Sel. math., New ser. 8 (2002) 341 – 371
1022–1824/02/030341–31

Selecta Mathematica, New Series

The Århus integral of rational homology 3-spheres II: Invariance and universality

Dror Bar-Natan, Stavros Garoufalidis, Lev Rozansky and Dylan P. Thurston

Abstract. We continue the work started in [Å-I], and prove the invariance and universality in the class of finite type invariants of the object defined and motivated there, namely the Århus integral of rational homology 3-spheres. Our main tool in proving invariance is a translation scheme that translates statements in multi-variable calculus (Gaussian integration, integration by parts, etc.) to statements about diagrams. Using this scheme the straightforward "philosophical" calculus-level proofs of [Å-I] become straightforward honest diagram-level proofs here. The universality proof is standard and utilizes a simple "locality" property of the Kontsevich integral.

Mathematics Subject Classification (2000). 57M27.

Key words. Finite type invariants, Gaussian integration, 3-manifolds, Kirby moves, holonomy.

1. Introduction

This paper is the second in a four-part series on "the Århus integral of rational homology 3-spheres". In the first part of this series, [Å-I], we gave the definition of a diagram-valued invariant $\AA$ of "regular pure tangles", pure tangles whose linking matrix is non-singular,[1] and gave "philosophical" reasons why $\AA$ should descend to an invariant of regular links (framed links with non-singular linking matrix), and as such satisfy the Kirby relations and hence descend further to an invariant of rational homology 3-spheres. Very briefly, we defined the pre-normalized Århus integral $\AA_0$ to be the composition

[1] A precise definition of regular pure tangles appears in [Å-I, Definition 2.2]. It is a good idea to have [Å-I] handy while reading this paper, as many of the definitions introduced and explained there will only be repeated here in a very brief manner.

$$\mathring{A}_0 : \left\{ \begin{array}{c} \text{regular} \\ \text{pure} \\ \text{tangles} \end{array} \right\} =$$

$$= RPT \xrightarrow[\substack{\text{the [LMMO]} \\ \text{version of the} \\ \text{Kontsevich integral}}]{\check{Z}} \mathcal{A}(\uparrow_X) \xrightarrow[\substack{\text{formal} \\ \text{PBW}}]{\sigma} \mathcal{B}(X) \xrightarrow[\substack{\text{formal} \\ \text{Gaussian} \\ \text{integration}}]{\int^{FG}} \mathcal{A}(\emptyset).$$

In this formula,

- RPT denotes the set of regular pure tangles whose components are marked by the elements of some finite set X (see [Å-I, Definition 2.2]).
- $\check{Z}$ denotes the Kontsevich integral normalized as in [LMMO] (check [Å-I, Definition 2.6] for the adaptation to pure tangles).
- $\mathcal{A}(\uparrow_X)$ denotes the completed graded space of chord diagrams for X-marked pure tangles modulo the usual $4T/STU$ relations (see [Å-I, Definition 2.4]).
- σ denotes the diagrammatic version of the Poincaré-Birkhoff-Witt theorem (defined as in [B-N1], [B-N2], but normalized slightly differently, as in [Å-I, Definition 2.7]).
- $\mathcal{B}(X)$ denotes the completed graded space of X-marked uni-trivalent diagrams as in [B-N1], [B-N2] and [Å-I, Definition 2.5].
- $\mathcal{A}(\emptyset)$ denotes the completed graded space of manifold diagrams as in [Å-I, Definition 2.3].
- $\int^{FG}$ is a new ingredient, first introduced in [Å-I, Definition 2.9], called "formal Gaussian integration". In a sense explained there and developed further here, it is a diagrammatic analogue of the usual notion of Gaussian integration.

Our main challenge in this paper is to prove that $\mathring{A}_0$ descends to an invariant of links which is invariant under the second Kirby move. As it turns out, this depends heavily on understanding properties of formal Gaussian integration, which are all analogues of properties of standard integration over Euclidean spaces. We develop the necessary machinery in Section 2 of this article, and then in Section 3 we move on and use this machinery to prove two of our main results, Proposition 1.1 and Theorem 1:

Proposition 1.1. *The regular pure tangle invariant $\mathring{A}_0$ descends to an invariant of regular links and as such it is insensitive to orientation flips (of link components) and invariant under the second Kirby move.*

Definition 1.2. Let $U_\pm$ be the unknot with framing ± 1, and let σ_+ (σ_-) be the number of positive (negative) eigenvalues of the linking matrix of a regular link L. Let the Århus integral $\mathring{A}(L)$ of L be

$$\mathring{A}(L) = \mathring{A}_0(U_+)^{-\sigma_+}\mathring{A}_0(U_-)^{-\sigma_-}\mathring{A}_0(L), \tag{1}$$

with all products and powers taken using the disjoint union product of $\mathcal{A}(\emptyset)$.

Theorem 1. *$\mathring{A}$ is invariant under orientation flips and under both Kirby moves, and hence (*[Ki]*) it is an invariant of rational homology 3-spheres.*

Our second goal in this article is to prove that $\mathring{A}$ is a universal Ohtsuki invariant, and hence that all $\mathbb{Q}$-valued finite-type invariants of integer homology spheres are compositions of $\mathring{A}$ with linear functionals on $\mathcal{A}(\emptyset)$. We present all relevant definitions and proofs in Section 4 below.

2. Formal diagrammatic calculus

In this section we study the theory of formal Gaussian integration, along with a neighboring theory of formal differential operators. The idea is that monomials can be represented by vertices of certain valences, and differentiation (almost always) and integration (at least in the case of Gaussian integration) are given by combinatorial formulas that can be viewed as manipulations done on certain kinds of diagrams built out of these vertices. This extracts some parts of good old elementary calculus, and replaces algebraic manipulations by a diagrammatic calculus. Now forget the interpretation of diagrams as functions and operators, and you will be left with a formal theory of diagrams in which there are formal diagrammatic analogs of various calculus operations and of certain theorems from classical calculus.

This diagrammatic theory is more general than what we need for this paper; it is not restricted to the diagrams (and relations) that make up the spaces that we use often, such as $\mathcal{A}$ and $\mathcal{B}$. We are sure such a general formal diagrammatic theory was described many times before and we make no claims of originality. This theory is implicit in many discussions of Feynman diagrams in physics texts, but we are not aware of a good reference that does everything that we need the way we need it. Hence in this section we describe in some detail that part of the general theory that we will use in the later sections.

2.1. The general setup

Our basic objects are diagrams with some internal structure (that we mostly do not care about), and some number of "legs", outward pointing edges that end in a univalent vertex. The legs are labeled by a vector space, or by a variable that lives in the dual of that vector space. We think (for the purpose of the analogy with standard calculus) of such a diagram as representing a tensor in the tensor product of the spaces labeled next to its legs. We assume that legs that are labeled the same way are interchangeable, meaning that our diagrams represent symmetric tensors whenever labels are repeating. Symmetric tensors can be identified with polynomials on the dual:

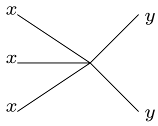

$\mapsto$ a polynomial of degree 3 in x and degree 2 in y.

We also allow legs labeled by dual spaces or by dual variables. By convention, the variable dual to x is denoted ∂_x. Just as the symmetric algebra $S(V^\star)$ can be regarded as a space of constant coefficient differential operators acting on $S(V)$, diagrams labeled by dual variables represent differential operators:

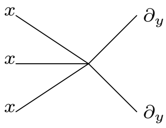

$\mapsto$ A second order differential operator acting on functions of y, with a coefficient cubic in x.

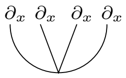

$\mapsto$ A fourth order constant coefficient differential operator.

We assume in addition that each diagram has an "internal degree", some non-negative half integer $(0, \frac{1}{2}, 1, \frac{3}{2}, \dots)$, associated with it. It is to be thought of as the degree in some additional (small or formal) parameter $\hbar$ that the whole theory depends upon. That is, the polynomials and differential operators that we imitate also depend on some additional parameter $\hbar$.

We also consider weighted sums of diagrams (representing not necessarily homogeneous polynomials and differential operators), and even infinite weighted sums of diagrams provided either their internal degree grows to infinity or their number of legs grows to infinity. These infinite sums represent power series (in $\hbar$ and/or in the variables labeled on the legs) and/or infinite order differential operators.

Finally, we allow some "internal relations" between the diagrams involved. That is, sometimes we mod out the spaces of diagrams involved by relations, such as the IHX and AS relation, that do not touch the external legs and the internal degree of a diagram. All operations that we will discuss below only involve the external legs and/or the internal degree, and so they will be well-defined even after moding out by such internal relations.

We then consider some operations on such diagrams. The operation of adding diagrams (whose output is simply the formal sum of the summands) corresponds to additions of polynomials or operators. The operation of disjoint union of diagrams (adding their internal degrees), extended bilinearly to sums of diagrams, corresponds to multiplying polynomials and/or composing differential operators (at least in the constant coefficients case, where one need not worry about the order of composition). Once summation and multiplication are available, one can define exponentiation and other analytic functions using power series expansions.

The most interesting operation we consider is the operation of contraction (or "gluing"). Two tensors, one in, say, $V^\star \otimes W$ and the other in, say, $V \otimes Z$, can be contracted, and the result is a new tensor in $W \otimes Z$. The graphical analog of this operation is the fusion of two diagrams along a pair (or pairs) of legs labeled by dual spaces or variables (while adding their internal degrees). In the case of legs labeled by dual variables, the calculus meaning of the fusion operation is the pairing of a derivative with a linear function. The laws of calculus dictate that when a differential operator D acts on a monomial f, the result is the sum of all possible ways of pairing the derivatives in D with the factors of f. Hence if D is a diagram representing a differential operator (i.e., it has legs labeled ∂_x, ∂_y, etc.) and f is a diagram representing a function (legs labeled $x, y, \dots$), we define

$$D \flat f = \begin{pmatrix} \text{sum of all ways of gluing all legs labeled } \partial_x \text{ on } D \\ \text{with some or all legs labeled } x \text{ on } f \text{ (and same} \\ \text{for } \partial_y \text{ and } y\text{, etc.)} \end{pmatrix}.$$

(This sum may be 0 if there are, say, more legs labeled ∂_x on D than legs labeled x on f). For example,

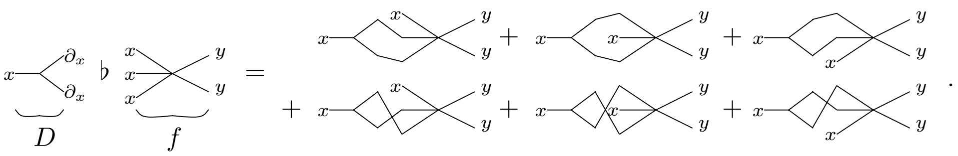

(In this figure 4-valent vertices are not real, but just artifacts of the planar projection). If this were calculus and the spaces involved were one-dimensional, we would call the above formula a proof that $x(\partial_x)^2\, x^3y^2 = 6x^2y^2$.

