follows: First observe that an analogue of [21, Proposition 1.3] holds for the AMR invariant: if two links L and L' in 1_Σ differ by a KII move, then $\check{V}_G(L)$ and $\check{V}_G(L')$ are related by a chord KII move, which is the move shown in [21, Fig. 6]. This is true because on one hand, $\check{V}_G$ satisfies (5) since we have chosen to work with an even associator, and on the other hand, we can always assume (up to isotopy of the link) that each handleslide occurs along a band whose projection to Σ is contained in a square S_i in the polygonal decomposition P_G. The invariance under KII is then shown purely at the diagrammatic level, and comes as a consequent property of the map ι_n, whose construction is precisely motivated by its behavior under a chord KII move; see [21, Section 3.1].

We note that using KII moves on the meridian components of K_G, we can alter any crossing of a longitude with any other link component, whence the value of the invariant does not depend on the particular embedding of K_G in M as long as L is disjoint from K_G and the homotopy classes of the longitudes of K_G are preserved.

We finally show that invariance under KIII is guaranteed by the presence of the system K_G of linking pairs. Let L' be obtained by adding a linking pair $l \cup m$ to the link L, where m is a 0-framed meridian of the knot l. Using the fact that the set of homotopy classes provided by the longitudes of K_G can be represented by a system of generating loops for $\pi_1(\Sigma)$, we use KII moves to successively slide l along longitude components of K_G until we obtain a linking pair with longitude null homotopic in 1_Σ and possibly linked with meridians in K_G. We can arrange by isotopy that this linking pair is contained in a 3-ball in 1_Σ and can assume by KII moves that it is unlinked with the meridians of K_G in that 3-ball; as noted earlier, any such linking pair in a 3-ball can be removed using Kirby KI–II moves.

Independence from the choice of orientation on L follows from properties of the map ι_n just as for the LMO invariant; see [21, Section 3.1]. □

Also just as for the LMO invariant, we unify the series $\widehat{\nabla}_n^G$ into a single invariant by setting

$$\widehat{\nabla}_G(M) := 1 + \big(\widehat{\nabla}_1^G(L)\big)_1 + \big(\widehat{\nabla}_2^G(L)\big)_2 + \cdots \in \mathcal{A}\big(\uparrow^h\big) \tag{7}$$

in order to define a map

$$\widehat{\nabla}_G : \mathcal{C}(\Sigma) \to \mathcal{A}\big(\uparrow^h\big).$$

In the case of a 2-disc $\Sigma_{0,1}$ with the convention that a single edge for the tail is allowed to be a fatgraph G, P_G is a disk, and both K_G and I_G are empty, then the invariant $\widehat{\nabla}_G$ exactly coincides with the LMO invariant.

Recall that the space of Jacobi diagrams $\mathcal{A}(\uparrow^h)$ on h intervals is isomorphic to the space $\mathcal{B}(h)$ of h-colored Jacobi diagrams via the Poincaré–Birkhoff–Witt isomorphism. Furthermore, there is the projection of $\mathcal{B}(h)$ onto $\mathcal{A}_h := \mathcal{B}^{\mathsf{Y}}(h)$, and we shall be equally interested in the value our invariant takes in the target space $\mathcal{A}_h$ and hence define

$$\nabla_G : \mathcal{C}(\Sigma) \to \mathcal{A}_h,$$

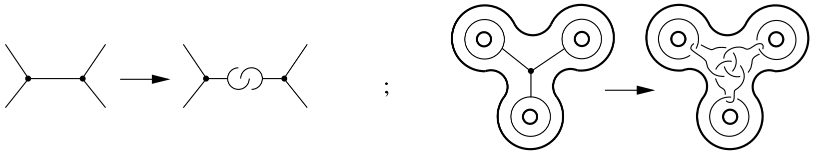

Fig. 5. The edge splitting rule and the surgery link associated to a Y-graph.

where ∇_G is the composition of $\widehat{\nabla}_G$ with the projection $\mathcal{A}(\uparrow^h) \cong \mathcal{B}(h) \to \mathcal{A}_h$. We wonder whether the strut part of $\hat{\nabla}_G(M)$ is related to the homology type of M.

3.2. Universality of ∇_G for homology cylinders

Homology cylinders are a special class of cobordisms which are important in the theory of finite type invariants, cf. [15,13,22]. In this section, we show that for any marked bordered fatgraph G in the surface Σ, the invariant ∇_G of cobordisms is universal among rational-valued finite type invariants of homology cylinders over Σ in the sense of Goussarov and Habiro [15,13]. We first recall the definition of these objects and review the theory of finite type invariants before stating our universality result.

3.2.1. Claspers and finite type invariants of homology cylinders

In this section, we briefly review the Goussarov–Habiro theory of finite type invariants for compact oriented 3-manifolds [13,11,15], which essentially generalizes Ohtsuki's theory [28] for integral homology spheres.

A *clasper* C in a 3-manifold M is an embedding in M of the skinny surface of a (possibly disconnected) Jacobi diagram having a framed copy of S^1 attached to each univalent vertex. The copies of S^1 are called the *leaves* of C, the trivalent vertices are called the *nodes* of C, and we still call the 4-gons associated to the edges of the graph the *edges* of C. We tacitly demand that each connected component of a clasper contains at least one node. The number of connected components of C is denoted $|C|$, and its *degree* is the total number of nodes. A connected clasper of degree 1 is often called a *Y-graph*.

A clasper C of degree k in M determines a framed link $L(C)$ in M, and *surgery along C* means surgery along $L(C)$. To construct $L(C)$ from C, first apply the *edge splitting rule* shown in the left-hand side of Fig. 5 until C becomes a disjoint union of k Y-graphs. Next in a regular neighborhood, replace each Y-graph by a 6-component framed link as shown in the right-hand side of Fig. 5.

The link $L(C)$ (that we sometimes also call a clasper) has $6k$ components if C has degree k. The $3k$ components coming from the k nodes are called the *Borromean components* of $L(C)$, and the remaining $3k$ components are called the *leaf components*. We may also sometimes write simply M_C for the surgery $M_{L(C)}$.

Surgery along a clasper in a 3-manifold produces a new 3-manifold with isomorphic homology and torsion pairing.

A *homology cylinder* over a compact surface Σ is a 3-manifold $M = (1_\Sigma)_C$ that arises from surgery along some clasper C in 1_Σ. Note that, if we identify the points $(x;t)$ and $(x;t')$ in $1_\Sigma = \Sigma \times I$ for $x \in \partial\Sigma$, the cylinder $M = (1_\Sigma)_C$ comes equipped with embeddings $i^\pm : \Sigma \to \partial M$ with respective images $\Sigma^\pm$, such that:

(i) $i^+|_{\partial\Sigma} = i^-|_{\partial\Sigma}$;
(ii) $\partial M = \Sigma^+ \cup (-\Sigma^-)$ and $\Sigma^+ \cap (-\Sigma^-) = \pm\partial\Sigma^\pm$, where $-\Sigma$ denotes reversal of orientation on Σ;
(iii) $i_*^\pm : H_1(\Sigma; \mathbf{Z}) \to H_1(M; \mathbf{Z})$ are identical isomorphisms.

In the special case where Σ has at most one boundary component, such a triple (M, i^+, i^-) satisfying (i)–(iii) conversely always arises from clasper surgery in 1_Σ, cf. [24].

The set of homology cylinders over Σ up to orientation-preserving diffeomorphism is denoted $\mathcal{HC}(\Sigma)$. There is a natural *stacking product* on $\mathcal{HC}(\Sigma)$ that arises by identifying the top of one homology cylinder with the bottom of another and reparameterizing the interval, i.e., by stacking one clasper on top of another. This induces a monoid structure on $\mathcal{HC}(\Sigma)$ with 1_Σ as unit element.

Let $\mathcal{H}_\Sigma$ be the $\mathbb{Q}$-vector space freely generated by elements of $\mathcal{HC}(\Sigma)$ with its descending *Goussarov–Habiro filtration* given by

$$\mathcal{H}_\Sigma \supset \mathcal{F}_1(\Sigma) \supset \mathcal{F}_2(\Sigma) \supset \cdots \tag{8}$$

where for $k \geqslant 1$, $\mathcal{F}_k(\Sigma)$ denotes the subspace generated by elements

$$[M; C] := \sum_{C' \subseteq C} (-1)^{|C'|} M_{C'},$$

with $M \in \mathcal{HC}(\Sigma)$, C a degree $\geqslant k$ clasper in M, and the sum running over all subsets C' of the set of connected components of C.

A *finite type invariant of degree* $\leqslant k$ is a map $f : \mathcal{HC}(\Sigma) \to V$, where V is a $\mathbb{Q}$-vector space, whose natural extension to $\mathcal{H}_\Sigma$ vanishes on $\mathcal{F}_{k+1}(\Sigma)$. Denote by $\mathcal{G}_k(\Sigma)$ the graded quotient $\mathcal{F}_k(\Sigma)/\mathcal{F}_{k+1}(\Sigma)$ and let

$$\overline{\mathcal{H}}_\Sigma := (\text{degree completion of } \mathcal{H}_\Sigma) \Big/ \left(\bigcap_k \mathcal{F}_k(\Sigma) \right).$$

A fundamental open question is whether $\bigcap_k \mathcal{F}_k(\Sigma)$ is trivial.

3.2.2. Universality of the invariant ∇_G

It is known that the LMO invariant is a universal invariant for homology spheres, i.e., every rational-valued finite type invariant of homology spheres factors through it [18]. As noted in Section 3.1, our invariant ∇_G coincides with the LMO invariant for $\Sigma = \Sigma_{0,1}$, and in this section, we prove the following generalization of the universality of LMO.

Theorem 3.2. *Let Σ be a compact connected oriented surface with boundary and G be a marked bordered fatgraph in Σ. Then the invariant ∇_G is a universal finite type invariant of homology cylinders over Σ.*

As an immediate consequence we have

Corollary 3.3. *For each marked fatgraph G in the surface Σ, there is a filtered isomorphism $\overline{\mathcal{H}}_\Sigma \xrightarrow{\cong} \mathcal{A}_h$ induced by the universal invariant ∇_G, where $h = 2g + n - 1$.*

Indeed, the corollary is simply a re-statement of the theorem, and our proof will proceed by exhibiting and checking the isomorphism in Corollary 3.3. This will occupy the remainder of the section and begins with the definition of the inverse map to ∇_G.

3.2.3. The surgery map

The graded quotient $\mathcal{G}_k(\Sigma)$ is generated by elements $[1_\Sigma; C]$, where C is a degree k clasper in 1_Σ since

$$\big[M; C \cup C'\big] = [M; C] - [M_{C'}; C],$$

where $M \in \mathcal{HC}(\Sigma)$ and $C \cup C'$ is a disjoint union of claspers in M with C' connected. Define a filtration

$$\mathcal{G}_k(\Sigma) = \mathcal{F}_{k,3k}(\Sigma) \supset \mathcal{F}_{k,3k-1}(\Sigma) \supset \cdots \supset \mathcal{F}_{k,1}(\Sigma) \supset \mathcal{F}_{k,0}(\Sigma),$$

where $\mathcal{F}_{k,l}(\Sigma)$ is generated by elements $[1_\Sigma; C]$ with C a degree k clasper in 1_Σ having $\leqslant l$ leaves. We also set

$$\mathcal{G}_{k,l}(\Sigma) := \mathcal{F}_{k,l}(\Sigma)/\mathcal{F}_{k,l-1}(\Sigma).$$

Denote by $\mathcal{B}^{\mathsf{Y}}_k(h)$ the i-degree k part of $\mathcal{A}_h = \mathcal{B}^{\mathsf{Y}}(h)$ and denote by $\mathcal{B}^{\mathsf{Y}}_{k,l}(h)$ the subspace generated by Jacobi diagrams of i-degree k with l univalent vertices. Note that $\mathcal{B}^{\mathsf{Y}}_k(h) = \bigoplus_{0\leqslant l\leqslant 3k} \mathcal{B}^{\mathsf{Y}}_{k,l}(h)$.

For any marked fatgraph $G \hookrightarrow \Sigma$ and for any pair k, l of integers with $k \geqslant 1$ and $0 \leqslant l \leqslant 3k$, we define a surgery map $\phi^G_{k,l}$ using claspers as follows. Let $D \in \mathcal{B}^{\mathsf{Y}}_{k,l}(h)$ be some Jacobi diagram. For each univalent vertex v of D labeled by i, consider an oriented framed knot in 1_Σ which is a parallel copy of the longitude l_i of the system K_G of linking pairs. For each trivalent vertex of D, consider an embedded oriented disk in 1_Σ. These choices are made subject to the constraint that the resulting annuli and disks are pairwise disjoint in 1_Σ. Connect these various embedded annuli and disks by disjoint bands as prescribed by the diagram D in a way which is compatible with their orientations and such that the cyclic order of the three attached bands (given by the orientation) at each disk agrees with the cyclic order at the corresponding trivalent vertex of D. The resulting surface is a degree k clasper with l leaves in 1_Σ denoted $C(D)$, and we set

$$\phi^G_{k,l}(D) := \big[1_\Sigma; C(D)\big] \in \mathcal{G}_{k,l}(\Sigma).$$

It follows from [10, Theorem 1] (see also [14]) that this assignment yields a well-defined surjection

$$\phi^G_{k,l} : \mathcal{B}^{\mathsf{Y}}_{k,l}(h) \to \mathcal{G}_{k,l}(\Sigma).$$

The proof makes use of the calculus of claspers; see [8,11,15,29] for similar results.

3.2.4. Proof of Theorem 3.2

We need to prove the following two facts:

Fact (1) The i-degree $\leqslant k$ part of ∇_G is a finite type invariant of degree k.

