ADV. THEOR. MATH. PHYS.
Volume 22, Number 5, 1289–1345, 2018

Cubic planar graphs and Legendrian surface theory

David Treumann and Eric Zaslow

We study Legendrian surfaces determined by cubic planar graphs. Graphs with distinct chromatic polynomials determine surfaces that are not Legendrian isotopic, thus giving many examples of non-isotopic Legendrian surfaces with the same classical invariants. The Legendrians have no exact Lagrangian fillings, but have many interesting non-exact fillings.

We obtain these results by studying sheaves on a three-ball with microsupport in the surface. The moduli of such sheaves has a concrete description in terms of the graph and a beautiful embedding as a holomorphic Lagrangian submanifold of a symplectic period domain, a Lagrangian that has appeared in the work of Dimofte-Gabella-Goncharov [DGG1]. We exploit this structure to find conjectural open Gromov-Witten invariants for the non-exact filling, following Aganagic-Vafa [AV, AV2].

1. Introduction and summary

An exact Lagrangian filling of a Legendrian in a cosphere bundle determines a family of constructible sheaves [NZ]. In this paper, we explore a curious counterpoint: Legendrian surfaces that give rise to beautiful moduli spaces of constructible sheaves, but have no exact fillings whatsoever.

For one-dimensional Legendrians, the families of fillings give the whole moduli space of constructible sheaves the rich structure of a cluster variety. This observation leads to strong new lower bounds on the number of Hamiltonian isotopy classes of exact Lagrangian surfaces filling Legendrian knots [STW, STWZ]. It is natural to wonder what structures are determined by Legendrian surfaces.

A fundamental example is related to the *Harvey-Lawson cone*, a singular exact Lagrangian in $\mathbf{R}^6$. Nadler studied the microlocal category of this cone in [N2], proving that it is equivalent to a category of constructible sheaves on $\mathbf{R}^3$ and furthermore equivalent to the category of coherent sheaves on the pair of pants $\mathbf{P}^1 \setminus \{0, 1, \infty\}$. We observe in this paper that this implies the non-fillability of the Legendrian boundary of the Harvey-Lawson cone. Nadler's example is fundamental for us: we prove similar results for a broad class of Legendrian surfaces of any genus.

We also use these moduli spaces to distinguish Legendrian surfaces with the same classical invariants. Generally, much less is known about Legendrian surfaces than Legendrian knots. The first examples of inequivalent Legendrian knots with the same classical invariants were obtained by Chekanov and distinguished by the Chekanov-Eliashberg differential graded algebra (dga), which can be computed combinatorially. In the case of Legendrian surfaces, the dga is much more difficult to compute. The best technology for enumerating the necessary holomorphic disks is Ekholm's gradient flow trees [E]. Very recent work of Rutherford and Sullivan [RS] exploits this technology to reduce the computation of the dga of many Legendrian surfaces to a (still difficult) combinatorial procedure. Our results, where the symplectic analysis is subsumed by the local, combinatorial nature of sheaves, demonstrate the strength of microlocal sheaf techniques in symplectic topology.

In addition, we are able to use the Lagrangian structure of this moduli space inside a period domain to perform an Aganagic-Vafa-style mirror symmetry [AV2] and compute conjectural open Gromov-Witten invariants of the obstructed non-exact Lagrangians which fill our Legendrian surfaces.

We now explain these results in some more detail.

1.1. Legendrian surfaces, fillings, and constructible sheaves

Our Legendrian surfaces lie in an open domain of S^5 contactomorphic to the cosphere bundle $T^\infty\mathbf{R}^3$ and to the jet bundle $J^1(S^2)$. They are genus-g surfaces double-covering their base projection to S^2 with $2g+2$ branch points. Such a Legendrian S can be defined from a cubic planar graph $\Gamma \subset S^2$, by constructing a wavefront projection of S in $S^2 \times \mathbf{R} \cong \mathbf{R}^3 \setminus \{0\}$ that is generically two-to-one over the base projection to S^2, but one-to-one over Γ. (The original S, not just its wavefront, is in fact still $2:1$ over the edges and only $1:1$ over the $2g+2$ vertices.)

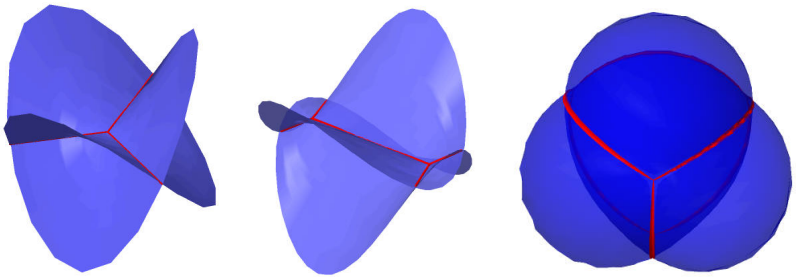

Figure 1.1.1: Wavefronts near a vertex, edge, and at right for Γ a tetrahedron.