Remark 2.1. The reader may show that Leibniz's formula, $D \flat (fg) = (D \flat f)g + f(D \flat g)$ holds in our context, whenever D is a first order differential operator. (Remember that multiplication is disjoint union. D can connect to the disjoint union of f and g either by connecting to a leg of f, or by connecting to a leg of g.)

Remark 2.2. The reader may prove the exponential Leibniz's formula, $(\exp D) \flat \prod f_i = \prod (\exp D) \flat f_i$, where D is first order in x and has no coefficients proportional to x (i.e., where D has one leg labeled ∂_x and no legs labeled x).

Below we will need at some technical points an extension of this remark to the case when the operators involved are not necessarily first order. The result we need is a bit difficult to formulate, and doing so precisely would take us too far aside. But nevertheless, the result is rather easy to understand in "chemical" terms, in which diagrams are replaced by molecules and exponentiations are replaced by substance-filled containers. Notice that the exponentiation of some object $\mathcal{O}$ is the sum $\sum_k \mathcal{O}^k/k!$ of all ways of taking "many" unordered copies of $\mathcal{O}$, so it can be thought of "taking a big container filled with (copies of) the molecule $\mathcal{O}$".

A "homogeneous reaction" (in chemistry) is a reaction in which a homogeneous mixture A of mutually inert reactants is mixed with another homogeneous mixture B of mutually inert reactants, allowing reactions to occur and products to be produced. The result of such a reaction is homogeneous mixture of substances, each of which produced by some allowed reaction between one (or many) of the reactants in A and one (or many) of the reactants in B.

In our context, the "mixture" A is the exponential $\exp\sum\alpha_i f_i$ of some linear combination of (diagrams representing) functions. The mixture B is the exponential $\exp\sum\beta_j D_j$ of some sum of (diagrams representing) mutually inert differential operators D_j. That is, all of the differentiations in the D_j's must act trivially on all of the coefficients of the D_j's. That is, the diagrams occurring in the D_j's have "coefficient legs" labeled by some set of variables X and "differentiations legs" labeled by the dual variables to some disjoint set of variables Y. Computing $B \flat A$ is in some sense analogous to mixing A and B and allowing them to react. The result is some "mixture" (exponential of a sum) of compounds produced by reactions in which the legs in some number of the diagrams in B are glued to some of the legs in some number of the diagrams in A. These compounds come with weights ("densities") that are (up to minor combinatorial factors) the products of the densities α_i and β_j of their ingredients.

Remark 2.3. It is possible to write $(\exp\sum\beta_j D_j) \flat (\exp\sum\alpha_i f_i)$ as an explicit exponential using the above terms.

We sometimes consider relabeling operations, where one takes (say) all legs labeled x in a given diagram f and replaces the x labels by, say, y's, calling the result $D/(x \to y)$. This corresponds to a simple change of variable in standard calculus. We wish to allow more complicated linear reparametrizations as well, but for that we need to add a bit to the rules of the game. The added rule is that we also allow labels that are linear combinations of the basic labels (such as $x+y$), with the additional provision that the resulting diagrams are multi-linear in the labels (so a diagram with a leg labeled $x+y$ and another leg labeled $z+w$ is set equal to a sum of four diagrams labeled (x,z), (x,w), (y,z), and (y,w)). Now reparametrizations such as $x \to \alpha+\beta$, $y \to \alpha-\beta$ make sense.

Remark 2.4. The reader may show that the operation of reparametrization is compatible with the application of a differential operator to a function, as in standard calculus. For instance, in standard calculus the change of variables $x \to \alpha+\beta$, $y \to \alpha-\beta$ implies an inverse change for partial derivatives: $\partial_x \to (\partial_\alpha+\partial_\beta)/2$, $\partial_y \to (\partial_\alpha-\partial_\beta)/2$. Show that the same holds in the diagrammatic context:

$$(D \flat f) \Big/ \begin{pmatrix} x \to \alpha + \beta \\ y \to \alpha - \beta \end{pmatrix} =$$
$$= \left(D \Big/ \begin{pmatrix} \partial_x \to (\partial_\alpha + \partial_\beta)/2 \\ \partial_y \to (\partial_\alpha - \partial_\beta)/2 \end{pmatrix} \right) \flat \left(f \Big/ \begin{pmatrix} x \to \alpha + \beta \\ y \to \alpha - \beta \end{pmatrix} \right).$$

2.2. Formal Gaussian integration

In standard calculus, Gaussians are the exponentials of non-degenerate quadratics, and Gaussian integrals are the integrals of such exponentials multiplied by polynomials or appropriately convergent power series. Such integrals can be evaluated using the technique of Feynman diagrams (see [Å-I, appendix]). The diagrammatic analogs of these definitions and procedures are described below.

To proceed we must have available the diagrammatic building blocks for quadratics. These are what we call struts. They are lines labeled on both ends, they come in an assortment of forms (Figure 1) and they satisfy simple composition laws (Figure 2), that imply that the strut x —— ∂_x acts like the identity on legs labeled by x, and that the struts $x \frown x$ and $\partial_x \smile \partial_x$ are inverses of each other and their composition is x —— ∂_x. Struts always have internal degree 0. To make convergence issues simpler below, we assume that there are no diagrams of internal degree 0 other than the struts, and that for any $j \geq 0$ there is only a finite number of strutless diagrams (diagrams none of whose connected components are struts) with internal degree $\leq j$.

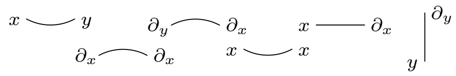

Figure 1. An assortment of struts.

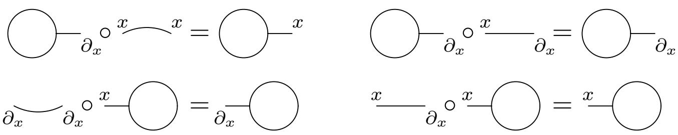

Figure 2. The laws governing strut compositions. The informal notation ∘ means: glue the two adjoining legs.

Definition 2.5. Let X be a finite set of variables. A quadratic Q in the variables in X is a sum of diagrams made of struts whose ends are labeled by these variables:

$$Q = \sum_{x,y \in X} l_{xy} \, {}^{x}\frown^{y},$$

where the matrix (l_{xy}) is symmetric. Such a quadratic is "non-degenerate" if (l_{xy}) is invertible. In that case, the inverse quadratic is the sum

$$Q^{-1} = \sum_{x,y \in X} l^{xy} \, {}_{\partial_x}\smile_{\partial_y},$$

where (l^{xy}) denotes the inverse matrix of (l_{xy}).

Definition 2.6. An infinite combination of diagrams of the form

$$G = P \cdot \exp Q/2$$

is said to be Gaussian with respect to the variables in X if Q is a quadratic (in those variables) and P is X-substantial, meaning that the diagrams in P have no components which are struts both of whose ends are labeled by members of X. Notice that P and Q are determined by G. In particular, the matrix $\Lambda = (l_{xy})$ of the coefficients of Q is determined by G. We call it the "covariance matrix" of G.

Definition 2.7. We say that a Gaussian $G = P \exp Q/2$ is non-degenerate, or integrable, if Q is non-degenerate. In such a case, we define the formal Gaussian integral of G with respect to X to be

$$\begin{aligned} \int^{FG} P \cdot \exp Q/2 \, dX &= \left\langle \exp -Q^{-1}/2, \, P \right\rangle_X \\ &= \left(\left(\exp -Q^{-1}/2 \right) \flat P \right) \Big/ \left(\begin{pmatrix} x \to 0 \\ \forall x \in X \end{pmatrix} \right), \end{aligned} \tag{2}$$

where

$$\langle D, P \rangle_X := \begin{pmatrix} \text{sum of all ways of glu-} \\ \text{ing the } \partial_x\text{-marked legs of} \\ D \text{ to the } x\text{-marked legs} \\ \text{of } P, \text{ for all } x \in X \end{pmatrix}. \qquad \begin{pmatrix} \text{This sum can be non-zero} \\ \text{only if the number of } \partial_x\text{-} \\ \text{marked legs of } D \text{ is equal to} \\ \text{the number of } x\text{-marked legs} \\ \text{of } P \text{ for all } x \in X. \end{pmatrix}$$

(Compare with [Å-I, Definition 2.9 and equation (6)]). The fact that P is X-substantial guarantees that for any given internal degree and any given number of legs, the computation of the Gaussian integral is finite.

Below we need to know some things about the relation between differentiation and integration. If a certain infinite order differential operator D contains too many struts, then the computation of (even a finite part of) $D \flat G$ may be infinite, or else, the result may be non-Gaussian and thus outside of our theory of integration. Both problems do not occur if D is "X-substantial", defined below:

Definition 2.8. Let X be a set of variables and D an differential operator. We say that D is X-substantial if it contains no struts both of whose ends are labeled by members of X or their duals.

2.3. Invariance under parity transformations

It is useful to know that formal Gaussian integrals, just like their real counterparts, are invariant under negation of one of the variables:

Proposition 2.9. *Let $G = P \exp Q/2$ be integrable with respect to X, let $y \in X$, and let $G' = P' \exp Q'/2$ be $G/(y \to -y)$. Then $\int^{FG} G\,dX = \int^{FG} G'\,dX$.*

Proof. One easily verifies that $P' = P/(y \to -y)$, $Q' = Q/(y \to -y)$ and $Q'^{-1} = Q'^{-1}/(\partial_y \to -\partial_y)$. Thus in each $(y \leftrightarrow \partial_y)$-gluing in the computation of $\int^{FG} G'\,dX$ two signs get flipped (relative to the computation of $\int^{FG} G\,dX$). Overall we flip an even number of signs, meaning, no signs at all. □

Remark 2.10. More generally, the reader may verify that Gaussian integration $\int^{FG}$ is compatible with linear reparametrizations such as in Remark 2.4. (Note that if the reparametrization matrix is M, then dual variables are acted on by M^{-1}. Each gluing in (2) is between one variable and one dual variable, and so occurrences of M and of M^{-1} come in pairs.)