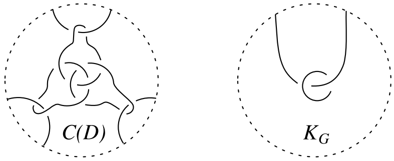

Fig. 6. Convenient positions for $C(D)$ and K_G.

Fact (2) For each pair (k,l) with $k \geqslant 1$ and $0 \leqslant l \leqslant 3k$, the invariant ∇_G induces the inverse to the surgery map $\phi^G_{k,l}$. More precisely, given a Jacobi diagram $D \in \mathcal{B}^{\mathsf{Y}}_{k,l}(h)$, we have

$$\nabla^G_{k,l}\big(\phi^G_{k,l}(D)\big) = (-1)^k D \in \mathcal{B}^{\mathsf{Y}}_{k,l}(h), \tag{9}$$

where $\nabla^G_{k,l}$ denotes the composition of ∇_G with the projection onto $\mathcal{B}^{\mathsf{Y}}_{k,l}(h)$.

In order to prove Fact (1), it is enough to consider an element $[1_\Sigma; C]$, where C is a disjoint union of k Y-graphs in 1_Σ (by construction using the edge splitting rule) and prove that the minimal i-degree of $\nabla_G([1_\Sigma; C])$ is k. To this end, we may up to isotopy assume that there are k disjoint 3-balls contained in the boxes of the polygonal decomposition P_G corresponding to G which intersect $C(D)$ as depicted on the left-hand side of Fig. 6. Note that

$$[1_\Sigma; C] = \sum_{C' \subseteq C} (-1)^{|C'|} (1_\Sigma)_{C'} = \sum_{C' \subseteq C} (-1)^{|C'|} (1_\Sigma)_{L_0(C')},$$

where $L_0(C')$ is obtained from the link $L(C)$ by replacing each Borromean linking corresponding to a node of $C - C'$ by a trivial linking, so in particular $L_0(C')$ is Kirby equivalent to $L(C')$. In the computation of $\nabla_G([1_\Sigma; C])$ at lowest i-degree, we thus obtain for each node of C a trivalent vertex attached to the three corresponding core components. This follows [29] from the property

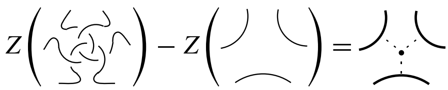 + terms with i-degree $\geqslant 2$ (⋆)

of the Kontsevich integral and implies that the minimal i-degree of $\nabla_G([1_\Sigma; C])$ is k as required to prove Fact (1).

Turning our attention now to Fact (2), let $D \in \mathcal{B}^{\mathsf{Y}}_{k,l}(h)$ be a Jacobi diagram of i-degree k with l univalent vertices and consider $\phi^G_{k,l}(D) = [1_\Sigma; C(D)]$. Denote by $J = \frac{1}{2}(k+l)$ the Jacobi degree of D. As before, we can assume that there are k disjoint 3-balls in the boxes of P_G each of which intersects $C(D)$ as depicted on the left-hand side of Fig. 6, and we can assume that there are a further h disjoint 3-balls that intersect the system K_G of linking pairs as illustrated on the right-hand side of the same figure.

Let us now compute the (relevant part of the) AMR invariant $\check{V}_G$ of the alternating sum $\phi^G_{k,l}(D)$. Since we are only computing the lowest i-degree part, the contributions of all associators and ν's can be ignored, see (25) and [5] respectively. In fact, we shall only need to consider the contributions arising from the trivalent vertices (⋆) and struts coming from crossings, cf. (⋆⋆). Since the value of $\nabla^G_{k,l}$ is of J-degree $J = \frac{k+l}{2}$, we must post-compose $\check{V}_G(\phi^G_{k,l}(D))$ with the map

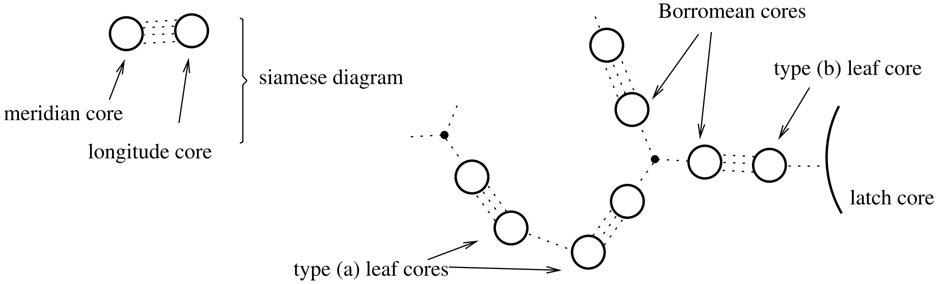

Fig. 7. Typical lowest i-degree term.

ι_J to establish the formula in Fact (2). We need only focus on those terms of $\mathcal{A}(\bigcirc^{6k+2h} \sqcup \uparrow^h)$ having exactly $2J$ vertices on each copy of S^1 by definition of ι_J since only those terms can contribute to the lowest i-degree part of $\nabla_G(\phi^G_{k,l}(D))$.

We call the core components corresponding to Borromean (leaf, meridian, longitude, latch respectively) components of $L(C(D)) \cup K_G \cup I_G$ the Borromean (leaf, meridian, longitude, latch respectively) cores of the Jacobi diagrams in the AMR invariant $\check{V}_G(L(C(D)) \cup K_G \cup I_G)$.

We first consider contributions of the linking pairs and recall [29] that

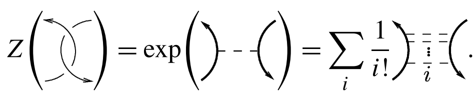

(★★)

Since the meridian component of a linking pair is isolated from every component other than its corresponding longitude, it follows that all $2J$ vertices on the meridian core must be the ends of distinct struts arising from the linking with this longitude. The resulting connected diagram, which arises from (★★) with a coefficient $\frac{1}{(2J)!}$, is called a *Siamese diagram*; see Fig. 7.

Each Borromean component of $L(C(D))$ on the one hand forms a Borromean linking with two other such components and, on the other hand links a leaf, cf. Fig. 6. On each Borromean core there is thus one vertex arising from (⋆), and the remaining $(2J-1)$ vertices are the ends of parallel struts arising from (★★) with their opposite ends on a leaf core.

On each leaf core, there is thus only room for one additional vertex. Furthermore, there is the following dichotomy on leaves of $C(D)$:

(a) the leaf forms a positive Hopf link with another leaf of $C(D)$ as in the left-hand side of Fig. 5;
(b) the leaf is a parallel copy of a longitude component l_i of K_G, pushed off so that it is unlinked from the meridian m_i, for some $1 \leqslant i \leqslant h$.

For a type (a) leaf, the only possible contribution is the linking with another type (a) leaf, which produces a strut by (★★). For a type (b) leaf, we must consider several cases: either the strut comes from a crossing with another type (b) leaf, or it comes from a crossing with a component of I_G (since a crossing with the longitude components of K_G cannot contribute, as we have noted previously). In the first case, we thus have a strut joining two type (b) leaf cores, and we say that a Jacobi diagram with such a strut is *looped*. In the second case, we have a strut joining the leaf core to a latch core. A typical example is (partially) represented in Fig. 7.

By construction, each type (b) leaf has linking number 1 with exactly one component of I_G and 0 with all others. The lowest i-degree terms in $\check{V}_G([1_\Sigma; C(D)])$ are therefore a sum of looped diagrams plus a single Jacobi diagram with each type (b) leaf core connected by a strut to a latch core.

We can now apply the map ι_J, and a computation shows that

$$\iota_k\Bigg($$ 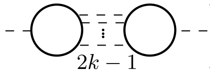

$$\Bigg) = (-1)^{k-1}(2k-1)! - - - - - .$$

We find two (and one respectively) such configurations for each edge incident (and not incident) on a univalent vertex of D, and each comes from $(\star\star)$ with a coefficient $\frac{1}{(2J-1)!}$. This formula also shows that ι_J maps each Siamese diagram to a factor $(-1)^J(2J)!$, and there are $h = 2g + n - 1$ such diagrams. We obtain that $\iota_J(\check{V}_J(\phi^G_{k,l}(D)))$ is given by

$$(-1)^{(J-1)l+Jh}D + \begin{cases} \text{terms of } i\text{-degree } k \text{ with less than } l \text{ univalent vertices,} \\ \text{terms of } i\text{-degree} > k, \end{cases}$$

where the terms of i-degree k with less than l univalent vertices arise from the looped Jacobi diagrams.

To conclude the computation, observe that the surgery link $L(C(D)) \cup K_G$ satisfies $\sigma_+^{L(C(D))\cup K_G} = \sigma_-^{L(C(D))\cup K_G} = 3k + h$. Since $\iota_J(\check{Z}(U_\pm)) = (\mp 1)^J +$ terms of i-degree $\geqslant 1$, we therefore find

$$\iota_J\big(\check{V}_P(U_+)\big)^{\sigma_+^{L(C(D))\cup K_G}} \iota_J\big(\check{V}_P(U_-)\big)^{\sigma_-^{L(C(D))\cup K_G}} = (-1)^{Jk+Jh} + \text{terms of } i\text{-degree} \geqslant 1.$$

It follows that

$$\nabla^G_{k,l}\big(\phi^G_{k,l}(D)\big) = (-1)^k D,$$

which concludes the proof of Theorem 3.2.

3.3. *The rigid ∇^r_G invariant*

In this section, we introduce a modified "rigid" version ∇^r_G of our invariant ∇_G, which is formulated in terms of the LMO invariant of tangles and again depends on the choice of a marked fatgraph for Σ. In this incarnation, ∇^r_G shares properties with the invariant defined in [8], which gives an extension of the LMO invariant to so-called Lagrangian cobordisms between once-bordered and closed surfaces. The invariant in [8] depends upon choices similar to certain of those determined by a fatgraph discussed here, and it induces a universal invariant for homology cylinders. Roughly, it is defined by first "capping off" a cobordism by attaching 2-handles along the boundary producing a tangle in a homology ball and then computing the LMO invariant (actually, the equivalent Århus integral) of this tangle.

3.3.1. *The rigid ∇^r_G invariant for homology cylinders*

Let G be a marked fatgraph in the bordered surface Σ. Recall that the bigon B in the polygonal decomposition P_G of Section 2.4 gives rise to the preferred box in 1_Σ, which is identified with the standard cube $C = [0, 1]^3$ so that the upper face $f = [0, 1]^2 \times \{1\}$ is the cutting face. Denote

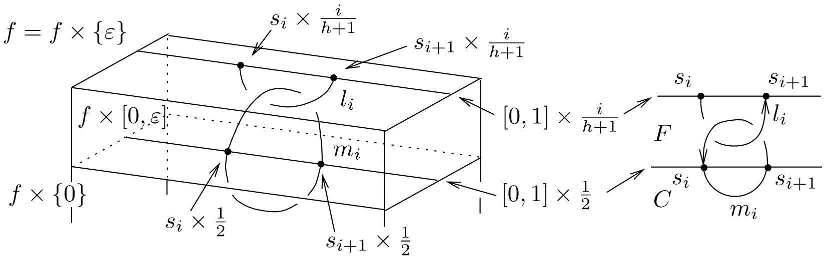

Fig. 8. System of linking pairs in rigid position.

by F a collar neighborhood $f \times [0, \varepsilon]$ of f in $(1_\Sigma - C)$, where f is identified with $f \times \{0\}$, and fix the standard points $s_i := \frac{i}{2h+1} \in [0, 1]$, for $1 \leqslant i \leqslant 2h$.

Consider in 1_Σ the system of linking pairs $K_G = \bigsqcup_{i=1}^{h}(l_i \sqcup m_i)$ determined by G as in Section 2.6. We assume that each l_i lies in the surface $\Sigma \times \{\frac{i}{h+1}\} \subset 1_\Sigma$, for $1 \leqslant i \leqslant h$ and meets ∂F only at the points $s_j \times \{\frac{i}{h+1}\} \times \{\varepsilon\}$, for $j = i, i+1$, that each meridian m_i intersects ∂F only at the points $s_j \times \{\frac{1}{2}\} \times \{0\}$, again for $j = i, i+1$, and that $(m_i \cup l_i) \cap (F \cup C)$ is in the standard position depicted in Fig. 8 up to isotopy. We say that K_G is in *rigid position* in 1_Σ in this case.

Suppose that $M = (1_\Sigma)_L$ is a cobordism over Σ for some framed link L in 1_Σ. By Lemma 2.2, we can use the system of linking pairs K_G in rigid position to obtain a reduced representative L_0 which lies in the preferred box C.

Letting I_G denote the system of latches determined by G (cf. Section 2.6), cut 1_Σ along f in order to split $I_G \cup K_G \cup L_0$ into two q-tangles

$$T_G := (I_G) \cup \big(K_G \cap (1_\Sigma - C)\big) \quad \text{and} \quad (K_G \cap C) \cup (L_0), \tag{10}$$

where the bracketing $(\bullet\bullet)((\bullet\bullet)((\bullet\bullet)\cdots((\bullet\bullet)(\bullet\bullet))\cdots))$ is taken on both sets of boundary points. Set

$$\widehat{\nabla}_n^{G,r}(L) := \frac{\iota_n(\check{Z}((K_G \cap C) \cup L_0))}{\iota_n(\check{Z}(U_+))^{\sigma_+^L}\iota_n(\check{Z}(U_-))^{\sigma_-^L}} \in \mathcal{A}_{\leqslant n}\big(\uparrow^h\big), \tag{11}$$

where we make use of the same notation as for (6), and define the *rigid* $\widehat{\nabla}_G^r$ *invariant* of M to be

$$\widehat{\nabla}_G^r(M) := 1 + \big(\widehat{\nabla}_1^{G,r}(L)\big)_1 + \big(\widehat{\nabla}_2^{G,r}(L)\big)_2 + \cdots \in \mathcal{A}\big(\uparrow^h\big). \tag{12}$$

As before, we define the corresponding *rigid* ∇_G^r *invariant* as the composition of $\widehat{\nabla}_G^r$ with the projection $\mathcal{A}(\uparrow^h) \to \mathcal{A}_h$.