Remark 1.1. The construction can be motivated as a higher-dimensional analogue of the Stokes Legendrian [STWZ] encoding the wild character varieties of complex curves near an irregular singularity of an holomorphic ODE. For differential equations of second order, such a Legendrian can be specified by a finite set of points on the boundary, the locations where the asymptotic behavior of the two solutions switch. In three dimensions, those points are replaced by cubic graphs (indeed the story in this paper has a generalization to arbitrary 3-manifolds with cubic graphs drawn on their boundary) — it would be very interesting to find a family of second-order PDEs whose asymptotic behavior exhibited the kind of Stokes phenomenon modeled on these graphs.

In this introduction, we will mainly restrict our attention to *simple* (no loops or multiple edges) cubic planar graphs, although the Legendrian surface associated to Γ is of interest more generally.

Given such a Legendrian $S \subset T^\infty\mathbf{R}^3$ we can take it as singular-support data for a category of constructible sheaves. This category of constructible

sheaves is equivalent [NZ, N] to a Fukaya category in $T^*\mathbf{R}^3$ with asymptotic conditions on geometric branes defined by S. We give an explicit description of the moduli space $\mathcal{M}_r$ of objects of this category that have microlocal rank r, focusing on $\mathcal{M} := \mathcal{M}_1$. In [KS], sheaves in $\mathcal{M}_r$ are called "pure," or "simple" if $r = 1$. In the Fukaya category, points of $\mathcal{M}_r$ correspond to stacks of r basic branes ending on S — see Remark 1.5.

A point of $\mathcal{M}$ is easy to describe: think of the faces of Γ as countries on a map of the globe S^2, and color each one with a point in $\mathbf{P}^1$ so that countries which share an edge border have different colors. The set of such choices provides a framed version of the moduli space: to get $\mathcal{M}$, we must quotient by the automorphism group PGL_2 of $\mathbf{P}^1$, which acts freely. If we put $\widehat{\Gamma}$ for the dual graph, this means $\mathcal{M}$ is the set of graph colorings of $\widehat{\Gamma}$, with colors chosen in $\mathbf{P}^1$, modulo PGL_2. As a variety, $\mathcal{M}$ is defined over the integers and we may count its points over finite fields:

Proposition 1.2. *Let Γ be a simple, cubic planar graph, $\widehat{\Gamma}$ its dual graph. Let $P_{\widehat{\Gamma}}$ denote the chromatic polynomial, whose value $P_{\widehat{\Gamma}}(c)$ is the number of colorings of $\widehat{\Gamma}$ with c colors. Let $\mathbf{F}_q$ be a field with q elements. Then*

$$\#\mathcal{M}/\mathbf{F}_q = \frac{1}{q^3 - q} \cdot P_{\widehat{\Gamma}}(q+1).$$

Proof. The argument $q+1$ is the number of $\mathbf{F}_q$-points of $\mathbf{P}^1$. The denominator is the order of $\mathrm{PGL}_2(\mathbf{F}_q)$. □

It is known [B] that if G has n vertices and e edges, then $P_G(x) = x^n - ex^{n-1} + O(x^{n-2})$. If Γ is a cubic planar graph with v vertices, e edges and f faces, then it is easy to show that

$$v = 2g + 2, \quad e = 3g + 3, \quad f = g + 3,$$

where g is the genus of the corresponding Legendrian surface S.

Theorem 1.3. *Let $S \subset T^\infty\mathbf{R}^3$ be the genus-g Legendrian surface defined by a simple, cubic planar graph Γ. Then S has no smooth oriented graded exact Lagrangian fillings in $\mathbf{R}^6$.*

Proof. An oriented 3-manifold L whose boundary is a genus-g surface has $b_1(L) \geq g$. One automatically has $w_2(L) = 0$, so that if L is a graded exact Lagrangian filling, the inverse of microlocalization [NZ] defines a torus chart $(\mathbb{G}_\mathrm{m})^{\times g} \hookrightarrow \mathcal{M}$ as in [STWZ]. Thus over $\mathbf{F}_q$, a filling gives $(q-1)^g = q^g -$

$gq^{g-1} + \cdots$ distinct points of $\mathcal{M}$. On the other hand, recalling that for $\widehat{\Gamma}$ we have $n = g + 3$ and $e = 3g + 3$, we get

$$\begin{aligned} \#\mathcal{M}/\mathbf{F}_q &= \frac{1}{q^3 - q}\left((q+1)^{g+3} - (3g+3)(q+1)^{g+2} + \cdots\right) \\ &= q^g - 2g \cdot q^{g-1} + \cdots \end{aligned}$$

Taking q large, we see $\mathcal{M}$ is not big enough to accommodate any torus chart! □

Nevertheless, we can construct smooth fillings that are not exact. Suppose that Γ is drawn on the surface of the ball D^3. A *foam* filling Γ is a singular surface $\mathbf{F} \subset D^3$ with codimension-one singularities that look like the letter "Y" times an interval, and codimension-two singularities (points) that look like the cone over the 1-skeleton of a tetrahedron, and such that $\mathbf{F} \cap \partial D^3 = \Gamma$.