2.4. Iterated integration

The classical Fubini theorem says that whenever all integrals involved are well defined, integration over a product space is equivalent to integration over one factor followed by integration over the other. We seek a similar iterated integration identity for formal Gaussian integrals.

Let $G = P \exp Q/2$ be a non-degenerate Gaussian with respect to a set of variables Z, and let $Z = X \cup Y$ be a decomposition of the set of variables into two disjoint subsets. Write the covariance matrix Λ of G and its inverse Λ^{-1} as block matrices with respect to this decomposition, taking the variables in X first and the variables in Y later:

$$\Lambda = \begin{pmatrix} A & B \\ B^T & C \end{pmatrix}, \qquad \Lambda^{-1} = \begin{pmatrix} D & E \\ E^T & F \end{pmatrix}.$$

The blocks A, C, D, and F are symmetric, and the fact that Λ and Λ^{-1} are inverses implies the following identities:

$$\begin{aligned} AD + BE^T &= I_X, & AE + BF &= 0, \\ B^T D + CE^T &= 0, & B^T E + CF &= I_Y. \end{aligned} \tag{3}$$

Proposition 2.11. *If the block A is invertible, then the formula*

$$\int^{FG} G\,dZ = \int^{FG} \left(\int^{FG} G\,dX \right) dY \tag{4}$$

makes sense and holds.

Proof. The left hand side of this formula is not problematic, and directly from the definition of formal Gaussian integration and from the decomposition of Λ^{-1} into blocks, we find that it equals

$$\left\langle \exp -\frac{1}{2}\left(\sum_{x,x'\in X} D^{xx'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} +2 \sum_{x\in X,\, y\in Y} E^{xy}\ {}_{\partial_x}\!\smile\!{}_{\partial_y} + \sum_{y,y'\in Y} F^{yy'}{}_{\partial_y}\!\smile\!{}_{\partial_{y'}} \right),\ P \right\rangle .$$

We usually suppress the summation symbols, getting

$$\left\langle \exp -\frac{1}{2}\left(D^{xx'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} +2E^{xy}\ {}_{\partial_x}\!\smile\!{}_{\partial_y} +F^{yy'}\ {}_{\partial_y}\!\smile\!{}_{\partial_{y'}} \right),\ P \right\rangle . \tag{5}$$

We need to prove that the right hand side of (4) is well defined and equals (5). Let us start with the inner integral. Rewriting the integrand in the form

$$\left(P \exp \frac{1}{2}\left(2B_{xy}\ {}^{x}\!\frown\!{}^{y} +C_{yy'}\ {}^{y}\!\frown\!{}^{y'} \right)\right) \exp \frac{1}{2} A_{xx'}\ {}^{x}\!\frown\!{}^{x'},$$

we find that the integrand is Gaussian with respect to X with covariance matrix A. This matrix was assumed to be invertible, and hence the inner integral G' is defined and equals (denoting the inverse of A by $\bar{A}$, and using (2))

$$G' = \left(\exp -\frac{1}{2}\left(\bar{A}^{xx'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} \right) \flat \left(P \exp \frac{1}{2}\left(2B_{xy}\ {}^{x}\!\frown\!{}^{y} +C_{yy'}\ {}^{y}\!\frown\!{}^{y'} \right)\right)\right) \Big/ (x\to 0).$$

Using Remark 2.3 and suppressing the automatic evaluation at $x=0$, this becomes

$$\left(\exp -\frac{1}{2}\left(\bar{A}^{xx'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} +2\bar{A}^{xx_1} B_{x_1 y}\ {}^{y}\!\!-\!\!-\!\!-\ {}_{\partial_x} \right) \flat P \right) \cdot$$
$$\cdot \exp \frac{1}{2}\left(C_{yy'}\ {}^{y}\!\frown\!{}^{y'} -B^{T}_{y'x'} \bar{A}^{x'x} B_{xy}\ {}^{y}\!\frown\!{}^{y'} \right).$$

We are now ready to evaluate the dY integral of G'. In the above formula G' is already written in the required format $P' \exp \frac{1}{2} Q'$, with $Q = C_{yy'}\ {}^{y}\!\frown\!{}^{y'} -B^{T}_{y'x'} \bar{A}^{x'x} B_{xy}\ {}^{y}\!\frown\!{}^{y'}$. Thus the covariance matrix is $\Lambda' = C - B^T \bar{A} B$. The relations (3) imply that Λ' is invertible, with inverse F. Thus the integral with respect to Y of G' is (suppressing the evaluation at $y=0$)

$$\left(\exp -\frac{1}{2} F^{yy'}\ {}_{\partial_y}\!\smile\!{}_{\partial_{y'}} \right) \flat \left(\exp -\frac{1}{2}\left(\bar{A}^{xx'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} +2\bar{A}^{xx_1} B_{x_1 y}\ {}^{y}\!\!-\!\!-\!\!-\ {}_{\partial_x} \right) \flat P \right).$$

Again using Remark 2.3, this is

$$= \exp -\frac{1}{2}\left(\begin{array}{c} \bar{A}^{xx'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} + \bar{A}^{xx_1}B_{x_1y}F^{yy'}B^T_{y'x'_1}\bar{A}^{x'_1x'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} \\ -\ 2\bar{A}^{xx'}B_{x'y'}F^{y'y}\ {}_{\partial_y}\!\smile\!{}_{\partial_x} + F^{yy'}\ {}_{\partial_y}\!\smile\!{}_{\partial_{y'}} \end{array}\right) \flat\, P.$$

Giving names to the coefficients and switching to matrix-talk, we find that this is

$$\exp -\frac{1}{2}\left(L^{xx'}\ {}_{\partial_x}\!\smile\!{}_{\partial_{x'}} + 2M^{xy}\ {}_{\partial_x}\!\smile\!{}_{\partial_y} + F^{yy'}\ {}_{\partial_y}\!\smile\!{}_{\partial_{y'}}\right) \flat\, P,$$

with $L = \bar{A} + \bar{A}BFB^T\bar{A}^T$ and $M = -\bar{A}BF$. A second look at the relations (3) reveals that $L = D$ and $M = E$, proving that the last formula is equal to (5), as required. □

2.5. Integration by parts

Let D be a diagram representing a differential operator with respect to the variable z (that is, it may have "differentiation legs" labeled ∂_z, "coefficient legs" labeled z, and possibly other legs labeled by other variables). Assume D has l differentiation legs labeled ∂_z ("D is of order l") and k coefficient legs labeled z.

Definition 2.12. The "divergence" $\operatorname{div}_z D$ of D with respect to z is the result of "applying D to its own coefficients". That is,

$$\operatorname{div}_z D = \begin{cases} 0 & \text{if } l > k, \\ \begin{pmatrix}\text{sum of all ways of attaching all legs labeled} \\ \partial_z \text{ with some or all legs labeled } z\end{pmatrix}, & \text{if } l \le k. \end{cases}$$

(Compare with the standard definition of the divergence of a vector field, where each derivative "turns back" and acts on its own coefficient).

In standard calculus, the following proposition is an easy consequence of integration by parts:

Proposition 2.13. *Let X be a set of variables, and let $z \in X$. If G is a non-degenerate Gaussian with respect to X and D is an X-substantial operator of order l, then*

$$\int^{FG} D \flat\, G\, dX = (-1)^l \int^{FG} (\operatorname{div}_z D) G\, dX. \tag{6}$$

Proof. Write $G = P \exp Q/2$ with $Q = \sum_{x,y\in X} l_{xy}\ {}^x\!\frown\!{}^y$. Let us pick one leg marked ∂_z in D, put a little asterisk $(*)$ on it, and follow it throughout the computation of the left hand side of (6). First, in computing $D \flat G$, the special leg gets glued either to one of the legs in P, or to one of the legs in $\exp Q/2$. The result looks something like

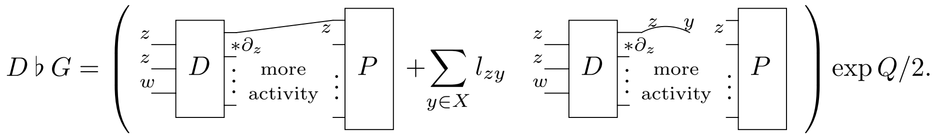

The next step is integration. The factor $\exp Q/2$ is removed, and struts labeled and weighted by the negated inverse covariance matrix are glued in. Of particular interest is the strut ${}_y\smile_{y'}$ glued to the marked leg in the right term. It comes with a coefficient like $-l^{yy'}$ from the negated inverse covariance matrix, which multiplies the coefficient l_{zy} already in place. The summation over y evaluates matrix multiplication of a matrix and its inverse, and we find that $y' = z$ and the overall coefficient is -1. The other end of this strut is glued to some leg (with a label in X), either on P or on D. Following that, all other negated inverse covariance gluings are performed. The result looks something like

$$\int^{FG} D \flat G dX = \left(\exp -\frac{1}{2} \sum_{x,y \in X} l^{xy} \, {}_{\partial_x}\smile_{\partial_y} \right) \flat$$

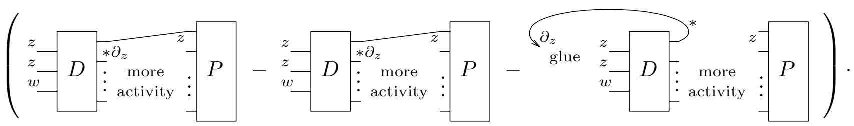

In this formula the first term cancels the second, and we are left only with the third. But the same argument can be made for all legs marked ∂_z in D, and hence in left hand side integral in equation (6) they all have to "turn back" and differentiate a coefficient of D. Counting signs, this is precisely the right hand side of equation (6). □

2.6. Our formal universe

Below we apply the formalism and techniques developed in this section in the case where the diagrams are X-marked uni-trivalent diagrams modulo the AS and IHX relations, for some label set X. The "internal degree" of a diagram is half the number of internal (trivalent) vertices it has, and the struts are simply the internal degree 0 diagrams — uni-trivalent diagrams that have no internal vertices. In the Kontsevich integral, the coefficients of struts measure linking numbers between the components marked at their ends (or self linkings, if the two ends are marked the same way).