In the next section, we prove the following:

Theorem 3.4. ∇_G^r *is an invariant of homology cylinders and induces a graded isomorphism*

$$\nabla_G^r : \overline{\mathcal{H}}_\Sigma \to \mathcal{A}_h$$

for any marked fatgraph G in Σ.

Though ∇_G^r is defined for any cobordism over Σ, we can at present only prove it is an invariant of homology cylinders; cf. the next section.

Remark 3.5. Recall that the LMO invariant extends naturally to q-tangles in homology balls. This is done in a similar manner to the extension to links in 3-manifolds [21], and more generally to framed graphs in 3-manifolds [27], via a formula similar to (11). Using this extension, we can reformulate (12) as

$$\widehat{\nabla}_G^r(M) := Z^{LMO}(B_M, \gamma_M) \in \mathcal{A}(\uparrow^h),$$

where (B_M, γ_M) denotes the result of surgery on $(C, K_G \cap C)$ along the link L_0. Note that B_M is indeed a homology ball since M is a homology cylinder.

3.3.2. Proof of Theorem 3.4

For $h \geqslant 1$, let $T_{h,0} = \bigotimes_h C_+$ be the q-tangle in C obtained by horizontal juxtaposition of h copies of the q-tangle C_+ of Fig. 2, with bracketing of the form $(\bullet\bullet)((\bullet\bullet)((\bullet\bullet)\cdots((\bullet\bullet)(\bullet\bullet))\cdots))$. See Fig. 10 below and set

$$\mathcal{T}(h) := \{(C, T_{h,0})_\Gamma\colon\ \Gamma \text{ is a clasper in } C \text{ disjoint from } T_{h,0}\}.$$

Let $\mathbb{Q}\mathcal{T}(h)$ denote the vector space freely generated by elements of $\mathcal{T}(h)$. In analogy to (8), we have the Goussarov–Habiro filtration

$$\mathbb{Q}\mathcal{T}(h) \supset \mathcal{F}_1(h) \supset \mathcal{F}_2(h) \supset \cdots,$$

where $\mathcal{F}_k(h)$ denotes the subspace generated by elements $[(B, \gamma); \Gamma]$ with $(B, \gamma) \in \mathcal{T}(h)$ and with Γ a degree $\geqslant k$ clasper in B disjoint from γ, for $k \geqslant 1$. This filtration serves to define a notion of finite type invariants for these objects as in Section 3.2.1.

Proposition 3.6. *For any $h \geqslant 1$, the LMO invariant induces a universal finite type invariant for tangles in $\mathcal{T}(h)$.*

Proof. The proof follows closely that of Theorem 3.2. In particular as in Section 3.2.3, we define for each pair (k, l) with $k \geqslant 1$ and $0 \leqslant l \leqslant 3k$ a surgery map $\phi_{k,l}^r$ as follows. Let $D \in \mathcal{B}_{k,l}^{\mathsf{Y}}(h)$. For each i-labeled univalent vertex, pick a parallel copy of a small 0-framed meridian of the ith component of $T_{h,0}$, and for each trivalent vertex of D, pick an embedded oriented disk in C. Connect these meridians and disks by disjoint bands as prescribed by the diagram D to obtain a degree k clasper with l leaves denoted $C^r(D)$. The assignment $\phi_{k,l}^r(D) := [(C, T_{h,0}); C^r(D)]$ yields a well-defined surjective map

$$\phi_{k,l}^r : \mathcal{B}_{k,l}^{\mathsf{Y}}(h) \to \mathcal{G}_{k,l}(h),$$

where $\mathcal{G}_{k,l}(h)$ is defined as in Section 3.2.3. The rest of the proof follows from the analogues of Facts (1) and (2) of Section 3.2.4, which hold according to exactly the same arguments. □

As a consequence, we have a graded isomorphism

$$Z^{LMO} : \overline{\mathcal{T}}(h) \xrightarrow{\cong} \mathcal{A}_h$$

induced by the LMO invariant, where $\overline{\mathcal{T}}(h)$ denotes the quotient

$$\big(\text{degree completion of } \mathbb{Q}\mathcal{T}(h)\big)\Big/\Big(\bigcap_{k\geqslant 1}\mathcal{F}_k(h)\Big).$$

The inverse isomorphism, denoted ϕ^r, is induced by the surgery maps $\phi^r_{k,l}$.

We can now proceed with the proof of Theorem 3.4. For any marked fatgraph G in Σ, define a map

$$J_G : \mathcal{T}(h) \to \mathcal{HC}(\Sigma)$$

as follows. If $(B,\gamma) = (C, T_{h,0})_\Gamma \in \mathcal{T}(h)$, where Γ is some clasper in C disjoint from $T_{h,0}$, then $J_G(B,\gamma)$ is the homology cylinder obtained by stacking the tangle T_G defined in (10) above (B,γ) and performing surgery along the $2h$-component link resulting from this stacking. We shall give in Remark 4.3 a purely diagrammatic version of the map J_G for any marked bordered fatgraph G. As a generalization of the Milnor–Johnson correspondence of Habegger [14], we wonder if J_G is invertible; if so, then it would follow that the rigid invariant ∇^r_G is indeed an invariant not just of homology cylinders but also of general cobordisms over Σ.

Since $J_G(\bigcap_{k\geqslant 1}\mathcal{F}_k(h)) \subset (\bigcap_k \mathcal{F}_k(\Sigma))$, there is an induced map

$$\overline{J}_G : \overline{\mathcal{T}}(h) \to \overline{\mathcal{H}}_\Sigma,$$

which is surjective according to Lemma 2.2.

Lemma 3.7. *The map $\overline{J}_G$ is a graded isomorphism.*

This implies Theorem 3.4 since (12) can thus be rewritten as $\nabla^r_G = Z^{LMO} \circ (\overline{J}_G)^{-1}$.

Proof of Lemma 3.7. For each pair (k,l) with $k \geqslant 1$ and $0 \leqslant l \leqslant 3k$, consider the surjective map $J^{k,l}_G : \mathcal{G}_{k,l}(h) \to \mathcal{G}_{k,l}(\Sigma)$ induced by J_G. It suffices to show that $J^{k,l}_G$ is a graded isomorphism, which follows from commutativity of the following diagram:

$$\begin{array}{ccc} \mathcal{B}^{\mathsf{Y}}_{k,l}(h) & & \\ {\scriptstyle \phi^r_{k,l}}\downarrow & \searrow^{\phi^G_{k,l}} & \\ \mathcal{G}_{k,l}(h) & \xrightarrow[J^{k,l}_G]{} & \mathcal{G}_{k,l}(\Sigma). \end{array}$$

To see that this diagram is in fact commutative, let $D \in \mathcal{B}^{\mathsf{Y}}_{k,l}(h)$ and $C^r(D) \subset C$ be the clasper obtained by the construction explained in the proof of Proposition 3.6, whence $\phi^r_{k,l}(D) = [(C, T_{h,0}); C^r(D)]$. Applying $J^{k,l}_G$ amounts to stacking the tangle T_G on $(T_{h,0} \cup C^r(D)) \subset C$. The result of this stacking is the system of linking pairs K_G in rigid position, together with a clasper with l leaves in 1_Σ, each leaf being a disjoint copy of a 0-framed meridian of the component m_i of K_G, for some i. By a Kirby KII move, we can slide each of these leaves along the corresponding longitude component l_i of K_G and denote by Γ the resulting clasper in 1_Σ. It follows from the definition of the surgery map $\phi_{k,l}$ (see Section 3.2.3) that $[1_\Sigma; \Gamma] = \phi_{k,l}(D)$ in $\mathcal{G}_{k,l}(\Sigma)$ as required. □

4. Diagrammatic formulations of topological gluings

Throughout this section, fix a non-negative integer g as well as a closed genus g surface Σ_g which we identify with the boundary of the standard genus g handlebody $H_g := \Sigma_{0,g+1} \times I$, where $\Sigma_{0,g+1}$ is a fixed disc in the plane with basepoint on its boundary having g holes ordered and arranged from left to right. We also fix a genus g surface with one boundary component $\Sigma_{g,1}$ and identify $\Sigma_g = \Sigma_{g,1} \cup D^2$ with the closed surface obtained by capping off $\Sigma_{g,1}$ with a disc D^2, so that $\partial\Sigma_{g,1} \subset (\partial\Sigma_{0,g+1}) \times I$.

4.1. Topological operations

4.1.1. Homology handlebodies

A genus g *homology handlebody* is a 3-manifold M with boundary a closed genus g surface Σ such that the inclusion $\Sigma \hookrightarrow M$ induces a surjection in integral homology with kernel a maximal integral isotropic subgroup $\Lambda \subset H_1(\Sigma_g; \mathbb{Z})$; in this definition, we always require an identification of the boundary Σ of M with the fixed surface Σ_g and call Λ the *Lagrangian* of the handlebody. For example, H_g is a homology handlebody, whose associated Lagrangian subspace Λ^{st} we call the *standard Lagrangian* of Σ_g.

We consider two homology handlebodies v_1 and v_2 equivalent if there is a diffeomorphism of v_1 to v_2 which restricts to the identity on Σ_g under the corresponding identifications. Denote by $V(\Sigma_g, \Lambda)$ the set of equivalence classes of homology handlebodies with Lagrangian Λ.

By a result of Habegger [14], any two homology handlebodies are related by clasper surgeries if and only if they have the same induced Lagrangian. In particular, any genus g homology handlebody with Lagrangian Λ^{st} can be obtained by clasper surgery in H_g. In other words, we have

$$V\left(\Sigma_g, \Lambda^{st}\right) = \mathcal{H}C(\Sigma_{0,g+1}).$$

4.1.2. Stacking, shelling and pairing

We have already defined in (Section 3.2.1) the natural stacking product for homology cylinders, which induces a map

$$\cdot : \overline{\mathcal{H}}_{\Sigma_{g,1}} \times \overline{\mathcal{H}}_{\Sigma_{g,1}} \to \overline{\mathcal{H}}_{\Sigma_{g,1}},$$

and we next similarly introduce two further products

$$\bigcup_\iota : \overline{\mathcal{H}}_{\Sigma_{0,g+1}} \times \overline{\mathcal{H}}_{\Sigma_{0,g+1}} \to \overline{\mathcal{H}}_{\Sigma_{0,1}},$$
$$* : \overline{\mathcal{H}}_{\Sigma_{g,1}} \times \overline{\mathcal{H}}_{\Sigma_{0,g+1}} \to \overline{\mathcal{H}}_{\Sigma_{0,g+1}}.$$

The pairing $\bigcup_\iota$ on the vector space $\overline{\mathcal{H}}_{\Sigma_{0,g+1}}$ is defined as follows. Consider the standard orientation-reversing map $\iota : \Sigma_g \to \Sigma_g$ which "takes longitudes to meridians" and vice versa so that gluing two copies of H_g along their boundaries via ι produces the standard 3-sphere S^3.

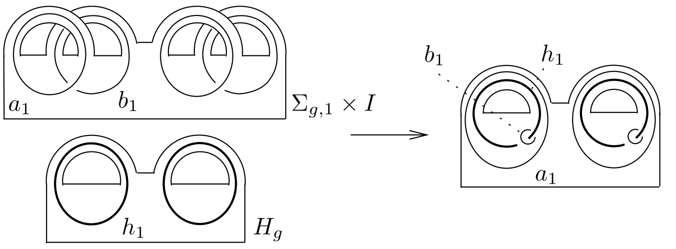

Fig. 9. The shelling $(\Sigma_{g,1} \times I) * H_g$ and the curves a_i, b_i, h_i.

By gluing two arbitrary handlebodies with boundary Σ_g along this map (i.e., their adjunction space collapsing fibers to points), we obtain a closed 3-manifold and refer to this operation as their *pairing*. Observe that the pairing of two homology handlebodies in $\overline{\mathcal{H}}_{\Sigma_{0,g+1}}$ is an integral homology 3-sphere, or equivalently, a homology cylinder over $\Sigma_{0,1}$, as required.

The *shelling product* $*$ is defined as follows. Given a genus g homology handlebody H and a homology cylinder (N, i^+, i^-) over Σ_g in the notation of Section 3.2.1, we can glue the boundaries via the identification $\partial H = i^-(\Sigma_g) \subset \partial N$ to obtain a new genus g homology handlebody with boundary $i^+(\Sigma_g)$, denoted by $N * H$. Similarly, given the identification $\Sigma_g = \Sigma_{g,1} \cup D^2$, we can glue a homology cylinder M over $\Sigma_{g,1}$ to the homology handlebody H to obtain a 3-manifold with boundary $i^+(\Sigma_{g,1}) \cup (S^1 \times I) \cup D^2$. By gluing a cylinder $D^2 \times I$ along $(S^1 \times I) \cup D^2$ in the standard way, we obtain a new genus g homology handlebody $M * H$.

To illustrate the shelling product, let $\{a_i, b_i\}_{i=1}^g$, respectively, $\{h_i\}_{i=1}^g$, be the collection of disjoint loops in $\Sigma_{g,1} \times I$, respectively, in the genus g handlebody H_g, shown in Fig. 9. Note that each collection induces a basis for the first homology group of the corresponding 3-manifold. The images of these loops under the shelling product $H_g = (\Sigma_{g,1} \times I) * H_g$, which we still denote by a_i, b_i and h_i, are shown on the right-hand side of the figure. In particular, note that each b_i is null-homotopic in H_g, and satisfies $|lk(b_i, h_i)| = 1$.