Construction 1.4. From a foam in D^3 we construct an exact *singular* Lagrangian filling of S. The singularities are of Harvey-Lawson type, and can be smoothed away (in different "phases") to smooth non-exact Lagrangian fillings.

A foam is the same combinatorial structure that configurations of soap bubbles have — the singularities are the Plateau borders of the soap film. It is interesting to speculate whether, if we were to impose Plateau's laws on the foam (that the sheets of soap have constant mean curvature and meet at equilateral angles), the Lagrangian we construct could be chosen to be sLags. The singular Lagrangians are easy to describe, they are branched double covers of the 1-skeleton of the foam. The smooth Lagrangians we construct are not just non-exact but *obstructed*: they are not objects in a Fukaya category if one does not introduce a "bounding cochain."

Remark 1.5. In the Fukaya category constructed in [NZ], only exact branes were considered, but in most of our examples such exact Lagrangians are not available to support points of $\mathcal{M}$ or $\mathcal{M}_r$. Instead, points of $\mathcal{M}_r$ morally correspond to geometric Lagrangians, usually non-exact and even obstructed, together with a $\mathrm{U}(r)$-connection. In the two-dimensional version of this analogy exploited in e.g. [STWZ], the connection can be taken to be flat, but in three dimensions when the Lagrangian is obstructed it must obey a more complicated Maurer-Cartan equation. This is part of the theory of bounding cochains and curved A_∞-categories, which so far has not influenced microlocal sheaf theory.

In any case, there is a map from $\mathcal{M}$ that corresponds to "restrict a local system to its boundary" when an exact Lagrangian does exists, and makes sense over the whole moduli space. It is a special case of the μhom-functor of [KS], that was called "microlocal monodromy" in [STZ2]. This map

$$\phi : \mathcal{M} \to H^1(S, \mathbf{C}^*) \tag{1.1.1}$$

is isomorphic to one considered in [DGG1], and is known to be Lagrangian with respect to the symplectic structure on $H^1(S, \mathbf{C}^*)$ induced by the intersection form — something that follows in this sheaf-theoretic context from the general results of [BD, ST1].

1.2. Generalized Aganagic-Vafa mirror symmetry

Recall that Aganagic-Vafa [AV] fix a Lagrangian brane, then use the equation of the moduli space of that brane in the resolved conifold to define conjectural open Gromov-Witten invariants. In [AKV] the construction was extended to different "phases" and "framings" of the brane, and to other toric Calabi-Yau three-folds, including $\mathbf{C}^3$. In [AV3] the authors apply this construction to conormals of knots and call this "generalized SYZ mirror symmetry." Our method generalizes this technique to three-dimensional branes that look nothing like tori, so we refer to it as "generalized AV mirror symmetry."

The original AV construction identifies a distinguished set of coordinates $x = e^u, y = e^v$ on an ambient symplectic $(\mathbf{C}^*)^2$ and uses the equation of the moduli space $F(u, v) = 0$ to solve $v = -\partial_u W$ (the choice of sign is historical), where W has a four-dimensional interpretation as a superpotential, but also as the generating function for open disk invariants. The procedure depends on a choice of "framing" which distinguishes the second coordinate, as the so-called "phase" only identifies one of the two.

The method can be interpreted in the following way. The moduli space $\mathcal{M}$ of the brane is a Lagrangian submanifold in $H^1(S, \mathbf{C}^*) = (\mathbf{C}^*)^{2g}$. A *phase* is a geometric Lagrangian, an oriented 3-manifold L, contained in $\mathbf{R}^6$ with boundary on S. A phase determines a Lagrangian map $H^1(L, \mathbf{C}^*) \to H^1(S, \mathbf{C}^*)$, and we define a framing to be an extension of this to a symplectic map $T^*H^1(L, \mathbf{C}^*) \to H^1(S, \mathbf{C}^*)$. Then it makes sense to try to express $\mathcal{M}$ as the graph of the differential of a multivalued transcendental function $W : H^1(L, \mathbf{C}^*) \to \mathbf{C}$.