The spaces of uni-trivalent diagrams that we consider are Hopf algebras, and the formal linear combinations of uni-trivalent diagrams that we take as inputs come from evaluating the Kontsevich integral, whose values are always grouplike. Thus our inputs are always exponentials. Splitting away the struts, we find that they are always Gaussian with the linking matrix of the underlying link as the covariance matrix. If we stick to pure tangles whose linking matrix is non-singular, our inputs are always integrable.

3. The invariance proof

Let us start with an easy warm-up:

Proposition 3.1. *$\mathring{A}_0$ is insensitive to orientation flips.*

Proof. Flipping the orientation of the component labeled x in some pure tangle L acts on $\sigma\check{Z}(L)$ by flipping the sign of all uni-trivalent diagrams that have an odd number of x-marked legs (check [B-N1, Section 7.2] for the case of knots; the case of pure tangles is the same). Namely, it acts by the substitution $x \to -x$. Now use Proposition 2.9. □

3.1. $\mathring{A}_0$ descends to regular links

Our first real task is to show that if two regular pure tangles have the same closures then they have the same pre-normalized Århus integral and hence the pre-normalized Århus integral $\mathring{A}_0$ descends to regular links. We first extend the definition of $\mathring{A}_0$ to some larger class of "closable" objects (Definition 3.2), the class of regular dotted Morse links. We then show that $\mathring{A}_0$ descends from that class to links (Proposition 3.4), and finally that regular pure tangles "embed" in regular dotted Morse links (Proposition 3.7). Taken together, these two propositions imply that $\mathring{A}_0$ descends to regular links also from regular pure tangles.

We should note that the Århus integral can be defined and all of its properties can be proved fully within the class of regular dotted Morse links, and that this is essentially what we do in this paper. The only reasons we also work with regular pure tangles are reasons of elegance.

Definition 3.2. A dotted Morse link L is a link embedded in $\mathbb{R}^3_{xyt}$ so that the third Euclidean coordinate t is a Morse function on it, together with a dot marked on each component. We assume that the components of L are labeled by the elements of some label set X. Notice that we do not divide by isotopies. The "closure" of a dotted Morse link is the (X-marked) link obtained by forgetting the dots and dividing by isotopies. These definitions have obvious framed counterparts.

Remark 3.3. Why so ugly a definition? Because all other choices are even worse. We have to "dot" the link components because we want the Kontsevich integral to be valued in $\mathcal{A}(\uparrow_X)$ (see below). But then we have to give up isotopy invariance at the time slices of the dots, and it is simpler to give it up altogether. See also the comment about q-tangles/non-associative tangles above Proposition 3.7.

The framed Kontsevich integral Z, as well as the variant $\check{Z}$ due to Le, H. Murakami, J. Murakami, and Ohtsuki [LMMO], both have obvious definitions in the case of framed dotted Morse links. The new bit is that each component has dot marked on it, which can serve as a cutting mark for scissors. In other words, every component can be regarded as a directed line, and thus the images of Z and $\check{Z}$ are in $\mathcal{A}(\uparrow_X)$. But now we can compose $\check{Z}$ with σ and then with $\int^{FG}$, and we find that the pre-normalized Århus integral $\AA_0$ can also be defined on regular dotted Morse links (framed dotted Morse links with a non-singular linking matrix).

Proposition 3.4. *$\AA_0$ descends from regular dotted Morse links to regular links.*

Proof. The usual invariance argument for the Kontsevich integral (see [Ko], [B-N1]) applies also in the case of (framed) dotted Morse links, provided the time slices of the dots are frozen. So the only thing we need to prove is that $\AA_0$ is invariant under sliding the dots along a component; once this is done, the frozen time slices melt and we have complete invariance.

A different way of saying that a dot moves on a framed dotted Morse link L is saying that we have two dots on one component (say z), cutting it into two subcomponents x and y. Each time we ignore one of the dots and compute $\check{Z}$, getting two results G_1 and G_2, and we wish to compare the integrals of G_1 and G_2. Alternatively, we can keep both dots on the z component and compute $\check{Z}$ in the usual way, only cutting the resulting chord diagrams open at both dots, getting a result G in the space[2] $\mathcal{A}(\uparrow_x\uparrow_y\uparrow_E)$. From G both G_1 and G_2 can be recovered by attaching the components x and y in either of the two possible orders. This process is made precise in Definition 3.5 below, and the fact that $\int^{FG}\sigma G_1 = \int^{FG}\sigma G_2$ follows from the "cyclic invariance lemma" (Lemma 3.6) below. We only need to comment that G is group-like like any evaluation of the Kontsevich integral. □

Definition 3.5. Let $\overrightarrow{m}^{xy}_z : \mathcal{A}(\uparrow_x\uparrow_y\uparrow_E) \to \mathcal{A}(\uparrow_z\uparrow_E)$ be the map described in Figure 3. The map $\overrightarrow{m}^{yx}_z$ is the same, only with the roles of x and y interchanged.

Lemma 3.6. (the cyclic invariance lemma)**.** *If $G \in \mathcal{A}(\uparrow_x\uparrow_y\uparrow_E)$ is group-like and $\sigma\overrightarrow{m}^{xy}_z G$ is an integrable member of $\mathcal{B}_n$, then $\sigma\overrightarrow{m}^{yx}_z G$ is also integrable and the two integrals are equal:*

$$\int^{FG} \sigma\overrightarrow{m}^{xy}_z G\, dzdE = \int^{FG} \sigma\overrightarrow{m}^{yx}_z G\, dzdE.$$

[2] The notation means: pure tangle diagrams whose skeleton components are labeled by the symbols x, y, and some additional $n-1$ symbols in some set E of "Extra variables". Below we will use variations of this notation with no further comment.

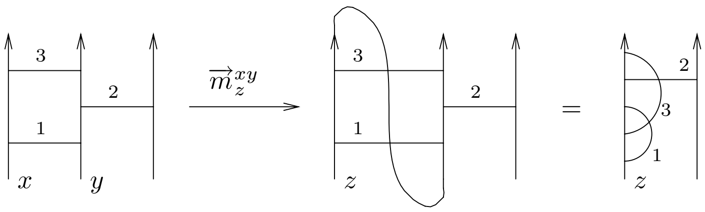

Figure 3. The map $\overrightarrow{m}^{xy}_z$ in the case $n=2$: Connect the strands labeled x and y in a diagram in $\mathcal{A}(\uparrow_x\uparrow_y\uparrow)$, to form a new "long" strand labeled z, without touching all extra strands.

Proof. It is easy to verify that $G_1 = \sigma\overrightarrow{m}^{xy}_z G$ and $G_2 = \sigma\overrightarrow{m}^{yx}_z G$ are both group-like and hence Gaussian, and that they have the same covariance matrix (when we apply this lemma as in Proposition 3.4, in both cases the covariance matrix is the linking matrix of the underlying link). Thus if one is integrable so is the other, and we have to prove the equality of the integrals.

The case of knots. If $n = 1$ then the fact that $\mathcal{A}(\uparrow)$ is isomorphic to $\mathcal{A}(\circlearrowright)$ (namely, the commutativity of $\mathcal{A}(\uparrow)$, see [B-N1]) implies that $\overrightarrow{m}^{xy}_z = \overrightarrow{m}^{yx}_z$ and there's nothing to prove.

The lucky case. If $G_{1,2}$ are integrable with respect to E we can use Proposition 2.11 and compute the integrals with respect to those variables first. The results are diagrams labeled by just one variable (z) (namely, functions of just one variable), and we are back in the previous case.

The ugly case. If $G_{1,2}$ are not integrable with respect to E, we can perturb them a bit by multiplying by some $\exp\sum_{i,j}\epsilon_{ij}\,{}^{e_i}\frown^{e_j}$ to get $G^\epsilon_{1,2}$. The integrals of $G^\epsilon_{1,2}$ (with respect to all variables) depend polynomially on the ϵ's in any given degree. For generic ϵ's we get $G^\epsilon_{1,2}$ that are integrable with respect to E, and we fall back to the lucky case. Thus the integrals of $G^\epsilon_{1,2}$ are equal as power series in the ϵ's, and in particular they are equal at $\epsilon_{ij} = 0$. □

Every (framed) pure tangle L defines a class of associated (framed) dotted Morse links, obtained by picking a specific Morse representative of L, marking dots at the tops of all strands, and closing to a link in some specific way making sure that the down-going strands used in the closure are very far (d miles away) from the original pure tangle. An example is in Figure 4. What's very far? In the infinite limit; meaning that whenever we refer to an associated (framed) dotted Morse link, we really mean "a sequence of such, with $d \to \infty$". To remind ourselves of that, we add the phrase "(at limit)" to the statements that are true only when this (or a similar, see below) limit is taken. If one is ready to sacrifice some simplicity, all of these statements can be formulated without limits if the technology of q-tangles ([LM1]) (or, what is nearly the same, non associative tangles ([B-N3])) is used instead of using specific Morse embeddings. Readers familiar with [LM1] and/or [B-N3] should have no difficulty translating our language to the more precise language of those papers.

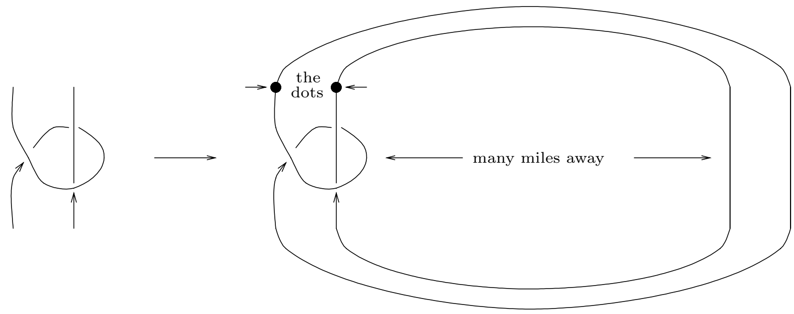

Figure 4. A pure tangle and an associated dotted Morse link.

Proposition 3.7 (at limit)**.** *If L is a pure tangle and $L^\bullet$ is an associated dotted Morse link, then $\check{Z}(L) = \check{Z}(L^\bullet)$.*

Proof (at limit). The dotted Morse link $L^\bullet$ is obtained from L by sticking L within a "closure element" C_X, shown on the right (for $|X| = 3$). Let C'_X be C_X with the two boxes at its ends removed. These two boxes denote "adapters" A and A^{-1} that only change the strand spacings to be uniform, from a possibly non-uniform spacings in C'_X.