The main goal of this section is to provide explicit diagrammatic formulas in Section 4.3 for these three topological operations.

We note that similar stacking and shelling operations can be defined for surfaces $\Sigma_{g,n}$ with $n > 1$ boundary components; however, we restrict here to the case $n = 1$ for simplicity.

4.2. A general gluing formula

We now introduce another more basic operation, which is a key tool for manipulating our diagrammatic formulas.

4.2.1. The contraction $\circ$ of labeled Jacobi diagrams

Let $D \in \mathcal{B}(S)$ and $D' \in \mathcal{B}(S')$ be diagrams, for some finite sets S and S', and let $R \subseteq S \cap S'$. Define the *contraction product* $D \circ_R D' \in \mathcal{B}((S \cup S') - R)$, as follows. If $R = \emptyset$ or if for some $x \in R$ the number of x-colored vertices of D and D' is not the same, then set $D \circ_R D' = 0$, and otherwise, $D \circ_R D'$ is defined to be the sum of all possible ways of gluing pairwise the univalent vertices of D and D' labeled by the same element of R. By linear extension, this defines a contraction map

$$\circ_R : \mathcal{B}(S) \times \mathcal{B}(S') \to \mathcal{B}\big((S \cup S') - R\big),$$

which we will call the *contraction over* R.

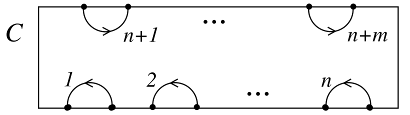

Fig. 10. The q-tangle $T_{m,n}$, where the bracketing on both sets of boundary points is of the form $(\bullet\bullet)((\bullet\bullet)\times((\bullet\bullet)\cdots((\bullet\bullet)(\bullet\bullet))\cdots))$.

Let $-_s \in \mathcal{A}(\uparrow,\{s\})$ be the Jacobi diagram consisting of a single strut with one vertex on $\uparrow$ and one vertex colored by s. Set

$$\lambda(s,u,v) := \chi_{\{v\}}^{-1}\big(\exp(-_s)\cdot\exp(-_u)\big) \in \mathcal{B}\big(\{s,u,v\}\big),$$

where the exponential is with respect to the stacking product of Jacobi diagrams.[1] If S, U and V respectively denote the sets $\{s_1,\dots,s_n\}$, $\{u_1,\dots,u_n\}$ and $\{v_1,\dots,v_n\}$, then define

$$\Lambda^n(S,U,V) := \bigsqcup_{i=1}^{n} \lambda(s_i,u_i,v_i) \in \mathcal{B}(S\cup U\cup V).$$

Proposition 4.1. *(See [4, Proposition 5.4], also [8, Claim 5.6].) For $n \geqslant 1$, let $D \in \mathcal{A}(X\cup\uparrow^n)$ and $E \in \mathcal{A}(X'\cup\uparrow^n)$, where X and X' are two (possibly empty) 1-manifolds. Let $D\cdot E \in \mathcal{A}(X\cup X'\cup\uparrow^n)$ be obtained from the stacking product of $\mathcal{A}(\uparrow^n)$. Then*

$$\begin{aligned}\chi_{\uparrow^n,V}^{-1}(D\cdot E) &= \Lambda^n(S,U,V)\circ_{S\cup U}\big(\chi_{\uparrow^n,S}^{-1}(D)\sqcup\chi_{\uparrow^n,U}^{-1}(E)\big)\\ &= \chi_{\uparrow^n,S}^{-1}(D)\circ_S \Lambda^n(S,U,V)\circ_U \chi_{\uparrow^n,U}^{-1}(E)\\ &\in \mathcal{A}\big(X\cup X',V\big).\end{aligned}$$

4.2.2. A gluing formula for the LMO invariant of tangles

For $m,n \geqslant 0$, denote by $T_{m,n}$ the q-tangle in the standard cube $C := [0,1]^3$ represented in Fig. 10. In the natural way, we consider the tangles $T_{m,0}$ and $T_{0,n}$ as subtangles of $T_{m,n}$.

Consider the q-tangle $\gamma' = T_{m,n}\cup L'$ in C, where L' is some framed link disjoint from $T_{m,n}$. For any element $D \in \mathcal{A}(\uparrow^n)$, define

$$D \overset{k}{\odot} \gamma' = \chi_S^{-1} D \circ_S \frac{\iota_k j_k^V(\Lambda^n(S,U,V)\circ_U \chi_{T_{0,n},U}^{-1}\hat{Z}(C,\gamma'))}{\iota_k(\check{Z}(U_+))^{\sigma_+(\gamma')}\iota_k(\check{Z}(U_-))^{\sigma_-(\gamma')}} \in \mathcal{A}\big(\uparrow^m\big), \tag{13}$$

and set

$$D\odot\gamma' := 1 + \sum_{k\geqslant 1}\big(D\overset{k}{\odot}\gamma'\big)_k \in \mathcal{A}\big(\uparrow^m\big), \tag{14}$$

where

[1] As explained in [4, Proposition 5.4] and [8, Remark 4.8], $\lambda(s,u,v)$ can also be defined in terms of the Baker–Cambell–Hausdorff series.

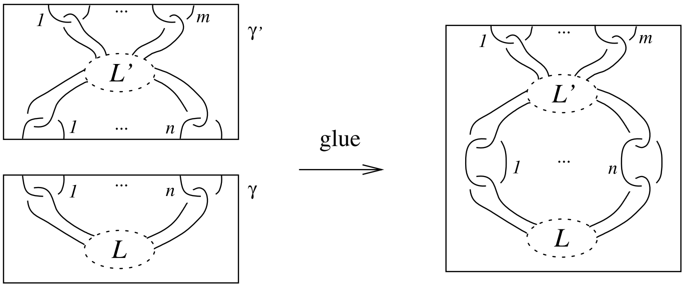

Fig. 11. The tangles γ and γ', and their composition.

- $\hat{Z}(C,\gamma')$ is obtained from $\check{Z}(C,\gamma')$ by taking connected sum of ν with each component of $T_{0,n} \subset \gamma'$.
- $\sigma_\pm(\gamma') := \sigma_\pm^{L'\cup T_{0,n}}$ denotes the number of positive and negative eigenvalues of the linking matrix of the tangle $(T_{0,n} \cup L') \subset \gamma'$.
- The map $j_k^V : \mathcal{B}(S \cup V) \to \mathcal{B}(S)$ is defined by applying j_k to the V-colored vertices, as in the definition of Section 3.1.1, and leaving the S-colored vertices unchanged.

It will be useful to define an analogous product for n-colored Jacobi diagrams, still denoted by $\odot$. For $E \in \mathcal{B}(n)$, we set

$$E \odot \gamma' := \chi^{-1}\big(\chi(E) \odot \gamma'\big) \in \mathcal{B}(m). \tag{15}$$

We now use this product to give a gluing formula for the LMO invariant. Let $\gamma = T_{n,0} \cup L$ be a q-tangle in C, where L is a framed link disjoint from $T_{n,0}$ so that $b = C_L$ is an integral homology ball (in particular, L can be chosen to be a clasper); see Fig. 11. Set

$$(b,t) = (C, T_{n,0})_L \in \mathcal{T}(n).$$

Let O_n denote the n component link arising as the composition of $T_{n,0}$ and $T_{0,n}$, so that $\gamma \cdot \gamma' = L' \cup O_n \cup L \cup T_{m,0}$.

Lemma 4.2. *Let γ and γ' be two q-tangles as described above. Then the LMO invariant of*

$$(B,T) := (C, T_{m,0})_{L\cup O_n \cup L'}$$

is given by

$$Z^{LMO}(B,T) = Z^{LMO}(b,t) \odot \gamma' \in \mathcal{A}\big(\uparrow^m\big).$$

Proof. Let δ denote the tangle in b obtained from $\gamma \cdot \gamma'$ by surgery along the link L. By definition, the degree k part of $Z^{LMO}(B,T)$ is given by

$$Z_k^{LMO}(B,T) = \left(\frac{\iota_k^{L'} \iota_k^{O_n} \iota_k^{L}(\check{Z}(C,\gamma\cdot\gamma'))}{(\iota_k(\check{Z}(U_+)))^{\sigma_+^{L'\cup O_n \cup L}} (\iota_k(\check{Z}(U_-)))^{\sigma_-^{L'\cup O_n \cup L}}} \right)_k,$$

where $\iota_k^{O_n}$, respectively, $\iota_k^{L'}$ and ι_k^{L}, denote the map ι_k applied only to the copies of S^1 corresponding to $O_n \subset \gamma \cdot \gamma'$, respectively, to $L', L \subset \gamma \cdot \gamma'$, as in the definition of Section 3.1.1.

By following (the proof of) [21, Theorem 6.6], we have

$$Z_k^{LMO}(B,T) = \left(\frac{\iota_k^{L'} \iota_k^{O_n}(Z^{LMO}(b,\delta))}{(\iota_k(\check{Z}(U_+)))^{\sigma_+^{L'\cup O_n}} (\iota_k(\check{Z}(U_-)))^{\sigma_-^{L'\cup O_n}}} \right)_k,$$

and

$$\iota_k^{L'} \iota_k^{O_n}\big(Z^{LMO}(b,\delta)\big) = \big(\iota_k^{L'} j_k^V \chi_{O_n,V}^{-1}\big(Z^{LMO}(b,\delta)\big)\big)_{\leqslant k}$$

from the definition of ι_k. By Proposition 4.1,

$$\chi_{O_n,V}^{-1}\big(Z^{LMO}(b,\delta)\big) = \chi_S^{-1} Z^{LMO}(b,t) \circ_S \Lambda^n(S,U,V) \circ_U \chi_{T_{0,n},U}^{-1} \hat{Z}\big(C,\gamma'\big).$$

Note that the only copies of S^1 in the core of the above quantity are those corresponding to the link L', so that applying $\iota_k^{L'}$ just amounts to applying the map ι_k of Section 3.1.1. Finally, note that by our assumption on the link L, the linking matrix of $L' \cup O_n$ is just the linking matrix of the tangle $(T_{0,n} \cup L') \subset \gamma'$, so $\sigma_\pm^{L'\cup O_n} = \sigma_\pm(\gamma')$ as required. □

Remark 4.3. A similar formula holds in general for the invariant ∇_G of q-tangles in cobordisms over Σ. The only requirement is that such a tangle decomposes as the stacking of some q-tangle with an element of $\mathcal{T}(n)$, for some integer n (such as γ in Lemma 4.2). In this case, there is a formula similar to (13), but the Kontsevich integral is replaced with the AMR invariant V_G.

To illustrate, we give a diagrammatic version of the map J_G of §3.3.2, which allows us to express $\nabla_G(M)$ in terms of $\nabla_G^r(M)$ for a homology cylinder M. Recall that T_G denotes the tangle $I_G \cup (K_G \cap (1_\Sigma - C))$ in $1_\Sigma - C$, cf. (10). The subtangle $m_i \cap (1_\Sigma - C) \subset T_G$ is just a copy of the tangle $T_{0,h}$, and we have the formula

$$\nabla_G(M) = \nabla_G^r(M) \odot T_G \in \mathcal{A}_h, \tag{16}$$

for any homology cylinder M over Σ. Though it is defined more generally for cobordisms, this expresses our universal invariant ∇_G for homology cylinders in terms of LMO since $\nabla_G^r(M)$ can be computed in terms of the LMO invariant of a q-tangle in a homology ball as in Eq. (11).

4.3. *Diagrammatic formulas for the topological gluings*

In this section, we finally give the explicit formulas for the pairing, stacking and shelling products.

4.3.1. *Model for preferred structures*

We begin by choosing preferred marked bordered fatgraphs in each of the surfaces $\Sigma_{g,1}$ and $\Sigma_{0,g+1}$. The specified fatgraphs each have the property that the greedy algorithm produces a line segment as maximal tree; such "linear chord diagrams" are studied in [6]. The first, denoted $\overline{C}_g$, consists of g edges attached along the line interval $T_{\overline{C}_g}$, creating g isolated humps

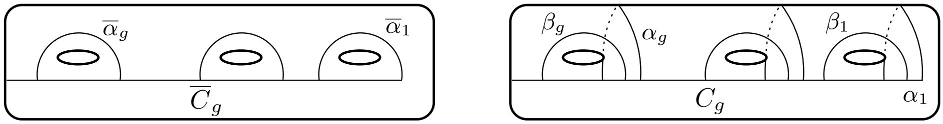

Fig. 12. Preferred marked fatgraphs $\overline{C}_g \hookrightarrow \Sigma_{0,g+1}$ and $C_g \hookrightarrow \Sigma_{g,1}$.