> *We conjecture that W is the generating function of open Gromov-Witten disk invariants.*

Currently, the Lagrangians considered in this paper fall outside the limited class for which open Gromov-Witten invariants are rigorously defined, so the conjecture is not strictly precise as it stands. However, there is recent progress in the work of Solomon and Tukachinsky — see Remark 4.7. Meanwhile, we can make a precise integrality conjecture, following Ooguri-Vafa [OV]:

> There are integers $a(d,f) \in \mathbf{Z}$ and an order-two element $\varepsilon_f \in H^1(L, \mathbf{C}^*)$, indexed by homology classes $d \in H_1(L, \mathbf{Z})$ and a framing f, such that

$$W(x) = \sum_d a(d,f) \mathrm{Li}_2\left((\varepsilon_f x)^d\right) \tag{1.2.1}$$

> where $(-)^d$ denotes the monomial function $H^1(L, \mathbf{C}^*) \to \mathbf{C}$ corresponding to d and $\mathrm{Li}_2(z) = \sum \frac{1}{n^2} z^n$ is the classical dilogarithm function.

Identities involving $\mathrm{Li}_2(z)$, $\mathrm{Li}_2(1/z)$, and other arguments related by Möbius transformations prevent equation (1.2.1) from determining $a(d,f)$ uniqely. We expect that by restricting the sum to a strictly convex ("Mori") cone of elements in $H_1(L, \mathbf{Z})$ that support holomorphic disks, the coefficients $a(d,f)$ become uniquely determined and are the physical BPS numbers. The translation by ε_f functions as a kind of mirror map (change of variables to write superpotential correctly), that we do not have a great explanation for.

Before giving some examples let us discuss coordinates. The choice of phase allows us to choose coordinates $(v_1, \ldots, v_g)$ on the universal cover of $H^1(S, \mathbf{C}^*)$ that cut out $H^1(L, \mathbf{C})$, and a framing nails down conjugate coordinates $u = (u_1, ..., u_g)$. These choices identify the universal cover with $\mathbf{C}^{2g} \cong T^*\mathbf{C}^g$, and W is the solution to $v_i = \pm \partial_{u_i} W$.

We implement this generalized AV mirror symmetry explicitly in several examples.

1.3. Examples

- When Γ is the complete graph with four vertices, we have $f = g + 3 = 4$, so $g = 1$ and the filling is a solid torus. It is the Aganagic-Vafa brane in $\mathbf{C}^3$. Then $\mathcal{M} \subset (\mathbf{C}^*)^2$ is defined by

$$x + y = 1,$$

a pair of pants, and here we have implicitly chosen a phase and (zero) framing to write $x = e^u, y = e^v$. Then $v = -\partial_u W$ identifies $W = \mathrm{Li}_2(x) := \sum_{n>0} \frac{1}{n^2} x^n$, as in the work of Aganagic-Vafa [AV]. The procedure is identical to theirs in this example, so different framings — making the replacement $u \to u + pv$ — give the framing-dependent disk invariants computed by Aganagic-Klemm-Vafa [AKV, Section 6.1].

- When Γ is a triangular prism, $g = 2$ and the moduli space $\mathcal{M}$ is cut out by two equations:

$$x_1 + y_1 = 1, \qquad x_2 + y_2 = 1.$$

Then $\mathcal{M}$ is a product of two copies of the pair of pants, and zero framing $(x_i = e^{u_i}, y_i = e^{v_i})$ gives $W = \mathrm{Li}_2(x_1) + \mathrm{Li}_2(x_2)$. "Diagonal" framings $u_1 \to u_1 + p_1 v_1$, $u_2 \to u_2 + p_2 v_2$ lead to different BPS numbers, but exactly as above. Here, however, we also have the possibility of non-diagonal framings such as

$$u_1 \to u_1 + pv_2, \qquad u_2 \to u_2 + pv_1.$$

For example, when $p = 1$ we find $W(x_1, x_2) = \mathrm{Li}_2(x_1) + \mathrm{Li}_2(x_2) - \mathrm{Li}_2(x_1 x_2)$. More generally, framings are determined by a symmetric 2×2 integer matrix $M_i{}^j$ via $u_i \to u_i + M_i{}^j v_j$. We consider these examples in Section 5.3. While we generally cannot write the superpotential W in closed form, we can perform strict integrality checks: in all cases considered, we derive integer BPS numbers $a(d, f)$ using the Ooguri-Vafa $1/d^2$ multiple cover formula [OV].