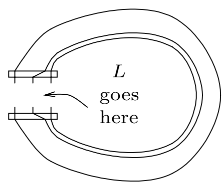

Inspecting the definitions of $\check{Z}$ for pure tangles (see [Å-I, Definition 2.6]) and for dotted Morse links (see [LMMO]), we see that we only need to show that $Z(C_X) = \Delta_X \nu$ in the space $\mathcal{A}(\uparrow_X)$ (check [Å-I, Definition 2.6] for the definition of Δ_X). Here C_X is itself regarded as a dotted Morse link (with the dots at the space allotted for L, which is assumed to be small relative to the size of C_X itself) and Z denotes the Kontsevich integral in its standard normalization. Clearly $Z(C_{\{x\}}) = \nu = \Delta_{\{x\}}\nu$, as $C_{\{x\}}$ is the dotted unknot and ν is by definition the Kontsevich integral of the unknot. Theorem 4.1 of [LM2], rephrased for dotted Morse links, says that doubling a component (so that the two daughter components are parallel and very close) and then computing Z is equal to Z followed by Δ. In other words, $Z(C_{\{x,y\}}) = \Delta_{\{x,y\}}(\nu)$. Iterating this argument, we find that $Z(C'_X) = \Delta_X\nu$, for some *specific* (at limit) choice of strand spacings in C'_X. But $\Delta_X\nu$ is central and hence $Z(C_X) = Z(A^{-1})Z(C'_X)Z(A) = Z(A)^{-1}(\Delta_X\nu)Z(A) = \Delta_X\nu$. □

3.2. $\AA_0$ is invariant under the second Kirby move

Definition 3.8. A tight Kirby move $L_1 \to L_2$ is a move between two framed dotted Morse links L_1 and L_2 as in Figure 5, in which

- Before the move the two parallel strands in the domain S are "tight". Namely, they are very close to each other relative to the distance between them and any other feature of the link.
- The doubling of the y component is done in a "very tight" fashion. Namely, the distance between the the copies of y produced is very small relative to the scale in which the rest of the link is drawn, even much smaller than the original distance between the x and y components.
- The dots on the x and y components are inside the domain S both before and after surgery, and they are placed as in the picture on the right.

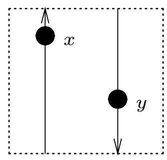

before surgery

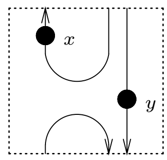

after surgery

We extend the notion of "at limit" to mean that "tightness" is also increased ad infinitum.

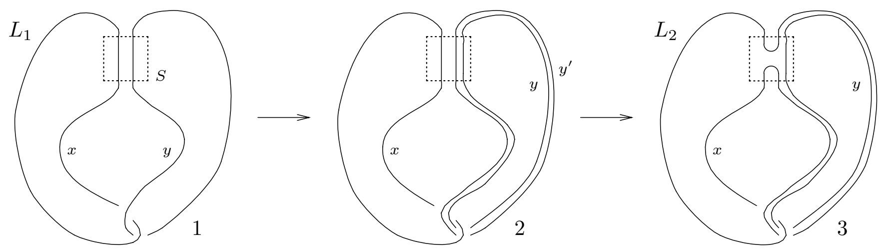

Figure 5. The second Kirby move $L_1 \to L_2$: (1) Some domain S in space in which some two components of the link (denoted x and y) are adjacent and nearly parallel is specified. (2) The component y is doubled (using its framing), getting a new component y'. (3) A surgery is performed in S combining y' into x, so that now the x component runs parallel to the y component in addition to running its own course. We say that the component x "slides" over the component y.

The following proposition is due to Le, H. Murakami, J. Murakami, and Ohtsuki [LMMO]. It holds for $\check{Z}$ and not for Z, and it is the reason why [LMMO] introduced $\check{Z}$. We present the "at limit" version, which is equivalent to the "q-tangle" version proved in [LMMO].

Proposition 3.9 (at limit, proof in [LMMO])**.** *Let $L_1 \to L_2$ be a tight Kirby move* *between two framed dotted Morse links L_1 and L_2 marked as in Figure 5. Then*

$$\check{Z}(L_2) = \overrightarrow{\Upsilon}\,\check{Z}(L_1),$$

where $\overrightarrow{\Upsilon} = \overrightarrow{m}_x^{xy'} \circ \Delta^y_{yy'}$ and $\Delta^y_{yy'}$ denotes the diagram-level operation of doubling the y strand (lifting all vertices on it in all possible ways, and calling the double y'). (Compare with [Å-I, equation (1)]*).*

Proof of Proposition 1.1*.* After propositions 3.1, 3.4 and 3.7 have been proved, all that remains is to show that $\mathring{A}_0$ is invariant under tight Kirby moves of framed dotted Morse links. (Notice that every Kirby move between links has a presentation as a tight Kirby move between dotted Morse links). Using Proposition 3.9 we find that it is enough to show that whenever G is a non-degenerate Gaussian (think $G = \sigma\check{Z}(L_1)$),

$$\int^{FG} G\,dE = \int^{FG} \overrightarrow{\Upsilon} G\,dE, \tag{7}$$

where we re-use the symbol $\overrightarrow{\Upsilon}$ to denote the same operation on the level of uni-trivalent diagrams.

Let Υ be the same as $\overrightarrow{\Upsilon}$, only with $m_x^{xy'}$ replacing $\overrightarrow{m}_x^{xy'}$, where $m_x^{xy'}G := G/(y' \to x)$. The operation Υ is a substitution operation of the form discussed in Section 2; $\Upsilon G = G/(y \to x+y)$. Remark 2.10 shows that equation (7) holds if $\overrightarrow{\Upsilon}$ is replaced by Υ. So we only need to analyze the difference $\overrightarrow{\Upsilon} - \Upsilon$. The difference $\overrightarrow{m}_x^{xy'} - m_x^{xy'}$ is given by gluing a certain sum D' of forests whose roots are labeled x and whose leaves are labeled ∂_x and $\partial_{y'}$, followed by the substitution $(y' \to x)$. Hence,

$$(\overrightarrow{\Upsilon} - \Upsilon)G = (D' \flat [G/(y \to y+y')])\,/(y' \to x) = (D \flat G)/(y \to x+y),$$

where D is D' with every $\partial_{y'}$ replaced by a ∂_y. (A precise formula for D can be derived from the results of Section 5.3, but we don't need it here). Clearly, $\operatorname{div}_y D = 0$; the coefficients of D are independent of y and every term in D is of positive degree in ∂_y. Now

$$\begin{aligned}\int^{FG} (\overrightarrow{\Upsilon} - \Upsilon)G\,dE &= \int^{FG} (D \flat G)/(y \to x+y)\,dE \\ &= \int^{FG} D \flat G\,dE && \text{by Remark 2.10} \\ &= 0 && \text{by Proposition 2.13.}\end{aligned}$$

□

3.3. $\AA$ and invariance under the first Kirby move

Proof of Theorem 1. Flipping the orientation of a component negates all linking numbers between it and any other component, and hence the linking matrix changes by a similarity transformation. The second Kirby moves adds all linking numbers involving the y-component (see Figure 5) to the corresponding ones with the x-component. This again is a similarity transformation. Similarity transformations do not change the numbers $\sigma_\pm$ of positive/negative eigenvalues. Thus $\AA$ is invariant under orientation flips and under the second Kirby move.

All that is left is to show that $\AA$ is invariant under the first Kirby move. Namely, that it is invariant under taking the disjoint union of a link with $U_\pm$, the unknot with framing ± 1.

Let L be an n-component regular link. Adding a far-away U_+ component to L multiplies $\sigma\check{Z}$ by $\sigma\check{Z}(U_+)$ (using the disjoint union product). The new linking matrix is block diagonal, with an additional $+1$ entry on the diagonal, and the same holds for the new inverse linking matrix. Thus the $(n+1)$-variable Gaussian integral of $\sigma\check{Z}(L \cup U_+)$ factors as the n-variable integral of $\sigma\check{Z}(L)$ times the 1-variable integral of $\sigma\check{Z}(U_+)$. We find that $\AA_0(L \cup U_+) = \AA_0(L) \cup \AA_0(U_+)$, and as σ_+ also increases by 1, $\AA(L \cup U_+) = \AA(L)$ as required. A similar argument works in the case of U_-. □

4. The Universality of the Århus Integral

4.1. What is universality?

Let us first recall the definition of universality, as presented in [Å-I, Section 2.2.2]].

Definition 4.1. An invariant U of integer homology spheres with values in $\mathcal{A}(\emptyset)$ is a "universal Ohtsuki invariant" if

(1) The degree m part $U^{(m)}$ of U is of Ohtsuki type $3m$ ([Oh]).

(2) If OGL denotes the Ohtsuki-Garoufalidis-Le map, defined in Figure 6, from manifold diagrams to formal linear combinations of unit framed algebraically split links in S^3, and S denotes the surgery map from such links to integer homology spheres, then

$$(U \circ S \circ OGL)(D) = D + \text{(higher degree diagrams)} \qquad (\text{in } \mathcal{A}(\emptyset))$$

whenever D is a manifold diagram (we implicitly linearly extend S and U, to make this a meaningful equation).

Theorem 2. *Restricted to integer homology spheres, $\AA$ is a universal Ohtsuki invariant.*

Some consequences of this theorem were mentioned in [Å-I, Corollary 2.13].

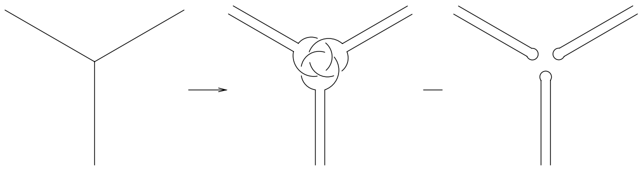

Figure 6. The OGL map: Take a manifold diagram D, embed it in S^3 in some fixed way of your preference, double every edge, replace every vertex by the difference of the two local pictures shown here, and put a $+1$ framing on each link component you get. The result is a certain alternating sum of 2^v links with e components each, where v and e are the numbers of vertices and edges of D, respectively.

4.2. $\mathring{A}$ is universal

The proofs of the two properties in the definition of universality are very similar and both depend on the same principle and the same observation. Both ideas have been used previously; see [Le], [B-N1].