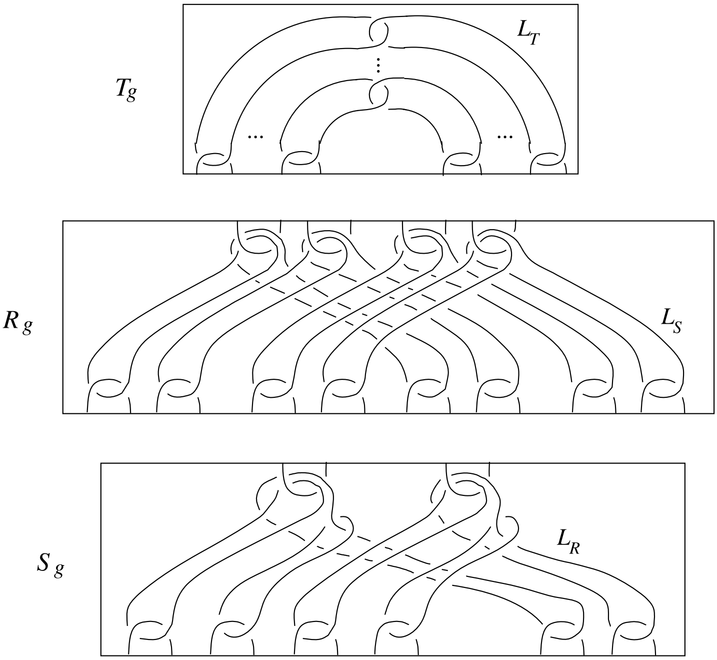

Fig. 13. The q-tangles T_g, S_g and R_g.

as shown in Fig. 12. The second fatgraph, which we call a genus g *symplectic fatgraph*[2] and denote by C_g, consists of $2g$ edges which appear along the interval T_{C_g} in g isolated overlapping pairs as illustrated in Fig. 12; see Figs. 14 and 15 for the skinny surfaces respectively associated to $\overline{C}_2$ and C_1. We choose the standard markings of C_g in $\Sigma_{g,1}$ and $\overline{C}_g$ in $\Sigma_{0,g+1}$ as shown in Fig. 12, where we have the identification of $\Sigma_g = \Sigma_{g,1} \cup D^2$ with the boundary of $H_g = \Sigma_{0,g+1} \times I$.

4.3.2. Diagrammatic pairing, stacking and shelling

Let $g \geqslant 1$ be an integer and define the three q-tangles in C

$$T_g = T_{0,2g} \cup L_T, \qquad S_g = T_{2g,4g} \cup L_S \quad \text{and} \quad R_g = T_{g,3g} \cup L_R,$$

where L_T, L_S and L_R are framed links as shown in Fig. 13.

[2] Note that this notation differs from that used in [7].

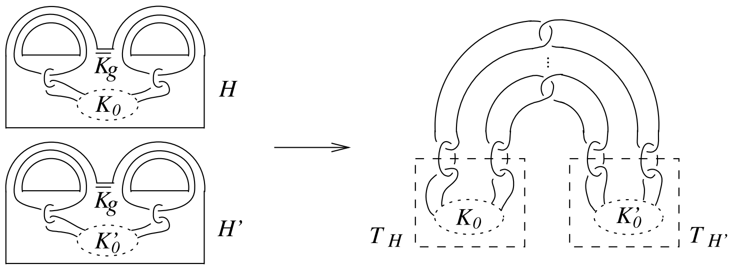

Fig. 14. Links for the pairing in the case $g = 2$.

Given Jacobi diagrams $D, D' \in \mathcal{A}_g$ and $E, E' \in \mathcal{A}_{2g}$, define

$$\langle , \rangle : \mathcal{A}_g \times \mathcal{A}_g \to \mathcal{A}(\emptyset),$$
$$\bullet : \mathcal{A}_{2g} \times \mathcal{A}_{2g} \to \mathcal{A}_{2g},$$
$$\star : \mathcal{A}_{2g} \times \mathcal{A}_g \to \mathcal{A}_g$$

by

$$\langle D, D' \rangle := (D \otimes D') \odot T_g,$$
$$E \bullet E' := (E \otimes E') \odot S_g,$$
$$E \star D := (E \otimes D) \odot R_g.$$

Theorem 4.4. *Let H and H' be two genus g homology handlebodies, and let M, M' be two homology cylinders over $\Sigma_{g,1}$. Then*

$$Z^{LMO}(H \cup_\iota H') = \langle \nabla^r_{\overline{C}_g}(H), \nabla^r_{\overline{C}_g}(H') \rangle, \tag{17}$$

$$\nabla^r_{C_g}(M \cdot M') = \nabla^r_{C_g}(M) \bullet \nabla^r_{C_g}(M'), \tag{18}$$

$$\nabla^r_{\overline{C}_g}(M * H) = \nabla^r_{C_g}(M) \star \nabla^r_{\overline{C}_g}(H). \tag{19}$$

Proof. Let K, K' be framed links in $1_{\Sigma_{0,g+1}}$ such that $H = (1_{\Sigma_{0,g+1}})_K$ and $H' = (1_{\Sigma_{0,g+1}})_{K'}$, and let L, L' be framed links in $1_{\Sigma_{g,1}}$ such that $M = (1_{\Sigma_{g,1}})_L$ and $M' = (1_{\Sigma_{g,1}})_{L'}$.

Denote by $\overline{K_g}$ and K_g, the system of linking pairs in rigid position in $1_{\Sigma_{0,g+1}}$ and $1_{\Sigma_{g,1}}$, respectively, induced by the preferred marked bordered fatgraphs $\overline{C}_g$ and C_g defined in Section 4.3.1. Let K_0, K'_0, L_0, and L'_0 be the reduced representatives of K, K', L and L' respectively with respect to the linking pairs $\overline{K_g}$ and K_g as provided by Lemma 2.2. See the left-hand side of Figs. 14 and 16. Note that surgery along these links in the preferred box always gives a homology ball since we are considering homology cylinders.

As to Eq. (17), it follows from straightforward Kirby calculus that the integral homology sphere $H \cup_\iota H'$ is obtained from S^3 by surgery along the framed link depicted on the right-hand side of Fig. 14. We see that this link can be decomposed as $(T_H \otimes T_{H'}) \cdot T_g$, where T_H and $T_{H'}$

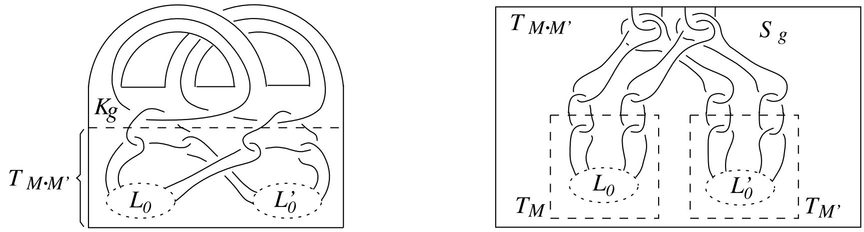

Fig. 15. Link and tangle for the stacking product in the case $g = 1$.

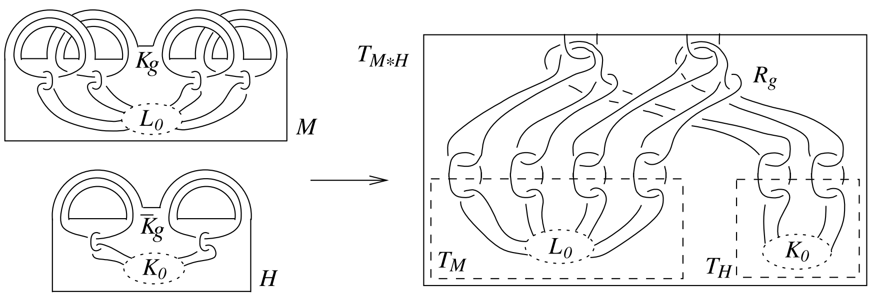

Fig. 16. Link and tangle for the shelling product in the case $g = 1$.

are the q-tangles in C defined in (10). The tangles $\gamma = T_H \otimes T_{H'}$ and $\gamma' = T_g$ indeed satisfy the hypotheses of Lemma 4.2, from which the result follows.

As to Eq. (18), the stacking product $M \cdot M'$ is obtained from $1_{\Sigma_{g,1}}$ by surgery along $L \cup L'$, where L and L' respectively occur in the lower and upper half of $1_{\Sigma_{g,1}}$. By Lemma 2.2, we can use the system of linking pairs K_g in rigid position to obtain a reduced representative, as shown in the left-hand side of Fig. 15.

Following (10), denote by $T_{M \cdot M'}$ the q-tangle in C obtained by cutting $1_{\Sigma_{g,1}}$ along the cutting face of the preferred box. This tangle, shown on the left-hand side of Fig. 15, is Kirby equivalent to the tangle shown on the right-hand side of the figure, which can be decomposed as $(T_M \otimes \check{T}_{M'}) \cdot S_g$. The result then follows from Lemma 4.2 with $\gamma = T_M \otimes \check{T}_{M'}$ and $\gamma' = S_g$.

Finally for Eq. (19), consider the link in $1_{\Sigma_{0,g+1}}$ obtained from L and K under the shelling product $1_{\Sigma_{g,1}} \star 1_{\Sigma_{0,g+1}}$. As in the previous case, we can use Lemma 2.2 and the system of linking pairs $\overline{K}_g$ in $1_{\Sigma_{0,g+1}}$ to obtain a reduced representative. One can check using Fig. 9 and Kirby calculus that the tangle T_{M*H} obtained by cutting $1_{\Sigma_{0,g+1}}$ along the cutting face of the preferred box is the tangle represented on the right-hand side of Fig. 16. Since the latter decomposes as $(T_M \otimes T_H) \cdot R_g$, and $\gamma = T_M \otimes T_H$ and $\gamma' = R_g$ satisfy the hypotheses of Lemma 4.2, the result follows. □

Remark 4.5. As mentioned earlier, both the stacking and shelling operations can be generalzed to the case of more than one boundary component. Versions of Theorem 4.4 can be established in these more general settings, by fixing appropriate analogues for the preferred marked fatgraph C_g and for the q-tangles R_g and S_g.

Remark 4.6. We can use (18) to define a multiplicative version of the rigid ∇_G^r invariant. This is done by renormalizing the invariant by a factor which uses the tangle R_g and a relative version of the contraction product $\odot$. A similar renormalization appears in Section 4.4 of [8].

5. Ptolemy groupoid representations

In this section, we exploit the dependence of our invariant ∇_G on the fatgraph G to construct representations of mapping class groups and their subgroups.

In Section 5.1, we discuss several actions of the mapping class group for general surfaces. However, throughout Section 5 we shall increasingly focus on the case of surfaces with one boundary component. While this restriction is primarily for convenience in Sections 5.2 and 5.3, it becomes necessary for technical reasons starting from Section 5.4.

5.1. Classical actions of subgroups of the mapping class group

The *mapping class group* $MC(\Sigma)$ of a compact orientable surface Σ, possibly with boundary $\partial\Sigma$ non-empty, is the group of isotopy classes relative to $\partial\Sigma$ of orientation-preserving self-diffeomorphisms of Σ which fix $\partial\Sigma$ pointwise. $MC(\Sigma)$ acts naturally on the integral homology groups of Σ, and we define the *Torelli group* $\mathcal{I}(\Sigma)$ of Σ to be the subgroup of $MC(\Sigma)$ acting trivially. Given a Lagrangian subspace $\Lambda \subset H_1(\Sigma_g;\mathbb{Q})$ for a closed surface Σ_g, we define the *Lagrangian preserving mapping class group* $MC(\Lambda) = \{\varphi \in MC(\Sigma_g)\colon \varphi_*(\Lambda) = \Lambda\}$. In particular, it is not difficult to see that the Torelli group is the intersection of all the Lagrangian preserving mapping class groups.

Consider the standard Heegaard decomposition of $S^3 = H_g \cup_\iota H_g$, where ι is the orientation-reversing involution of ∂H_g. Any mapping class $f \in \mathcal{I}(\Sigma_{g,1})$ gives rise to a corresponding mapping class $f \in \mathcal{I}(\Sigma_g)$ by capping off and extending by the identity. We may construct the homology sphere $S_f^3 = H_g \cup_{\iota\circ f} H_g$ by re-gluing the handlebodies using $\iota \circ f$. More generally for any Heegaard decomposition of a homology 3-sphere $M = H \cup_\iota H'$ into two genus g homology handlebodies, we obtain a similar map

$$f \mapsto M_f = H \cup_{\iota\circ f} H'. \tag{20}$$

Composing with the LMO invariant of the resulting homology 3-sphere M_f, we obtain a map $\mathcal{I}(\Sigma_{g,1}) \to \mathcal{A}(\emptyset)$, which is of some importance [25]. Similar maps can be defined on $\mathcal{I}(\Sigma_{g,n})$ for $n > 1$, but we shall only consider the case $n = 1$.

This kind of action of the Torelli group on the set of integral homology spheres with Heegaard splitting can equivalently be described in the context of homology cylinders via the mapping cylinder construction and the topological products described in Section 4.1.2. Indeed, the mapping cylinder of $\varphi \in MC(\Sigma)$, denoted $C(\varphi) = (1_\Sigma, \varphi, Id)$, is a special case of cobordism over Σ, and restricting to $\varphi \in \mathcal{I}(\Sigma)$, we obtain a homomorphism of monoids

$$\begin{aligned} \mathcal{I}(\Sigma) &\to \mathcal{HC}(\Sigma), \\ \varphi &\mapsto C(\varphi). \end{aligned} \tag{21}$$

Using this construction, we may reformulate (20) as $f \mapsto M_f = H \cup_\iota (C(f) * H')$, thus making precise the sense in which (21) describes an action on the set of integral homology spheres.

More generally, we can view the homomorphism (21) as an action of $\mathcal{I}(\Sigma)$ on the vector space generated by homology cylinders over Σ by stacking, i.e.,

$$M \mapsto M \cdot C(\varphi)$$

for $M \in \mathcal{HC}(\Sigma)$ and $\varphi \in \mathcal{I}(\Sigma)$. Similarly, we have the conjugation action

$$M \mapsto C(\varphi) \cdot M \cdot C(\varphi^{-1})$$

of $\varphi \in MC(\Sigma)$ on homology cylinders over Σ, where if $M = (1_\Sigma)_L$ is a homology cylinder over Σ, then

$$C(\varphi) \cdot (1_\Sigma)_L \cdot C(\varphi^{-1}) = C(\varphi) \cdot C(\varphi^{-1}) \cdot (1_\Sigma)_{\varphi^{-1}(L)} = (1_\Sigma)_{\varphi^{-1}(L)} \in \mathcal{HC}(\Sigma). \tag{22}$$

Analogously, we have a shelling action

$$H \mapsto C(\varphi) * H$$

of the Lagrangian preserving subgroup $MC(\Lambda)$ on the set $V(\Sigma_g, \Lambda)$ of genus g homology handlebodies with Lagrangian Λ.