- When Γ is the union of vertices and edges of a cube, $g = 3$ and we can similarly define multi-coordinates (u, v) for $\mathbf{C}^6 \cong T^*\mathbf{C}^3_u$. Then (the universal cover of) $\mathcal{M}$ is a graph of $-dW(u)$, where

$$W = \mathrm{Li}_2(x_1) + \mathrm{Li}_2(x_2) + \mathrm{Li}_2(x_3) - \mathrm{Li}_2(x_1 x_2) - \mathrm{Li}_2(x_1 x_3),$$

with $x_i = e^{-u_i}$. The answer is written in terms of integer linear combinations of dilogarithms, meaning the conjectural disk invariants obey integrality.

Remark 1.6. Cubic planar graphs define triangulations of the plane, and graph mutations correspond to Pachner 2-2 moves. This and the appearance of dilogarithms comes from an intimate relationship between the present

work and cluster theory. In future work with Linhui Shen [STZ1], we will describe the process of mutation and the relationship of W to the DT series in cluster theory.

1.4. Relation to previous works in physics

While preparing this document, we learned of prior constructions with signifcant overlap to our own. We will try to explain the connections, similarities and differences to the present work. While the works in Section 1.4.2 appeared first, we describe the relations in reverse chronological order due to the greater similarity in the approach of the latter works.

1.4.1. The work of Cecotti, Córdova, Espahbodi, Haghighat, Neitzke, Rastogi and Vafa. In the series of works [CCV, CNV, CEHRV], the above-named authors consider the dimensional reduction of the theory on a stack of two overlapping M5-branes wrapping a three-dimensional spacetime cross a Lagrangian three-fold. The two recombine into a single-M5 brane, the Lagrangian being a branched double cover over $\mathbf{R}^3$. The authors also consider the Harvey-Lawson cone and singular and smooth tangles — see in particular Section 5.1.2 of [CCV] and Section 3.1 of [CEHRV].

Those authors consider Seifert surfaces for the tangles, which they use to construct actions for the effective three-dimensional theory.[1] This is an interacting gauge theory, which in the Coulomb branch has a U(1) gauge field for each homology loop of the three-manifold, with Chern-Simons levels determined by the self-linking matrix of the tangle. Chiral matter and superpotentials have M-theoretic descriptions via holomorphic maps with various boundary conditions. The guiding principle is that the effective physics should be independent of the Seifert surface used to describe the three-manifold and construct the action: different Seifert surfaces thus determine dual theories. Equivalence of these theories is shown to give a physical explanation of the wall-crossing formula of Kontsevich-Soibelman.

In the three-dimensional theory, BPS states are formed by M2-branes ending on the M5-brane, so holomorphic disks bounding circles in the Lagrangian. The circles surround two strands of the tangle, so massless states arise for singular tangles, when the disk shrinks to zero size. The fields of the three-dimensional theory that create such states are chiral multiplets.

M2-brane instantons give rise to superpotential terms which couple the massless particles. The boundary of the three-ball M2 brane is a two-sphere

[1] We use the word "action" to avoid overloading the word "Lagrangian."

which projects to a polygon, the sides of which label the different coupled chiral multiplets.

1.4.2. The work of Dimofte, Gabella, Gaiotto, Goncharov, Gukov and Hollands. As just mentioned in Section 1.4, the above-named authors also consider M5-branes defined by a three-dimensional Lagrangian [DGG1, DGG2, DGH], usually a hyperbolic manifold such as a branched double cover of a knot. (These authors also consider n-fold covers, but not through a description of its branch locus.) They define an isomorphic moduli space $\mathcal{M}$ and map (1.1.1). They investigate the relationship between, on the one hand, Chern-Simons theory and hyperbolic three-manifolds, and on the other, three-dimensional supersymmetric field theories, their supersymmetric indices, superpotentials, and relation to four-dimensional BPS states. Some sketches are given in Appendix 5.4.

The authors construct a moduli space of G-local systems on a three-manifold L with an ideal triangulation — i.e. a decomposition into hyperbolic tetrahedra with boundary surfaces — where $G = \mathrm{PGL}_n(\mathbf{C})$. Gluing these together, one arrives at either a closed three-manifold or one with a boundary surface. For example, a boundary torus is relevant to the knot complement of a hyperbolic knot, and this is a prime example. The tetrahedra are truncated and thus have four triangles near the vertices, and one assigns an element of $\mathbf{P}^1$ at each triangle. The space $\mathrm{Loc}_G(\partial L)$ is known to be symplectic, and the restriction of a local system to the boundary embeds $\mathrm{Loc}_G(L)$ as a Lagrangian submanifold.

In forthcoming work with Shen [STZ1], we use cluster theory to determine the wavefunctions/partition functions appearing in these physical settings, including the dependence on framings. The present work gives a conjectural relationship between these functions and the framing-dependent superpotentials encoding open Gromov-Witten invariants.