The observation is that the degree m part of $\mathring{A}(L)$ comes from the internal degree m part of $\check{Z}^+(L)$, the strut-free part of $\sigma\check{Z}(L)$ in $\mathcal{B}_n$. Formal Gaussian integration acts by connecting all legs of a uni-trivalent diagram to each other using struts. All univalent vertices disappear in this process, while the trivalent ones are untouched. And so the degree m part of $\mathring{A}(L)$ is determined by the internal degree m part of $\check{Z}^+(L)$ (and the linking matrix).

The principle we use is a certain "locality" property of the Kontsevich integral. Recall how the Kontsevich integral of a link L is computed. One sprinkles the link in an arbitrary way with chords, and takes the resulting chord diagrams with weights that are determined by the positions of the end points of the chords sprinkled. This means that if a localized site on the link get modified, only the weights of chord diagrams that have ends in that site can change. Suppose one marks k localized sites, designates a modification to be made to the link on each one of them, and computes the alternating sum of $\check{Z}$ evaluated on the 2^k links obtained by performing any subset of these modifications. The result Z must have a chord-end in each of the k sites, and this bounds from below the complexity of any diagram appearing in Z and constrains the form of the diagrams of least complexity that appear in Z. If more is known about the nature of the modifications performed, more can be said about the parts of a diagram D in Z that originate from the sites of the modifications, and thus more can be said about D altogether.

A very simple application of this principle is the proof of the universality of the Kontsevich integral in, say, [B-N1]. Two more applications prove Theorem 2.

Proof of Theorem 2. *$\mathring{A}^{(m)}$ is of Ohtsuki type $3m$:* Take a unit framed $(k+3m+1)$-component algebraically split link L. (That is, the linking matrix of L is a $(k+3m+1)$-dimensional diagonal matrix with diagonal entries ± 1). We think of the first k components of L as representing some "background" integral homology sphere, and of the last $3m+1$ components as "active" components, over which the alternating summation in Ohtsuki's definition of finite-type [Oh] is performed. Let L^{alt} denote that alternating summation. Namely, it is the alternating sum of the 2^{3m+1} sublinks of L in which some of the active components are removed. We have to show that $\mathring{A}^{(m)}(L^{\mathrm{alt}})=0$. By the principle, every diagram in $\check{Z}(L^{\mathrm{alt}})$ must have a chord-end on every active component of L. The map σ never 'disconnects' a diagram from a component, and so every diagram D in $\check{Z}^{+}(L^{\mathrm{alt}})$ must have at least one leg per active component. But the linking matrix is a diagonal matrix, and hence the struts that are glued in the Gaussian integration are of form $\partial_x \smile \partial_x$ (both ends labeled the same way). So for the Gaussian integration to be non-trivial, there have to be at least two legs per active component of L, bringing the total to at least $2(3m+1)=6m+2$ legs. Each such leg must connect to some internal vertex, and there are at most three legs connected to any internal vertex. So there must by at least $2m+1$ internal vertices, and so the internal degree of D must be higher than m. By the observation, this means that $\mathring{A}(L^{\mathrm{alt}})$ vanishes in degrees up to and including m.

$\mathring{A}\circ OGL$ is the identity mod higher degrees: Let D be a manifold diagram. We aim to show that

$$(\mathring{A}\circ OGL)(D)=D+(\text{higher degree diagrams}) \qquad (\text{in } \mathcal{A}(\emptyset)). \tag{8}$$

If D is of degree m, it has $2m$ vertices and $L^{\mathrm{alt}}:=OGL(D)$ is an alternating summation over modifications in $2m$ sites. By the principle, there must be a contribution to $\check{Z}(L^{\mathrm{alt}})$ coming from each of those sites. Had there been just one such site, we would have been looking at the difference B between (a tangle presentation of) the Borromean rings and a 3-component untangle. As the Borromean linking numbers are equal to those of the untangle (both are 0), there are no struts in $\check{Z}(B)$, and the leading term is proportional to a Y diagram connecting the three components, looking like 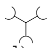. A simple computation shows that the constant of proportionality is 1 (cf. [Le]).

Thus, the leading term in $\check{Z}(L^{\mathrm{alt}})$ has a Y piece corresponding to every vertex of D, and the overall coefficient is 1. The map σ into uni-trivalent diagrams drops the loops corresponding to the link components and replaces them by labels on the univalent vertices thus created. (It also adds terms that come from gluing trees; these terms have a higher internal degree, so, by the observation, at lowest degree we can ignore them). Gaussian integration (with an identity covariance matrix, as we have here) simply connects legs with equal labels using struts, and the result is back again the diagram D we started with. This process is summarized in Figure 7. The renormalization in (1) doesn't touch any of that, and hence equation (8) holds. □

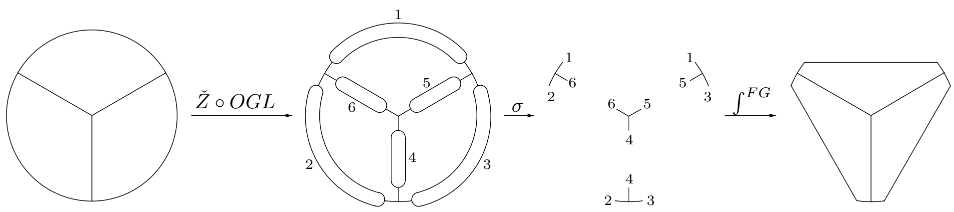

Figure 7. The computation of the leading order term in $\AA_0(OGL(D))$

5. Odds and ends

5.1. Homology spheres with embedded links

Everything said in the invariance section of this paper (Section 3) holds (or has an obvious counterpart) in the case of rational homology spheres with embedded links. Most changes required are completely superficial — wherever "components" are mentioned, of links, string links, dotted Morse links, skeletons of chord diagrams, etc., one has to label some of the components as "surgery components" (indexed by some set Y) and the rest as "embedded link" components (indexed by X). Surgeries are performed only on the components so labeled, σ is only applied on those components, and Gaussian integrations is carried out only with respect to the variables corresponding to the surgery components. Only the surgery components count for the purpose of determining $\sigma_\pm$ in (1). The embedded link components correspond to the embedded link in the post-surgery manifold. The only action taken on embedded link skeleton components (of chord diagrams in $\check{Z}(L)$) is to take their closures. The target space of the link-enhanced Århus integral is a mixture $\mathcal{A}'(\circlearrowleft_X)$ of the space $\mathcal{A}(\uplus \circlearrowleft_X)$ of chord diagrams (mod $4T/STU$) whose skeleton is a disjoint union of X-marked circles (see [Å-I, Figure 3]) and the space $\mathcal{A}(\emptyset)$ of manifold diagrams (modulo AS and IHX). The diagrams in $\mathcal{A}'(\circlearrowleft_X)$ are the disjoint unions of diagrams in $\mathcal{A}(\circlearrowleft_X)$ and diagrams in $\mathcal{A}(\emptyset)$, and the relations are all the relations mentioned above.

The only (slight) difficulty is that one should also prove invariance under the second Kirby move (Figure 5) in the case where an embedded link component x

slides over a surgery component y. A careful reading of the proof of Proposition 1.1 shows that it covers this case as well, as it uses only the integration with respect to y, the surgery component.

While the link-enhanced target space $\mathcal{A}'(\circlearrowright_X)$ suggests what universality should be like in the case of invariants of integer homology spheres with embedded links, the necessary preliminaries on finite-type invariants of such objects where never worked out in detail. So at this time we do not attempt to generalize the results of Section 4 to the case where embedded links are present.

5.2. The link relation

We (the authors) are not terribly happy about Section 3.1. Rather than showing that $\mathring{A}_0$ descends to regular links, we would have much preferred to be able to define it directly on regular links. The problem is that the Kontsevich invariant of X-component links is valued in the space $\mathcal{A}(\circlearrowright_X)$ of chord diagrams (mod $4T/STU$) whose skeleton is a disjoint union of circles marked by the elements of X (see [Å-I, Figure 3]). This space is not isomorphic to $\mathcal{B}(X)$, but rather to a quotient space $\mathcal{B}^{\mathrm{links}}(X)$ thereof, and we don't know how to define $\int^{FG}$ on $\mathcal{B}^{\mathrm{links}}(X)$. Let us write a few more words. First, a description of $\mathcal{B}^{\mathrm{links}}(X)$:

Definition 5.1. A "link relation symbol" is an X-marked uni-trivalent diagram R^* one of whose legs is singled out and carries an additional $*$ mark. If the other mark on the special leg of R^* is, say, x, we say that R^* is an "x-flavored link relation symbol". The "link relation" R corresponding to an x-flavored link relation symbol R^* is the sum of all ways of connecting the $*$-marked leg near the ends of all other x-marked legs. It is an element of $\mathcal{B}(X)$. An example appears in Figure 8. Finally, let $\mathcal{B}^{\mathrm{links}}(X)$ be the quotient of $\mathcal{B}(X)$ by all link relations.

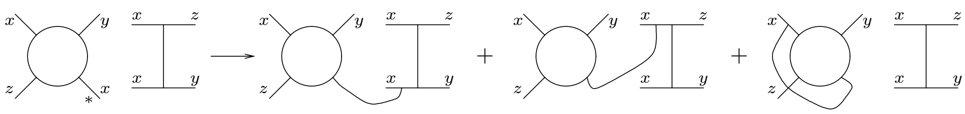

Figure 8. An x-flavored link relation symbol and the corresponding link relation.

Theorem 3. *The isomorphism* $\chi : \mathcal{B}(X) \to \mathcal{A}(\uparrow_X)$ *descends to a well defined isomorphism* $\chi : \mathcal{B}^{\mathrm{links}}(X) \to \mathcal{A}(\circlearrowright_X)$.