Recall that the preferred marked bordered fatgraphs C_g and $\overline{C}_g$ defined in Section 4.3.1 induce isomorphisms

$$\nabla_{C_g} : \overline{\mathcal{H}}_{\Sigma_{g,1}} \xrightarrow{\cong} \mathcal{A}_{2g} \quad \text{and} \quad \nabla_{\overline{C}_g} : V(\Sigma_g, \Lambda^{st}) = \overline{\mathcal{H}}_{\Sigma_{0,g+1}} \xrightarrow{\cong} \mathcal{A}_g,$$

and using these, we thus obtain representations

$$\xi : MC(\Sigma_{g,1}) \to \mathrm{Aut}(\mathcal{A}_{2g}) \quad \text{and} \quad \zeta : MC(\Lambda^{st}) \to \mathrm{Aut}(\mathcal{A}_g),$$

respectively induced by conjugation and the shelling action. Analogues of the representation ξ for surfaces $\Sigma_{g,n}$ with $n > 1$ can be similarly defined once a choice of a preferred marked bordered fatgraph for the surface is made. This section relies on the fundamental relationship between fatgraphs and mapping class groups provided by the Ptolemy groupoid of decorated Teichmüller theory to describe these various actions in a purely combinatorial way.

5.2. *Ptolemy groupoid*

We shall restrict for convenience to surfaces with only one boundary component. However, note that all the definitions and results given in this section can be given in more general settings. Given a bordered fatgraph G, define the *Whitehead move* W on a non-tail edge e of the uni-trivalent fatgraph G to be the modification that collapses e to a vertex of valence four and then expands this vertex in the unique distinct way to produce the uni-trivalent fatgraph G'; see Fig. 17. We shall write either $W : G \to G'$ or $G \xrightarrow{W} G'$ under these circumstances.

Not only do markings of fatgraphs evolve in a natural way under Whitehead moves, so that we can unambiguously speak of Whitehead moves on marked fatgraphs, but also there is a natural identification of the edges of G and G'. Furthermore, there are three families of finite sequences of Whitehead moves, called the involutivity, commutativity, and pentagon relations, which leave invariant each marked fatgraph G, cf. [31,26].

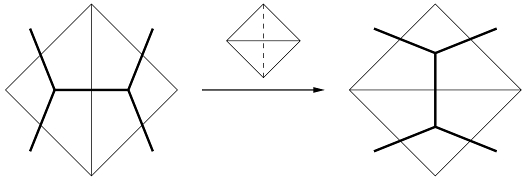

Fig. 17. A Whitehead move of a fatgraph superimposed on top of the diagonal exchange of its dual triangulation.

The *Ptolemy groupoid* $\mathfrak{Pt}(\Sigma)$ of a surface Σ with one boundary component is defined as the groupoid with objects given by marked bordered uni-trivalent fatgraphs $G \hookrightarrow \Sigma$ and morphisms given by sequences of Whitehead moves modulo the involutivity, commutativity, and pentagon relations.[3]

$\mathfrak{Pt}(\Sigma)$ provides a $MC(\Sigma)$-equivariant combinatorial model of the fundamental path groupoid of the decorated Teichmüller space of Σ, cf. [30–32]. As such, given any "point" in $\mathfrak{Pt}(\Sigma)$, i.e., any marked bordered fatgraph $G \hookrightarrow \Sigma$, each mapping class $\varphi \in MC(\Sigma)$ is represented by a unique morphism from G to $\varphi(G)$ in $\mathfrak{Pt}(\Sigma)$, where $\varphi(G)$ is the marked fatgraph that arises by postcomposing the marking $G \hookrightarrow \Sigma$ of G with φ.

For any marked bordered fatgraph $G \hookrightarrow \Sigma$, we may thus think of $MC(\Sigma)$ as being a set of equivalence classes of paths beginning at the point $G \hookrightarrow \Sigma$ and ending at a fatgraph combinatorially equivalent to G but potentially with a different marking in Σ. In this way, we get a presentation of the mapping class group of Σ:

Theorem 5.1. *(See [32].) For a surface Σ with one boundary component, the mapping class group $MC(\Sigma)$ has a presentation with generators given by sequences of Whitehead moves on marked fatgraphs in Σ beginning and ending at combinatorially isomorphic fatgraphs. The relations in this groupoid are given by identifying two sequences if they differ by a finite number of insertions or deletions of involutivity, commutativity, and pentagon relations.*

In a similar way, the Torelli group $\mathcal{I}(\Sigma)$ (and indeed each term of the Johnson filtration [16]) likewise admits an analogous combinatorial presentation as in [26].

By a *representation* $\mathfrak{Pt}(\Sigma) \to K$ of the Ptolemy groupoid in some group K, we mean a composition-preserving map $\mathrm{Mor}(\mathfrak{Pt}(\Sigma)) \to K$ from the morphisms of $\mathfrak{Pt}(\Sigma)$. In other words, a representation of $\mathfrak{Pt}(\Sigma)$ is a morphism that assigns an element of K to each Whitehead move such that the composition is trivial for the involutivity, commutativity, and pentagon relations.

5.3. The explicit Ptolemy action on $\mathcal{A}_h$

Here, as in Section 5.2, we restrict to the case $\Sigma = \Sigma_{g,1}$ of a surface with one boundary component, so that $h = 2g$. Note however that an analogue of Theorem 5.2 holds for surfaces with more than one boundary component.

Our first representation of $\mathfrak{Pt}(\Sigma)$ captures the dependence of ∇_G on the choice of marked bordered fatgraph $G \hookrightarrow \Sigma$ giving a representation as automorphisms of an appropriate space of

[3] The term "Ptolemy groupoid" is sometimes used to refer to the groupoid whose objects are $MC(\Sigma)$-orbits of uni-trivalent fatgraphs and whose morphisms are $MC(\Sigma)$-orbits of pairs of such with the natural composition. We prefer to call this the *mapping class groupoid* since it gives a combinatorial model for the fundamental path groupoid of Riemann's moduli space.

Jacobi diagrams extending the conjugation action of the mapping class group on homology cylinders. Recall from Corollary 3.3 that for a genus g surface Σ with $n \geqslant 1$ boundary components, any marked bordered fatgraph $G \hookrightarrow \Sigma$ provides a graded isomorphism

$$\nabla_G : \overline{\mathcal{H}}_\Sigma \xrightarrow{\cong} \mathcal{A}_h,$$

where $h = 2g + n - 1$, as a consequence of the universality of the invariant ∇_G. Thus, for any marked bordered fatgraphs G and G', we get isomorphisms ∇_G and $\nabla_{G'}$ of $\overline{\mathcal{H}}_\Sigma$ with $\mathcal{A}_h$. As a formal consequence, we obtain an explicit representation of the Ptolemy groupoid:

Theorem 5.2. *The map*

$$\bigl(G \xrightarrow{W} G'\bigr) \mapsto \nabla_{G'} \circ \nabla_G^{-1}$$

defines a representation of the Ptolemy groupoid acting on $\mathcal{A}_h$

$$\hat{\xi} : \mathfrak{Pt}(\Sigma) \to \mathrm{Aut}(\mathcal{A}_h),$$

which extends the representation $\xi : MC(\Sigma) \to \mathrm{Aut}(\mathcal{A}_h)$ *induced by the conjugation action in the sense that for any sequence of Whitehead moves* $C_g \xrightarrow{W_1} \cdots \xrightarrow{W_k} \varphi(C_g)$ *representing* $\varphi \in MC(\Sigma)$*, we have the identity* $\hat{\xi}(W_1) \circ \cdots \circ \hat{\xi}(W_k) = \xi(\varphi)$.

Before giving the proof, we first give the following topological interpretation of the automorphism associated to a morphism from G to G' in $\mathfrak{Pt}(\Sigma)$. Given an element in $\mathcal{A}_h$, we can pull it back via ∇_G^{-1} to an element of $\overline{\mathcal{H}}_\Sigma$, represented by a formal series L of framed links in 1_Σ in admissible position with respect to the polygonal decomposition P_G. We then evolve G by a sequence of Whitehead moves to a new marked fatgraph G', and isotope the links in L accordingly to put them in admissible position with respect to the new polygonal decomposition $P_{G'}$. Evaluating $\nabla_{G'}$ on the resulting series of links then provides a new element of $\mathcal{A}_h$.

Proof of Theorem 2.2. The fact that the above action defines a representation of the Ptolemy groupoid follows easily since any sequence of Whitehead moves representing a trivial morphism of $\mathfrak{Pt}(\Sigma)$ begins and ends at identical marked fatgraphs and thus must give the trivial action.

For any link $L \subset 1_\Sigma$, we have $\nabla_n^G(L) = \nabla_n^{\varphi(G)}(\varphi(L))$ by construction, so that

$$\nabla_{\varphi(G)}\bigl((1_\Sigma)_L\bigr) = \nabla_G\bigl((1_\Sigma)_{\varphi^{-1}(L)}\bigr).$$

Thus by (22) for any $M \in \overline{\mathcal{H}}_\Sigma$, we have

$$\xi(\varphi)\bigl(\nabla_G(M)\bigr) = \nabla_G\bigl(C(\varphi) \cdot M \cdot C\bigl(\varphi^{-1}\bigr)\bigr) = \nabla_{\varphi(G)}(M),$$

and setting $G = C_g$, the result follows. □

5.4. The Ptolemy groupoid action on handlebodies

A similar extension of the shelling action ζ arises as follows. By Lemma 7.4 of [1], there is an algorithm which produces a representation

$$\mathfrak{Pt}(\Sigma_{g,1}) \to MC\big(\Lambda^{st}\big)$$

of the Ptolemy groupoid extending the identity on $MC(\Lambda^{st})$. Thus, we obtain:

Proposition 5.3. *Let Σ_g be a closed genus g surface. Fix a disc in Σ_g and let $\Sigma_{g,1}$ be its complement. Then we have an explicit algorithmically defined representation*

$$\hat{\zeta}: \mathfrak{Pt}(\Sigma_{g,1}) \to \mathrm{Aut}(\mathcal{A}_g)$$

which extends the shelling action $\zeta : MC(\Lambda^{st}) \to \mathrm{Aut}(\mathcal{A}_g)$.

Owing to its dependence on the complicated algorithms in [1], the action on $\mathcal{A}_g$ obtained in this way is more complicated than the action on $\mathcal{A}_h$ described in the previous section.

5.5. Extension of the LMO invariant to the Ptolemy groupoid

In this section, we give a kind of Ptolemy groupoid action on finite type invariants of integral homology spheres which extends the usual action of the Torelli group via Heegaard decomposition. It is enough for our purpose to restrict ourselves to the case of a once-bordered surface. More precisely, we give a Ptolemy groupoid action on finite type invariants of homology cylinders over $\Sigma_{g,1}$ which extends the stacking action of $\mathcal{I}(\Sigma_{g,1})$ on $\overline{\mathcal{H}}_{\Sigma_{g,1}}$, and which induces a map from $\mathfrak{Pt}(\Sigma_{g,1})$ to $\mathcal{A}(\emptyset)$ extending the analogous map of the Torelli group $\mathcal{I}(\Sigma_{g,1})$.

We begin by recalling Corollary 7.1 and Theorem 8.1 of [1], which together give[4] that any choice of a marked bordered fatgraph $G \hookrightarrow \Sigma_{g,1}$ determines a representation

$$\hat{id}_G : \mathfrak{Pt}(\Sigma_{g,1}) \to \mathcal{I}(\Sigma_{g,1})$$

of the Ptolemy groupoid which extends the identity homomorphism of $\mathcal{I}(\Sigma_{g,1})$. Let $\hat{id} = \hat{id}_{C_g}$ be the representation provided by the marked fatgraph $C_g \hookrightarrow \Sigma_{g,1}$ defined in Section 4.3.1. Define a representation of the Ptolemy groupoid of $\Sigma_{g,1}$

$$\rho : \mathfrak{Pt}(\Sigma_{g,1}) \to \mathcal{A}_{2g},$$

to be a composition-preserving map, where the target space is imbued with the stacking product $\bullet$, by setting

$$\rho(W) := \nabla^r_{C_g}\big(C\big(\hat{id}(W)\big)\big).$$

[4] The proof in [1] relies on a sequence of algorithms, beginning with the greedy algorithm, which produces a sequence of Whitehead moves taking a given fatgraph to a symplectic one, followed by an algorithm which manipulates the homological information associated to each edge of the symplectic fatgraph; this last algorithm apparently has a paradigm in K-theory.

Theorem 5.4. *The representation $\rho : \mathfrak{Pt}(\Sigma_{g,1}) \to \mathcal{A}_{2g}$ of the Ptolemy groupoid of $\Sigma_{g,1}$ provides an extension of the LMO invariant of integral homology spheres to the Ptolemy groupoid in the following sense: Let $f \in \mathcal{I}_{g,1}$ and let*

$$G \xrightarrow{W_1} G_1 \xrightarrow{W_2} \cdots \xrightarrow{W_k} G_k = f(G)$$

be a sequence of Whitehead moves representing f. Let $M = H \cup_\iota H'$ be a genus g Heegaard splitting of an integral homology sphere M. Then the LMO invariant of the integral homology 3-sphere $M_f = H \cup_{\iota \circ f} H'$ is given by

$$Z^{LMO}(M_f) = \big\langle v, \big(\rho(W_1) \bullet \rho(W_2) \bullet \cdots \bullet \rho(W_k)\big) \star v' \big\rangle,$$

where $v = \nabla^r_{\overline{C}_g}(H) \in \mathcal{A}_g$ and $v' = \nabla^r_{\overline{C}_g}(H') \in \mathcal{A}_g$.