Acknowledgements. It is a pleasure to thank Roger Casals, Bohan Fang, Xin Jin, Melissa Liu, Emmy Murphy, David Nadler, and Linhui Shen for sharing their insights, time and suggestions. Roger Casals and Emmy Murphy were quickly able to settle many questions raised in an early draft of this paper. We also thank Tudor Dimofte and Cumrun Vafa for helpful discussions about their respective joint works. We thank Jake Solomon and Sara Tukachinsky for discussing the role of framings in their work. D.T. is supported by NSF grant DMS-1510444. E.Z. is supported by NSF grant DMS-1406024.

2. The hyperelliptic wavefront

As in the introduction, let $\Gamma \subset S^2 \subset \mathbf{R}^3$ be a cubic planar graph. We assume each edge is smoothly embedded, and that in the tangent space to a vertex the three edges are linearly independent and do not lie in a half-plane. We write v, e, f for the number of vertices, edges, and faces of Γ. Euler's relation gives $v - e + f = 2$ and the cubicness of Γ gives $3v = 2e$, so there is an integer g such that

$$f = g + 3, \qquad v = 2g + 2, \qquad e = 3g + 3 \tag{2.0.1}$$

Let $\pi : S \to S^2$ be the connected oriented double cover of S^2, branched over the vertices of Γ — the Riemann-Hurwitz formula shows S has genus g. Write ι_S for the nontrivial Deck transformation of π.

Definition 2.1. We define a *hyperelliptic wavefront* modeled on Γ to be a map $i : S \to \mathbf{R}^3$ with the following properties:

1) The image of S does not contain the origin so that i can be written in spherical coordinates $S \to S^2 \times \mathbf{R}_{>0}$.
2) In those coordinates we have $i(x) = (\pi(x), r(x))$ where π is the covering map and $r : S \to \mathbf{R}_{>0}$.
3) The map r obeys $r(\iota_S(x)) = 1/r(x)$, with $r(x) = 1$ exactly on the preimage of Γ.
4) Near each vertex of Γ, we may find coordinates (x, y) on S^2 and (u, v) on S such that (π, r) is given by

$$x(u,v) = u^2 - v^2, \quad y(u,v) = 2uv, \quad r(u,v) = \exp\left(\frac{2}{3}u^3 - 2uv^2\right). \tag{2.0.2}$$

   In other words, after setting $w = u + iv$, we have $x + iy = w^2$ and $\log(r) = \frac{2}{3}\mathrm{Re}(w^3)$, which give the front projection of the exact complex Lagrangian $\{(w^2, w)\} \subset T^*\mathbf{C}$.
5) r has exactly $(2g + 2) + 2(g + 3)$ critical points: the $2g + 2$ critical points of π, and a unique maximum and unique minimum over each face of Γ.

There is a hyperelliptic wavefront modeled on every Γ. One way to construct it is to identify S^2 with the Riemann sphere (with holomorphic coordinate z, say), and choose a Strebel differential $f(z)dz^2$ whose

non-closed trajectories are the edges of Γ. Such a differential will have a unique quadratic pole p_i in each face, and the Strebel condition implies that $\tilde{r}(z) = \exp(\pm \mathrm{Re} \int \sqrt{f(z)}dz)$ obeys conditions (1)-(5) — except it takes the values $\{0, \infty\}$ on the poles p_i. We obtain r by damping $\tilde{r}$ near each pole p_i. We call this the "Strebel model" for the hyperelliptic wavefront.

The local model at a vertex (2.0.2) gives an immersed hypersurface with double points along $\mathrm{Re}(w^3) = 0$, shown here in red.

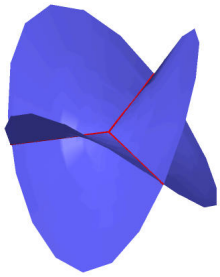

It fails to be an immersion at $w = 0$, but it has a well-defined normal direction there: radially outwards. The oriented Legendrian lift (also called the oriented Nash blowup) gives a smooth Legendrian embedding to $T^{\infty}\mathbf{R}^3 \cong \mathbf{R}^3 \times S^2$. The Nash blowup construction shows that one may recover a hyperelliptic wavefront $i : S \to \mathbf{R}^3$ from the image $\Psi := i(S)$ as a subset of $\mathbf{R}^3$.

Remark 2.2. A few comments are in order.

1) The wavefront immersion is not generic among front projections — it is possible to "push three swallowtails" out of any vertex, either on the top or the bottom of the figure. In other words it is the critical front of a "bifurcation of fronts." In low dimensions Arnold gave a classification of such bifurcations, and this one is called D_4^-.