Proof (sketch). The fact that the link relations get mapped to 0 by χ composed with the projection on $\mathcal{A}(\circlearrowright_X)$ is easy — after applying χ, use the STU relation near every leg touched by the link relation. On a circular skeleton, the result is an ouroboros[3] summation, namely, it is 0. Suppose now you have a pair of diagrams in $\mathcal{A}(\uparrow_X)$ that get identified upon closing one of the skeleton components, say y. Use STU relations as here,

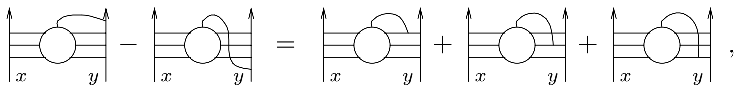

to turn their difference into a sum S of diagrams with a lower number of y-legs. Dropping the y component of the skeleton and forgetting the order of the y legs, the result is a y-flavored link relation. If trees are glued (as $\sigma = \chi^{-1}$ dictates) after the y component of the skeleton is dropped, then using the IHX relation one can show that the result is still a link relation:

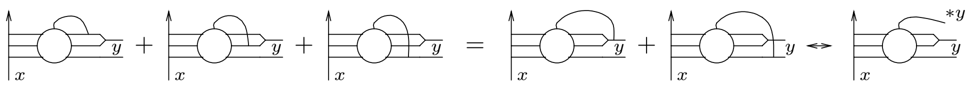

□

Problem 5.2. Is there a good definition for $\int^{FG}$ on (a domain in) $\mathcal{B}^{\text{links}}(X)$?

It makes no sense to ask if $\int^{FG}$ is well defined modulo the link relation; if G is a Gaussian and R is a link relation, $G+R$ would no longer be a Gaussian. We are mostly interested in integrating group-like G's. Maybe there's a more restrictive "group-like link relation" that relates any two group-like elements of $\mathcal{A}(\uparrow_X)$ that are equal modulo the usual link relation (namely, whose projections to $\mathcal{A}(\circlearrowleft_X)$ are the same)?

Problem 5.3. If $G_{1,2}$ are group-like elements of $\mathcal{A}(\uparrow_z\uparrow_E)$ that are equal modulo the link relation (applied only on the z component), is there always a group-like $G \in \mathcal{A}(\uparrow_x\uparrow_y\uparrow_E)$ so that $G_1 = \overrightarrow{m}^{xy}_z G$ and $G_2 = \overrightarrow{m}^{yx}_z G$? (notation as in Definition 3.5).

5.3. An explicit formula for the map $\overrightarrow{m}^{xy}_z$ on uni-trivalent diagrams

Let x and y be two elements in a free associative (but not-commutative) completed algebra. The Baker-Campbel-Hausdorf (BCH) formula (see e.g. [Ja]) measures the failure of the identity $e^{x+y} = e^x e^y$ to hold, in terms of Lie elements, or, what is the same, in terms of trees modulo the IHX and AS relation. The first few terms in the BCH formula are:

[3] The medieval symbol of holism depicting a snake that bites its own tail.

$$\log e^x e^y = \tag{9}$$

$$= x \quad +y \quad +\frac{1}{2}[x,y] \quad +\frac{1}{12}[x,[x,y]] \quad -\frac{1}{12}[y,[x,y]] \quad -\frac{1}{24}[x,[y,[x,y]]] \quad +\cdots$$

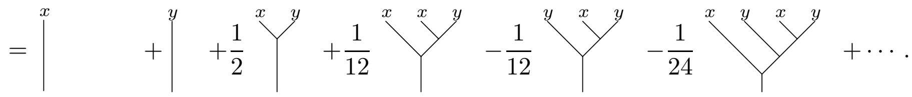

The proposition below states that as an operation on uni-trivalent diagrams, the map $\overrightarrow{m}^{xy}_z : \mathcal{A}(\uparrow_x\uparrow_y\uparrow_E) \to \mathcal{A}\ \uparrow_z\uparrow_E)$ of Definition 3.5 is given by gluing the disjoint-union exponential of the trees in the BCH formula (9). Precisely, let Λ be the sum of trees in the BCH formula, only with ∂_x replacing x, with ∂_y replacing y, and with a z marked on each root:

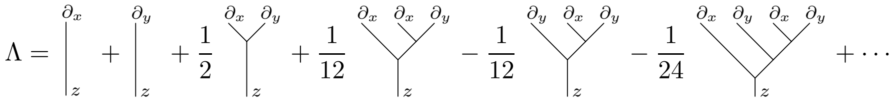

(10)

Proposition 5.4. *For any $C \in \mathcal{B}(\{x,y\}\dot\cup E)$,*

$$\sigma_n \overrightarrow{m}^{xy}_z \chi_{n+1} C = \langle \exp_{\dot\cup} \Lambda, C\rangle_{x,y}, \tag{11}$$

where χ_{n+1} denotes the standard isomorphism $\mathcal{B}(\{x,y\}\dot\cup E) \to \mathcal{A}(\uparrow_x\uparrow_y\uparrow_E)$ (whose inverse is σ_{n+1}) and σ_n denotes the standard isomorphism $\mathcal{A}(\uparrow_z\uparrow_E) \to \mathcal{B}(\{z\}\dot\cup E)$ (whose inverse is χ_n).

A noteworthy special case of this proposition is the case where $n = 1$ and C is a disjoint union $C_x \dot\cup C_y$ of a uni-trivalent diagram C_x whose legs are labeled only by x and a uni-trivalent diagram C_y whose legs are labeled only by y. In this case $\overrightarrow{m}^{xy}_z$ is (up to leg labelings) the product $\times_A : \mathcal{B}\otimes\mathcal{B} \to \mathcal{B}$ that $\mathcal{B}$ inherits from $\mathcal{A}$, and equation (11) becomes a specific formula for this product in terms of gluing forests. The existence of such a formula is immediate from the definition of $\sigma : \mathcal{A} \to \mathcal{B}$, and this existence was used in several places before (see e.g. [B-N2]), but we are not aware of a previous place where this formula was written explicitly. A similar formula is the "wheeling formula" of [BGRT].

Proof of Proposition 5.4. Let $\mathcal{A}^{xy}$ be the space of "planted forests" whose leaves are labeled ∂_x and ∂_y, modulo the usual STU (and hence AS and IHX) relations.

A planted forest is simply a forest in which the roots of the trees are "planted" along a directed line:

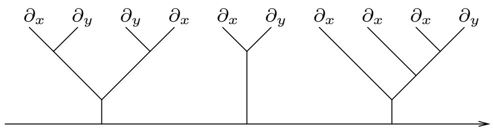

.

$\mathcal{A}^{xy}$ is similar in nature to $\mathcal{A}$; in particular, it is an algebra by the juxtaposition product $\times_A$ and it is graded, and hence an exponential $\exp_A$ and a logarithm $\log_A$ can be defined on it using power series.

Let ξ and η be the elements 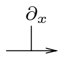 and 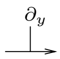 of $\mathcal{A}^{xy}$, respectively. Clearly,

$$\overrightarrow{m}^{xy}_z \chi_{n+1} C = \langle (\exp_A \xi) \times_A (\exp_A \eta), C \rangle_{x,y}$$
$$= \langle \exp_A \log_A \left((\exp_A \xi) \times_A (\exp_A \eta) \right), C \rangle_{x,y} .$$

But $\log_A ((\exp_A \xi) \times_A (\exp_A \eta))$ can be evaluated using the BCH formula (9). The result is $\chi^{xy}_z \Lambda$, where Λ was defined in equation (10) and $\chi^{xy}_z : \mathcal{B}^{xy}_z \to \mathcal{A}^{xy}$ is the natural isomorphism (whose inverse is σ^{xy}_z) of the space $\mathcal{B}^{xy}_z$ of forests with trees as in equation (10) (modulo AS and IHX) and the space $\mathcal{A}^{xy}$. Therefore $\overrightarrow{m}^{xy}_z \chi_{n+1} C = \langle \exp_A \chi^{xy}_z \Lambda, C \rangle_{x,y}$, and hence

$$\sigma_n \overrightarrow{m}^{xy}_z \chi_{n+1} C = \langle \sigma^{xy}_z \exp_A \chi^{xy}_z \Lambda, C \rangle_{x,y}. \tag{12}$$

The only thing left to note is that Λ is a sum of *trees*, namely forests in which z appears only once. On such forests $\exp_A \circ \chi^{xy}_z = \chi^{xy}_z \circ \exp_{\cup}$, and we see that equation (12) proves equation (11). ☐

Corollary 5.5. *(Compare with* [Å-I, Remark 1.7]*) Let* d_{BCH} *be* Λ *with the first two terms removed,*

$d_{\mathrm{BCH}} = \frac{1}{2}$ 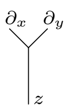 $+ \frac{1}{12}$ 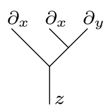 $- \frac{1}{12}$ 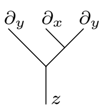 $- \frac{1}{24}$ 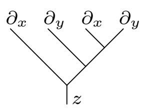 $+ \cdots ,$

and let $D_{\mathrm{BCH}} = \exp_{\cup} d_{\mathrm{BCH}}$. *Then*

$$\sigma_n \overrightarrow{m}^{xy}_z \chi_{n+1} C = (D_{\mathrm{BCH}} \flat\, C)/(x, y \to z).$$

Proof. Gluing the exponentials of the struts $|_z^{\partial_x}$ and $|_z^{\partial_y}$ is equivalent to applying the change of variables $(x, y \to z)$. □

5.4. A stronger form of the cyclic invariance lemma

As stated and proved in Section 3.1, the cyclic invariance lemma holds only when the integration is carried out over *all* the available variables. Otherwise, the argument given in the proof in the "lucky case" would not be a reduction to the case of knots. Here is an alternative statement and proof, that apply even if integration is carried out only over some subset F of the relevant variables.