Proof. Since $\hat{id}$ extends the identity homomorphism of $\mathcal{I}(\Sigma_{g,1})$, we therefore have $\hat{id}(W_1) \circ \hat{id}(W_2) \circ \cdots \circ \hat{id}(W_k) = f$, hence $M_{W_1} \cdot M_{W_2} \cdot \cdots \cdot M_{W_k} = C(f)$, where M_W denotes $C(\hat{id}(W))$. Since $M_f = H \cup_\iota (C(f) * H')$, the formula follows from Theorem 4.4. $\square$

Considering the map $f \mapsto S^3_f$ induced by the standard Heegaard decomposition of S^3, Theorem 5.4 shows that for a sequence $G \xrightarrow{W_1} G_1 \xrightarrow{W_2} \cdots \xrightarrow{W_k} G_k = f(G)$ of Whitehead moves representing f, we have

$$Z^{LMO}\big((S^3)_f\big) = \big\langle v_0, \big(\rho(W_1) \bullet \rho(W_2) \bullet \cdots \bullet \rho(W_k)\big) \star v_0 \big\rangle,$$

where the diagrammatic constant $v_0 = \nabla^r_{\overline{C}_g}(H_g) \in \mathcal{A}_g$ can easily be computed as follows. By definition, we have $\nabla^r_{\overline{C}_g}(H_g) = Z^{LMO}(C, T_{g,0})$, where $T_{g,0}$ is the q-tangle of Fig. 10. By (1), the Kontsevich integral $Z(T_{g,0}) \in \mathcal{A}(\uparrow^g)$ of this tangle is thus given by including a $\sqrt{\nu}$ on each copy of $\uparrow$. It follows that

$$v_0 = \bigsqcup_{i=1}^{g} \big(\chi^{-1}_{\{i\}} \sqrt{\nu}\big) \in \mathcal{A}_g,$$

where an explicit formula for ν is given in [5].

5.6. *Extension of the first Johnson homomorphism*

In [26], a representation of the Ptolemy groupoid was introduced using the notion of an H-marking of a fatgraph G and shown to be an extension of the first Johnson homomorphism τ_1 to the Ptolemy groupoid. In this section, we show how a variation of the invariant ∇_G can be used to realize this extension.

5.6.1. General latches

Let $G \hookrightarrow \Sigma$ be a marked bordered fatgraph in a surface Σ. We begin by introducing a generalized notion of a system of latches in 1_Σ and thus of the invariant ∇_G. In fact, the main property of the system of latches I_G we used so far in this paper, besides the fact that it is determined by the fatgraph G, is that it provides a dual basis in homology.

We define a general latch for G as an embedded interval in the boundary of $\Sigma \times I$ with endpoints lying in $(\partial\Sigma) \times \{\frac{1}{2}\}$ such that it can be isotoped relative to its boundary to be in admissible position with respect to the polygonal decomposition P_G. A collection of h disjoint latches in the boundary of 1_Σ whose homotopy class relative to the boundary induces a free basis for $H_1(1_\Sigma, \partial 1_\Sigma; \mathbb{Q})$ is a *general system of latches* for G.

It is clear that substituting for I_G in (6) any general system of latches yields an invariant of cobordisms. In fact, such an invariant is also universal for homology cylinders. We shall not make use of this result and omit the proof, which essentially follows Section 3.2.2 (the main difference being in the definition of the surgery map).

5.6.2. Extending τ_1 via the invariant ∇

We restrict our attention to the once-bordered surface $\Sigma = \Sigma_{g,1}$ of genus g, set $H = H_1(\Sigma_{g,1}; \mathbb{Z})$ and $H_\mathbb{Q} = H \otimes \mathbb{Q}$. Recall [16] that the first Johnson homomorphism

$$\tau_1 : \mathcal{I}(\Sigma_{g,1}) \to \Lambda^3 H$$

takes its values in the third exterior power of H.

Denote by I_g the $2g$-component q-tangle in $\Sigma_{g,1} \times I$ represented below:

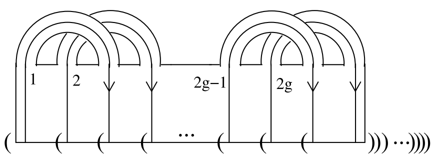

Note that by isotoping I_g so that it is contained in $(\Sigma_{g,1} \times \{1\}) \cup (\partial\Sigma_{g,1} \times I)$, we may consider I_g as a general system of latches for any choice of marked bordered fatgraph G in $\Sigma_{g,1}$. Indeed, one can unambiguously arrange the endpoints of I_g so that under the projection of $\Sigma_{g,1} \times I$ to $\Sigma_{g,1}$ they lie in a neighborhood of the fixed point q where the tail of G is attached, so that ∂I_g lies on the boundary of the preferred box in the polygonal decomposition P_G.

Let G be a marked bordered fatgraph in $\Sigma_{g,1}$ and let L be a framed link in $\Sigma_{g,1} \times I$ which is disjoint from both K_G and I_g. Set

$$\nabla_n^{G,I_g}(L) := \frac{\iota_n(\check{V}_G(L \cup K_G \cup I_g))}{\iota_n(\check{V}_G(U_+))^{\sigma_+^{L\cup K_G}}\,\iota_n(\check{V}_G(U_-))^{\sigma_-^{L\cup K_G}}} \in \mathcal{A}_{\leqslant n}\big(\uparrow^{2g}\big), \tag{23}$$

where we make use of the notation of (6). This quantity is an invariant of the surgered manifold $M = (\Sigma_{g,1} \times I)_L$, and following (7), we set

$$\nabla_G^{I_g}(M) := 1 + \big(\nabla_1^{G,I_g}(L)\big)_1 + \cdots + \big(\nabla_n^{G,I_g}(L)\big)_n + \cdots \in \mathcal{A}\big(\uparrow^{2g}\big).$$

Next, consider a Whitehead move $W: G \mapsto G'$. We can then compare the value of the invariants $\nabla_G^{I_g}$ and $\nabla_{G'}^{I_g}$ on the trivial element $1_{\Sigma_{g,1}}$ and assign the quotient to the Whitehead move W to define a map

$$\mathcal{J}(W) = \nabla_{G'}^{I_g}(1_{\Sigma_{g,1}})/\nabla_G^{I_g}(1_{\Sigma_{g,1}}) \in \mathcal{A}(\uparrow^{2g}).$$

More generally, for any two marked fatgraphs in $\Sigma_{g,1}$, not necessarily related by a Whitehead move, we can similarly take the quotient, and in the case that these two fatgraphs are equal, we get a trivial contribution by definition. This guarantees that this map $\mathcal{J}$ is the identity for the involutivity, commutativity, and pentagon relations, and hence we obtain a representation

$$\mathcal{J}: \mathfrak{Pt}(\Sigma_{g,1}) \to \mathcal{A}(\uparrow^{2g}).$$

Recall that the groups $H_1(\Sigma_{g,1};\mathbb{Q})$ and $H_1(\Sigma_{g,1},\partial\Sigma_{g,1};\mathbb{Q})$ are isomorphic via Poincaré duality. Define a map $\mathfrak{h}:\{1,\dots,2g\} \to H$ by taking i to the element of H dual to the class of the ith component of I_g in $H_1(\Sigma_{g,1},\partial\Sigma_{g,1};\mathbb{Z})$. More concretely, if we let $\{A_i, B_i\}_{i=1}^{2g}$ denote the standard symplectic basis of $\Sigma_{g,1}$ with $A_i \cdot B_j = \delta_{ij}$, then

$$\mathfrak{h}(2k) = A_k, \qquad \mathfrak{h}(2k-1) = B_k \quad \text{for } k = 1,\dots,g.$$

Also recall that $\mathcal{B}_1^{\mathsf{Y}}(2g) = \mathcal{B}_{1,3}^{\mathsf{Y}}(2g)$ is the space of $2g$-colored Y-shaped Jacobi diagrams and that we have the well-known and elementary isomorphism $\mathcal{B}_1^{\mathsf{Y}}(2g) \cong \Lambda^3 H_{\mathbb{Q}}$ defined by sending a Y-shaped diagram colored by i, j, k (following the vertex-orientation) to $\mathfrak{h}(i) \wedge \mathfrak{h}(j) \wedge \mathfrak{h}(k) \in \Lambda^3 H$.

In order to extend the first Johnson homomorphism τ_1, we restrict the target of our representation $\mathcal{J}$ by composing it with the series of maps given by

$$Y: \mathcal{A}(\uparrow^{2g}) \to \mathcal{B}(2g) \to \mathcal{B}^{\mathsf{Y}}(2g) \to \mathcal{B}_1^{\mathsf{Y}}(2g) \cong \Lambda^3 H_{\mathbb{Q}}. \tag{24}$$

From this, we obtain a representation of the Ptolemy groupoid

$$\mathcal{J}^{\mathsf{Y}}: \mathfrak{Pt}(\Sigma_{g,1}) \to \Lambda^3 H_{\mathbb{Q}}.$$

The first map in (24) is the inverse χ^{-1} of the Poincaré–Birkhoff–Witt isomorphism, and the second and third maps are the natural projections.

Theorem 5.5. *The representation $\mathcal{J}^{\mathsf{Y}}$ extends the first Johnson homomorphism τ_1 to the Ptolemy groupoid. More precisely, given a sequence*

$$G \xrightarrow{W_1} G_1 \xrightarrow{W_2} \cdots \xrightarrow{W_k} G_k = \varphi(G)$$

of Whitehead moves representing $\varphi \in \mathcal{I}_{g,1}$, we have $\tau_1(\varphi) = 4\sum_{i=1}^{k} \mathcal{J}^{\mathsf{Y}}(W_i)$.

5.6.3. Proof of Theorem 5.5

The computation of the invariant $\mathcal{J}^{\mathsf{Y}}$ is considerably simplified by the following observation.

Lemma 5.6. *For any marked bordered fatgraph G in $\Sigma_{g,1}$, we have*

$$Y\big(\nabla_G^{I_g}(1_{\Sigma_{g,1}})\big) = Y\big(\check{V}_G(I_g)\big) \in \Lambda^3 H_{\mathbb{Q}},$$

where Y is the sequence of maps in (24).

In other words, the Y-shaped part of $\nabla_G^{I_g}(1_{\Sigma_{g,1}})$ comes purely from the tangle I_g, and the system of linking pairs K_G can simply be ignored in the computation.

Proof of Lemma 5.6. We shall freely make use of the terminology introduced in the proof of Theorem 3.2. In computing $\nabla_G^{I_g}(1_{\Sigma_{g,1}})$, we can choose L to be empty in (23). By [29, pp. 283], we have that $\iota_2(\check{Z}(U_\pm)) = 1+$terms of i-degree $\geqslant 2$, and it follows that

$$Y\big(\nabla_G^{I_g}(1_{\Sigma_{g,1}})\big) = Y\big(\iota_2\big(\check{V}_G(K_G \cup I_g)\big)\big).$$

We now consider the linking pairs K_G. We may assume that there are $2g$ disjoint 3-balls in $1_{\Sigma_{g,1}}$ that intersect the system K_G of linking pairs as illustrated on the right-hand side of Fig. 6. The Kontsevich integral of the tangle contained in these balls is computed in [5, Theorem 4], from which it follows that the only terms in $\check{V}_G(K_G \cup I_g)$ that can contribute to $Y(\nabla_G^{I_g}(1_{\Sigma_{g,1}}))$ have exactly 4 vertices on each meridian core, which are the ends of 4 parallel struts connecting each to the corresponding longitude core.

Suppose that some longitude core has k additional vertices attached. It follows from the definition that applying the map ι_2 produces $\geqslant k$ univalent vertices, which imposes the constraint that $k \leqslant 1$. For $k = 1$, the diagram is also sent to zero by the map ι_2 since we obtain a sum of Jacobi diagrams each having a looped edge, which vanish by the AS relation. Thus, the only terms which can possibly contribute are Siamese diagrams with 4 struts, cf. Fig. 7, which come with a coefficient $\frac{1}{4!}$. As seen in Section 3.2.4, ι_2 maps each Siamese diagram to a factor $(-1)^2 4!$ as required. □

We can now proceed with the proof of Theorem 5.5 and calculate the representation $\mathcal{J}^{\mathsf{Y}}$ on a Whitehead move W. To this end, for any marked fatgraph G in $\Sigma_{g,1}$, we can assume by Lemma 2.1 that the q-tangle I_g is in admissible position and intersects each box except the preferred one in a trivial q-tangle. For each oriented edge of G, we may equip each strand of the trivial q-tangle in the corresponding box with a sign, according to whether its orientation agrees (plus sign) or disagrees (minus sign) with the specified one. For each oriented edge of G, assign an element of H to each box except the preferred one as follows: use the map $\mathfrak{h}$ to label all the strands of I_g intersecting the box by elements of the symplectic basis $\{A_i, B_i\}_{i=1}^{2g}$ and take the signed sum of these labels in H. We remark that this assignment is precisely an H-marking as described in [26,7].