2) When Γ is not cubic, but the tangential angles between adjacent edges at a degree d vertex are $(360/d)^\circ$, it is still possible to construct S and a wavefront immersion to $\mathbf{R}^3$, but it is no longer smooth over the vertices. The local model (4) at a degree d vertex looks like the Legendrian lift of the plane curve singularity $w^2 = z^{d-2}$.

3) We define the "Lagrangian projection" of a hyperelliptic wavefront to be the projection to T^*S^2 whose cotangent coordinates are given by the gradient of $\log(r)$. The Lagrangian projection is a Lagrangian immersion with $g + 3$ double points. In the Strebel model, there is a neighborhood $U_\Gamma \subset S^2$ of Γ for which the Lagrangian projection

$\pi^{-1}(U_\Gamma) \to T^*S^2 = T^*\mathbf{P}^1$ is a holomorphic embedding, but the immersion $S \to T^*S^2$ cannot be made holomorphic.

4) Suppose M is a 3-manifold, $\Sigma \subset M$ is a two-sided smooth surface in M, and $\Gamma \subset \Sigma$ is a cubic graph. If S denotes the double cover of Σ branched over the vertices of Γ, we have an associated wavefront $i : S \to M$, with $i(S) \cap \Sigma = \Gamma$. Of course, the name "hyperelliptic" is less appropriate when Σ is not a sphere. An interesting general class of examples are suggested by the Legendrian knots in the cocircle bundles over surfaces, considered in [STWZ]. If $\partial M \times [0,1] \to M$ is a collar neighborhood of the boundary, then we can set $\Sigma = \partial M \times \{\frac{1}{2}\}$ — we refer to these as "Stokes wavefronts of degree 2." The wavefronts of Definition 2.1 are included in the class of Stokes wavefronts, they are just the case when M is a ball.

When Γ is the edge graph of a tetrahedron, the hyperelliptic wavefront is as shown in Figure 2.0.1 below.

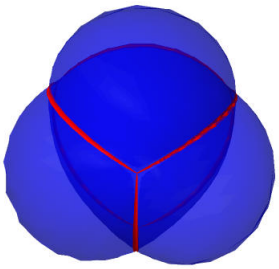

Figure 2.0.1: The hyperelliptic wavefront defined by the edge graph (in red) of a tetrahedron.

2.1. The extended Legendrian S^+

In [STZ2, §2.2.1], we associated to each front diagram $\Phi \subset \mathbf{R}^2$ a Legendrian graph Λ^+, obtained from the Legendrian lift Λ of Φ by attaching a Legendrian chord joining the two points projecting to a crossing. The cone over Λ^+ is the smallest Lagrangian containing the cone over Λ that is closed under addition. (Thus Λ^+ is "cragged" in the language of [SS].)

We may consider a similar extended form of S, which we denote by S^+. If i is an edge of Γ, its preimage is parametrized by a circle $\gamma_i : S^1 \to S$, with a distinguished orientation we discuss in §4.6. If the edge i is not a loop, then γ_i is an embedded circle, while if i is a loop from vertex v to itself, γ_i is

an immersion with a single double point at $\pi^{-1}(v)$. We attach a Legendrian disk to S whose boundary is along γ_i.

To define D_i, first note that each point p along i has two associated outward conormal rays, one along each branch of the wavefront immersion. D_i is the subset of the cosphere bundle over i which, at a point $p \in i$, consists of the acute angle of the conormal rays between those two. Note this angle shrinks to zero at the endpoints of i, where D_i is just a point. Thus D_i is the disk formed by the family of intervals over the interval i, collapsing at the boundary.

2.2. Edge and disk moves

If i is an edge of Γ that is not a loop, we may construct a new graph Γ' by changing Γ in a neighborhood of i according to the following diagram

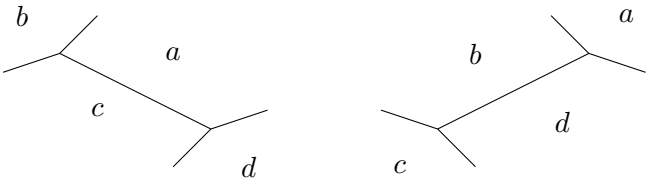

The faces of Γ and of Γ' are in natural bijection with each other. By performing this move at different edges one can get from any graph to any other graph. Indeed this is the sense in which cubic planar graphs label the top-dimensional cells in the Harer-Penner-Mumford-Thurston complex for $M_{0,n}$, which is connected.