Proposition 5.6 (Strong cyclic invariance). *If $G \in \mathcal{A}(\uparrow_x\uparrow_y\uparrow_E)$ is group-like and $\sigma\overrightarrow{m}_z^{xy}G$ is integrable with respect to some set F of variables satisfying $z \in F \subset \{z\}\cup E$, then $\sigma\overrightarrow{m}_z^{yx}G$ is also integrable and the two integrals are equal:*

$$\int^{FG} \sigma\overrightarrow{m}_z^{xy}G\,dF = \int^{FG} \sigma\overrightarrow{m}_z^{yx}G\,dF.$$

Proof. Let us describe the idea of the proof before getting into the details. Recall from [Å-I, Section 1.3] that we like to think of $\mathcal{A}(\uparrow_x\uparrow_y\uparrow_E)$ as a parallel of $\hat{U}(\mathfrak{g})^{\otimes n+1}$ and of $\mathcal{B}(\{x,y\}\cup E)$ as a parallel of $\hat{S}(\mathfrak{g})^{\otimes n+1}$ or of the function space $F(\mathfrak{g}^\star\oplus\cdots(n+1)\cdots\oplus\mathfrak{g}^\star)$. In this model, $\overrightarrow{m}_z^{xy}$ and $\overrightarrow{m}_z^{yx}$ become maps $\hat{U}(\mathfrak{g})^{\otimes n+1} \to \hat{U}(\mathfrak{g})^{\otimes n}$, defined by multiplying the first two tensor factors in the two possible orders, using the product $\times_U$ of $U(\mathfrak{g})$. If we had done the same with $\hat{S}(\mathfrak{g})$, using its product $\times_S$, the picture would have been a lot simpler. The product $\times_S$ is commutative, and m_z^{xy} and m_z^{yx} are the same. In fact, in the function space picture both maps become the "evaluation on the diagonal" map m:

$$F(\mathfrak{g}^\star\oplus\cdots(n+1)\cdots\oplus\mathfrak{g}^\star) \to F(\mathfrak{g}^\star\oplus\cdots(n)\cdots\oplus\mathfrak{g}^\star)$$
$$m : G(x,y,\dots) \mapsto G(z,z,\dots).$$

The difference between the two products $\times_U$ and $\times_S$ was discussed in [Å-I, Claim 1.5]. From that discussion it follows that

$$\overrightarrow{m}_z^{xy} - \overrightarrow{m}_z^{yx} = m \circ D,$$

where D is some differential operator acting on the variables x and y.

We wish to study the integral with respect to z of mDG; that is, the integral of DG on the diagonal $x = y$. We do this below by changing coordinates to the parallel coordinate $\alpha = (x+y)/2$ and the transverse coordinate $\beta = (x-y)/2$. In these coordinates the diagonal is given by $\beta = 0$, and z-integration can be replaced by α-integration at $\beta = 0$. We show that this α-integral vanishes by showing that the divergence (with respect to α) of D vanishes.

Let us turn to the details now. Let σG denote the (sum of) uni-trivalent diagrams corresponding to G via the isomorphism $\sigma : \mathcal{A}(\uparrow_x\uparrow_y\uparrow_E) \to \mathcal{B}(\{x,y\}\cup E)$. By abuse of notation, we re-use the symbols $\overrightarrow{m}^{xy}_z$ and $\overrightarrow{m}^{yx}_z$ to denote the maps $\mathcal{B}(\{x,y\}\cup E) \to \mathcal{B}(\{z\}\cup E)$ corresponding to the original $\overrightarrow{m}^{xy}_z$ and $\overrightarrow{m}^{yx}_z$ via the isomorphism σ (really, σ and its homonymic brother $\sigma : \mathcal{A}(\uparrow_z\uparrow_E) \to \mathcal{B}(\{z\}\cup E)$).

Corollary 5.5 implies that $\overrightarrow{m}^{xy}_z$ and $\overrightarrow{m}^{yx}_z$ both act on σG in the following manner:

- Glue some forest of non-trivial trees[4] whose leaves are labeled ∂_x and ∂_y and whose roots are labeled z to the x- and y-labeled legs of σG, gluing only ∂_x's to x's and ∂_y's to y's, and making sure that all leaves get glued.
- Relabel all remaining x- and y-labeled legs with z.

Gluing a forest of non-trivial trees does not touch the quadratic part of a group-like element, and hence this description implies that $\overrightarrow{m}^{xy}_z\sigma G$ and $\overrightarrow{m}^{yx}_z\sigma G$ have the same quadratic part. Hence if one is integrable so is the other, and they can be integrated together under the same formal integral sign, and we need to prove that

$$\int^{FG} (\overrightarrow{m}^{xy}_z - \overrightarrow{m}^{yx}_z)\sigma G\, dF = 0.$$

Using Proposition 5.4, we see that $(\overrightarrow{m}^{xy}_z - \overrightarrow{m}^{yx}_z)(\sigma G) = \langle D_1, \sigma G\rangle_{x,y}$, where D_1 is sum of forests of the general form

$D_1 = \cdots +$ 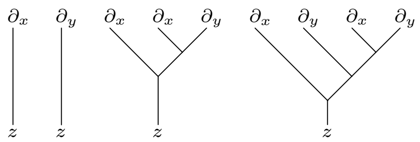

$+ \cdots .$

We now change variables to $\alpha = (x+y)/2$, $\beta = (x-y)/2$, $\partial_\alpha = \partial_x + \partial_y$, and $\partial_\beta = \partial_x - \partial_y$ (see Remarks 2.4 and 2.10), and rename the dummy integration variable z to be α. In the new variables, D_1 becomes a sum of forests D_2 of the general form

$D_2 = \cdots +$ 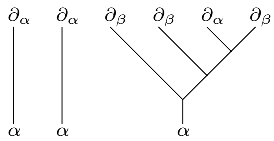

$+ \cdots ,$

[4] That is, a forest in which each tree has at least one internal vertex.

and the statement we need to prove is

$$\int^{FG} \langle D_2, \sigma G \rangle_{x,y} \, d\alpha d(F - \{z\}) = 0.$$

Notice that struts labeled ∂_x or ∂_y appear in D_1 in a symmetric role, and hence D_2 has only struts labeled by ∂_α and no struts labeled ∂_β. This implies that if D_3 is the strutless part of D_2, then $\langle D_2, \sigma G \rangle_{x,y} = D_3 \flat (\sigma G)$. Using integration by parts (see Section 2.5), we see that our job will be done one we can prove that $\operatorname{div}_\alpha D_3 = 0$.

Each component of each forest in D_3 must have at least one ∂_α leaf, for otherwise it would have only ∂_β leaves and it would vanish by the AS relation. Forests in which some tree has more than one ∂_α leaf contribute nothing to $\operatorname{div}_\alpha D_3$, as their ∂_α order is higher than their α degree. Thus we only care about forests of the form

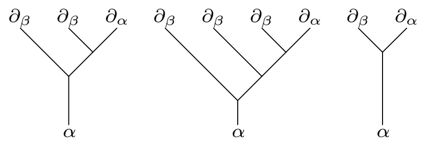

,

in which each component carries exactly one ∂_α.

The α-divergence of such a forest is obtained by having the α-derivatives act on the coefficients — namely, by summing over all possible ways of connecting the ∂_α's at the top of the picture to the α's at the bottom. The result is a sum W of diagrams like

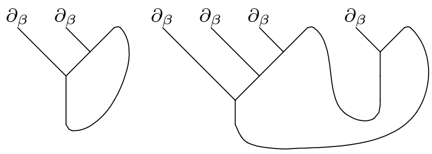

,

that are disjoint unions of ∂_β-wheels.

The next thing to notice is that $\overrightarrow{m}_z^{xy} - \overrightarrow{m}_z^{yx}$ is odd under reversal of the roles of x and y, and hence under flipping the sign of β. This means that ∂_β appears an odd number of times in each forest in D_1, D_2, and D_3, and thus in each union of wheels in W. But this means that every term in W contains a wheel with an odd number of legs, and such wheels vanish by the AS relation. Hence $W = 0$ and thus $\operatorname{div}_\alpha D_3 = 0$, as required. $\square$

References

[B-N1] D. Bar-Natan. On the Vassiliev knot invariants. *Topology* **34** (1995), 423–472.

[B-N2] D. Bar-Natan. Vassiliev homotopy string link invariants. *Jour. of Knot Theory and its Ramifications* **4** (1995), 13–32.

[B-N3] D. Bar-Natan. Non-associative tangles. in *Geometric topology* (proceedings of the Georgia international topology conference), (W. H. Kazez, ed.), 139–183, Amer. Math. Soc. and International Press, Providence, 1997.

[BGRT] D. Bar-Natan, S. Garoufalidis, L. Rozansky and D. P. Thurston. Wheels, wheeling, and the Kontsevich integral of the unknot. *Israel Journal of Mathematics* **119** (2000), 217–237, arXiv:q-alg/9703025.

[Ja] N. Jacobson. *Lie algebras.* Dover, New-York, 1979.

[Ki] R. Kirby. A calculus of framed links in S^3. *Invent. Math.* **45** (1978), 35–56.

[Ko] M. Kontsevich. Vassiliev's knot invariants. *Adv. in Sov. Math.,* **16** (1993), 137–150.

[Le] T. Q. T. Le. An invariant of integral homology 3-spheres which is universal for all finite type invariants. in Solitons, geometry and topology: on the crossroad (V. Buchstaber and S. Novikov, eds.). AMS Translations Series 2, Providence.

[LMMO] T. Q. T. Le, H. Murakami, J. Murakami and T. Ohtsuki. A three-manifold invariant via the Kontsevich integral. Max-Planck-Institut Bonn preprint, 1995.

[LM1] T. Q. T. Le and J. Murakami. The universal Vassiliev-Kontsevich invariant for framed oriented links. *Compositio Math.* **102** (1996), 41–64, arXiv:hep-th/9401016.

[LM2] Parallel version of the universal Vassiliev-Kontsevich invariant. *J. Pure and Appl. Algebra* **121** (1997), 271–291.

[Oh] T. Ohtsuki. Finite type invariants of integral homology 3-spheres. *Jour. of Knot Theory and its Ramifications* **5** (1996), 101–115.

[Å-I] D. Bar-Natan, S. Garoufalidis, L. Rozansky and D. P. Thurston. The Århus integral of rational homology 3-spheres I: A highly non trivial flat connection on S^3. *Selecta Math.*, arXiv:q-alg/9706004 (to appear).

D. Bar-Natan
The Hebrew University
Institute of Mathematics
Giv'at-Ram, Jerusalem 91904
Israel
e-mail: drorbn@math.huji.ac.il
http://www.ma.huji.ac.il/~drorbn

S. Garoufalidis
Brandeis University
Department of Mathematics
Waltham, MA 02254-9110
USA

current address:
School of Mathematics
Georgia Institute of Technology
Atlanta, GA 30332-0160
USA
e-mail: stavros@math.gatech.edu
http://www.math.gatech.edu/~stavros

L. Rozansky
University of Illinois at Chicago
Department of Mathematics
Statistics and Computer Science
Chicago, IL 60607-7045
USA

current address:
UNC-CH
Department of Mathematics
CB 3250 Phillips Hall
Chapel Hill, NC 27599-3250
USA
e-mail: rozansky@email.unc.edu

D. P. Thurston
University of California at Berkeley
Department of Mathematics
Berkeley, CA 94720-3840
USA

current address:
Harvard University
Department of Mathematics
Cambridge, MA 02138
USA
e-mail: dpt@math.harvard.edu

To access this journal online:
http://www.birkhauser.ch