Thus, we have a situation as in the upper part of Fig. 18, where each of the three strands depicted there represents a collection of parallel strands of $I_g \cup K_G$ and where $A, B, C \in H$ are the labels of the box as just explained. Note that the bracketing $(C, (B, A))$ in the bottom-left box is imposed by the condition on hexagons, see Section 2.4. After the Whitehead move, we have

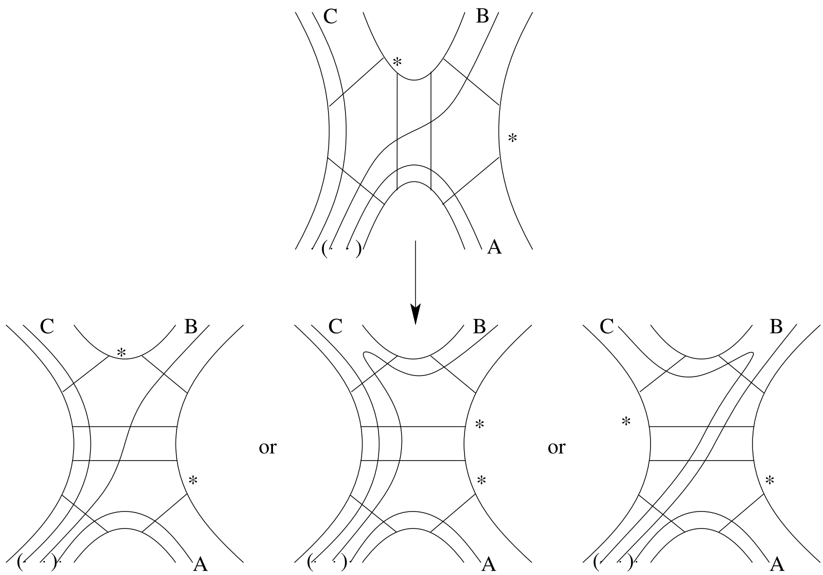

Fig. 18. The three possible evolutions of the forbidden sectors under a Whitehead move.

one of the three situations represented in the lower part of Fig. 18 depending on the ordering of the sectors associated to the edge on which the move has been performed. In each case, we see that the bracketing of the three strands in the bottom left box is changed to $((C, B), A)$. Also, in the last two cases, we get an extra cap or cup due to the evolution of the forbidden sectors, and these are the only changes; in particular, there are no crossing changes.

It follows from the computation [5] of ν that a cup or a cap cannot contribute to $\mathcal{J}^{\mathsf{Y}}$, so in all three cases, we get the same value for $\mathcal{J}^{\mathsf{Y}}(W)$ coming from the evolution in the bracketing, i.e., from the associator. Recall that an even associator is always of the form

$$\Phi = 1 + \frac{1}{24} \; + \text{terms of } J\text{-degree } > 3. \tag{25}$$

Furthermore, the value of the Kontsevich integral on a q-tangle obtained from Λ_+ (see Fig. 2) by taking parallel copies of each strand with arbitrary orientation is obtained from Φ by the comultiplication and antipode maps defined in Section 2.1.2.

By Lemma 5.6 and Eq. (25), we find $\mathcal{J}^{\mathsf{Y}}(W) = \frac{1}{24} A \wedge B \wedge C \in \Lambda^3 H_{\mathbb{Q}}$, and this formula coincides with one fourth of the Morita–Penner extension of the first Johnson homomorphism τ_1. The fact that it indeed is a multiple of an extension of τ_1 follows as in [26], upon which our determination of the factor 4 currently relies, and completes the proof of Theorem 5.5.

6. Concluding remarks and questions

There are several obvious questions regarding the Ptolemy representations derived in Section 5. Most notably, one may ask for a geometric interpretation of the mapping class group action arising from the representation $\mathcal{J}$. In particular, does $\mathcal{J}$ provide an extension of the full LMO invariant in the same TQFT spirit as in Theorem 5.4?

Also, a natural and interesting issue is the faithfulness of the action of the mapping class group $MC(\Sigma_{g,1})$ on $\mathcal{A}_{2g}$ induced by $\mathcal{J}$. As the groupoid formulas for these representations seem simpler to analyze than their corresponding mapping class group expressions, our techniques here may prove pivotal in providing such an answer. The facts that the pronilpotent representation of an automorphism of a free group is faithful and that the Johnson theory presumably corresponds to the tree-like part of LMO by [12,14,23] together suggest that the induced representation of the mapping class group may be faithful.

6.1. Magnus expansions and Johnson homomorphisms

In recent beautiful work, Gwénaël Massuyeau [23] has introduced the notion of symplectic Magnus expansions and proved their existence by giving explicit formulas in terms of the LMO invariant. Nariya Kawazumi [17] has asked the interesting question if such Magnus expansions might be computed directly in terms of suitably marked fatgraphs as in [7]. Our computation here of the LMO invariant provides such a formula but again a very complicated one. Moreover, it seems likely that a construction analogous to Massuyeau's using our invariants ∇_G or $\nabla_G^{I_g}$ will lead to a directly computable version, and it would be an interesting prospect to derive formulas for the various Johnson homomorphisms in terms of such a symplectic Magnus expansion.

6.2. Relation to triangulations of 3-manifolds

The dual in a surface Σ of a marked uni-trivalent fatgraph G is a triangulation Δ_G of the surface Σ, where a k-valent vertex of G gives rise to a $2k$-gon whose alternating sides correspond to incident half-edges and whose complementary sides correspond to arcs in the boundary, cf. [30, 32]. The dual of a Whitehead move on a uni-trivalent fatgraph corresponds to a diagonal exchange on its dual ideal triangulation as illustrated in Fig. 17. We may imagine this diagonal flip as exchanging the front and the back pair of faces of a tetrahedron in the obvious way. It is thus natural to regard a morphism in the Ptolemy groupoid as a sequence of adjoined tetrahedra starting from the corresponding fixed ideal triangulation Δ of the surface, i.e., a morphism provides a triangulated cobordism between one copy of the surface with triangulation Δ and another copy of the surface with potentially another triangulation. This is especially natural for a mapping cylinder, where the Ptolemy morphism connects Δ to its image under the corresponding mapping class; this has indeed been the point of view in [26,7].

Conversely, suppose that we have ideal triangulations of two bordered surfaces Σ and Σ' and suppose that M is a 3-manifold whose boundary contains $\Sigma \sqcup \Sigma'$. We may ask for a triangulation of M extending those given on the boundary all of whose vertices lie in $\Sigma \sqcup \Sigma'$. In the spirit of a TQFT, we are led to the following questions. Do finite compositions of Whitehead moves acting as before on triangulated cobordisms in fact act transitively on such triangulations of M? Which 3-manifold invariants can be computed that depend upon the ideal triangulations of $\Sigma \sqcup \Sigma'$ but not the triangulation of M? What type of state-sum model corresponds to this?

6.3. The original AMR invariant

As pointed out in Section 2.5, the AMR invariant V_G employed in the construction of ∇_G is actually only a weak version of the one in [2]. Indeed, we are post-composing the original invariant with the map that forgets the homotopy class of chord diagrams on surfaces. It is a natural and important problem to try to build a 3-manifold invariant from the full Andersen–Mattes–Reshetikhin invariant that would retain this homotopy information and thus non-trivially extend finite type invariants to all 3-manifolds. We shall return to this study in a forthcoming paper, where we also discuss how constructions inspired by those in this paper can be used to define universal perturbative invariants of closed 3-manifolds and more generally universal perturbative TQFTs.

Finally note that in the proof of Theorem 5.5, computations were made amenable by Lemma 5.6 in avoiding the complex maps ι_n from LMO theory, i.e., our calculation of τ_1 is performed at the "AMR level" rather than at the "LMO level", cf. [14]. The AMR-valued version of our invariant, or its homotopy analogue just discussed, may be suited to other explicit computations as well. Indeed, the original AMR invariant provides a graded isomorphism between the Vassiliev-filtered free vector space generated by links in the cylinder over a surface with a non-empty boundary and the algebra of chord diagrams on the surface [2], and this isomorphism is determined once a suitable fatgraph is chosen in the surface as discussed here. We therefore get an action of the Ptolemy groupoid on the algebra of chord diagrams on any surface with non-empty boundary just as in Section 5.3. We shall study this representation of the Ptolemy groupoid in a forthcoming publication.

Acknowledgment

The authors are indebted to Gwénaël Massuyeau for seminal discussions at the beginning and helpful ones at the end of this project.

References

[1] J.E. Andersen, A.J. Bene, R.C. Penner, Groupoid extensions of mapping class representations for bordered surfaces, Topology Appl. 156 (17) (2009) 2713–2725.

[2] J.E. Andersen, J. Mattes, N. Reshetikhin, Quantization of the algebra of chord diagrams, Math. Proc. Cambridge Philos. Soc. 124 (1998) 451–467.

[3] D. Bar-Natan, On the Vassiliev knot invariants, Topology 34 (2) (1995) 423–472.

[4] D. Bar-Natan, S. Garoufalidis, L. Rozansky, D.P. Thurston, The Århus integral of rational homology 3-spheres II. Invariance and universality, Selecta Math. (N.S.) 8 (3) (2002) 341–371.

[5] D. Bar-Natan, T.T.Q. Le, D.P. Thurston, Two applications of elementary knot theory to Lie algebras and Vassiliev invariants, Geom. Topol. 7 (2003) 1–31.

[6] A.J. Bene, A chord diagrammatic presentation of the mapping class group of a once bordered surface, Geom. Dedicata 144 (2010) 171–190.

[7] A.J. Bene, N. Kawazumi, R.C. Penner, Canonical extensions of the Johnson homomorphisms to the Torelli groupoid, Adv. Math. 221 (2009) 627–659.

[8] D. Cheptea, K. Habiro, G. Massuyeau, A functorial LMO invariant for Lagrangian cobordisms, Geom. Topol. 12 (2) (2008) 1091–1170.

[9] D. Cheptea, T.T.Q. Le, A TQFT associated to the LMO invariant of three-dimensional manifolds, Comm. Math. Phys. 272 (3) (2007) 601–634.

[10] S. Garoufalidis, The mystery of the brane relation, J. Knot Theory Ramifications 11 (5) (2002) 725–737.

[11] S. Garoufalidis, M. Goussarov, M. Polyak, Calculus of clovers and finite type invariants of 3-manifolds, Geom. Topol. 5 (2001) 75–108.

[12] S. Garoufalidis, J. Levine, Tree-level invariants of three-manifolds, Massey products and the Johnson homomorphism, in: Graphs and Patterns in Mathematics and Theoretical Physics, in: Proc. Sympos. Pure Math., vol. 3, Amer. Math. Soc., Providence, RI, 2005, pp. 173–203.

[13] M. Goussarov, Finite type invariants and n-equivalence of 3-manifolds, C. R. Math. Acad. Sci. Paris 329 (Série I) (1999) 517–522.

[14] N. Habegger, Milnor, Johnson and tree-level perturbative invariants, preprint, 2000.

[15] K. Habiro, Claspers and finite type invariants of links, Geom. Topol. 4 (2000) 1–83.
[16] D. Johnson, An Abelian quotient of the mapping class group $\mathcal{T}_g$, Math. Ann. 249 (1980) 225–242.
[17] N. Kawazumi, private communication.
[18] T.T.Q. Le, An invariant of integral homology 3-spheres which is universal for all finite type invariants, in: V. Buchstaber, S. Novikov (Eds.), Solitons, Geometry and Topology: On the Crossroad, in: Amer. Math. Soc. Transl. Ser. 2, vol. 179, Amer. Math. Soc., Providance, RI, 1997, pp. 75–100.
[19] T.T.Q. Le, On denominators of the Kontsevich integral and the universal perturbative invariant of 3-manifolds, Invent. Math. 135 (1999) 689–722.
[20] T.T.Q. Le, J. Murakami, Parallel version of the universal Vassiliev–Kontsevich invariant, J. Pure Appl. Algebra 121 (3) (1997) 271–291.
[21] T.T.Q. Le, J. Murakami, T. Ohtsuki, On a universal perturbative quantum invariant of 3-manifolds, Topology 37 (1998) 539–574.
[22] J. Levine, Homology cylinders: an enlargement of the mapping class group, Algebr. Geom. Topol. 1 (2001) 243–270.
[23] G. Massuyeau, Infinitesimal Morita homomorphisms and the tree-level of the LMO invariant, preprint, 2008, arXiv:0809.4629.
[24] G. Massuyeau, J.B. Meilhan, Characterization of Y_2-equivalence for homology cylinders, J. Knot Theory Ramifications 12 (4) (2003) 493–522.
[25] S. Morita, Casson's invariant for homology 3-spheres and characteristic classes of surface bundles I, Topology 28 (1989) 305–323.
[26] S. Morita, R.C. Penner, Torelli groups, extended Johnson homomorphisms, and new cycles on the moduli space of curves, Math. Proc. Cambridge Philos. Soc. 144 (3) (2008) 651–671.
[27] J. Murakami, T. Ohtsuki, Topological quantum field theory for the universal quantum invariant, Comm. Math. Phys. 188 (3) (1997) 501–520.
[28] T. Ohtsuki, Finite type invariants of integral homology spheres, J. Knot Theory Ramifications 5 (1996) 101–115.
[29] T. Ohtsuki, Quantum Invariants. A Study of Knots, 3-Manifolds, and Their Sets, Ser. Knots Everything, vol. 29, World Scientific, 2002.
[30] R.C. Penner, The decorated Teichmüller space of punctured surfaces, Comm. Math. Phys. 113 (1987) 299–339.
[31] R.C. Penner, Universal constructions in Teichmüller theory, Adv. Math. 98 (1993) 143–215.
[32] R.C. Penner, Decorated Teichmüller theory of bordered surfaces, Comm. Anal. Geom. 12 (2004) 793–820.
[33] J. Roberts, Kirby calculus in manifolds with boundary, Turkish J. Math. 21 (1) (1997) 111–117.