Let $U_i \subset S^2$ be a neighborhood of i. If i is not a loop, the preimage of U_i in S is an annulus, and its preimage in S^+ is an annulus with a disk attached — let us denote them by A and A^+. The Lagrangian projection of A^+ is an embedding into T^*U_i — it is an exact, *singular* Lagrangian of the kind that has been considered by [Y]. Yau considers a "disk move" which is part of our model for cluster transformations in [STWZ, STW]. In the present context these are related:

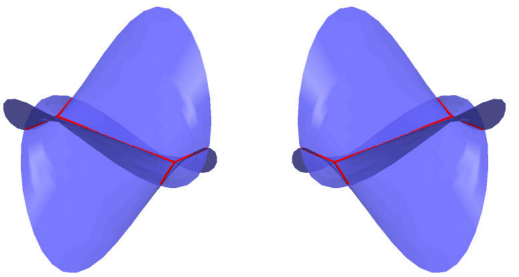

The two exact Lagrangian surfaces on either side of a disk move are part of the "focus-focus" family of Lagrangian annuli. But this family not provide for a good local model, as it does not fix the boundary (we thank Emmy Murphy for alerting us to this). Now if there is another edge i' meeting i in exactly one vertex (this is always possible if the graph before and after the edge move has no multiple edges), then we may find a one-parameter family of (not necessarily exact) Lagrangian immersions $S_t \to T^*S^2$ with the following properties: 1) $S_t = S$ at $t = 0$ and S' at $t = 1$; 2) S_t is independent of t outside of a neighborhood of $i \cup i'$; 3) for each t, $S_t \to T^*S^2$ is an embedding over the neighborhood of $i \cup i'$.

Proposition 2.3. *Let S and S' be two Legendrian surfaces associated to Γ and Γ', with the same number of vertices. Let j and j' denote the exact Lagrangian immersions into T^*S^2, each with $g + 3$ double points. Then j and j' are isotopic to each other through (non-exact) Lagrangian immersions that do not increase or decrease the number of double points.*

In particular, the Legendrians S and S' must have the same classical invariants.

2.3. The Ekholm-Honda-Kálmán Lagrangians

Let n be an integer, which for simplicity we will assume is odd. Let $\Lambda \subset S^3$ be a Legendrian $(2, n)$-torus knot. Two families of $\frac{1}{n+1}\binom{2n}{n}$ exact Lagrangian fillings of Λ have been produced, first in [EHK] and later in [STWZ]. Roughly speaking, [EHK] produces this family of fillings by a "elementary Lagrangian cobordism" technique, and [STWZ] by an "alternating strand diagram" technique. We sketch here a third construction of a family of $\frac{1}{n+1}\binom{2n}{n}$ fillings, in terms of hyperelliptic wavefronts, and an equivalence between these fillings and the EHK fillings.

2.3.1. Rainbow, spirograph and Ng projection. As in [STWZ] it is convenient to draw the $(2, n)$-torus knot in its "spirograph projection," with $n + 2$ crossings around a circle — let us recall the correspondence here. Let Λ_n be the Legendrian $(2, n)$ torus knot. There are two contactomorphisms

$$T^{\infty,-}\mathbf{R}^2 \hookrightarrow S^3 \qquad\qquad T^{\infty}\mathbf{R}^2 \hookrightarrow S^3,$$

onto open subsets of S^3, and Λ_n is contained in the image of both. The front projection to $\mathbf{R}^2$ has n crossings and 4 cusps in the first projection,

and $n+2$ crossing but no cusps in the second projection. For example when $n = 3$:

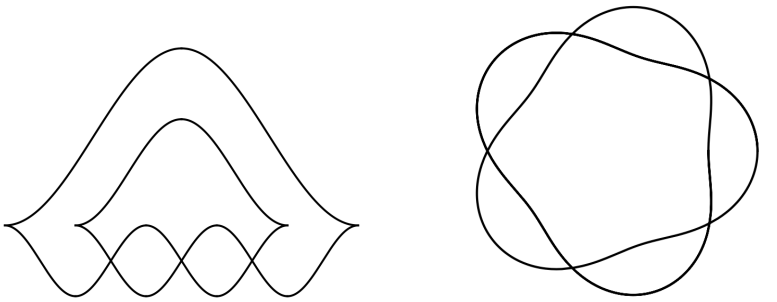

The rainbow projection is Reidemeister-equivalent to the following reflection-invariant front, which we will call the *Ng projection*:

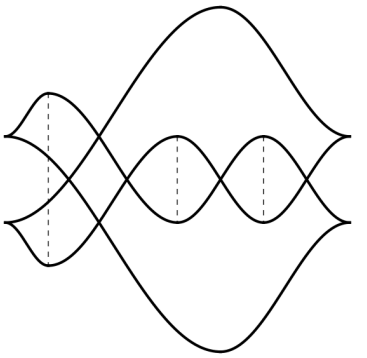

This front is homeomorphic to one whose Lagrangian projection is [EHK, Fig. 2, p. 5] The dashed lines indicate the Reeb chords that loc. cit. uses to construct saddle cobordisms.

2.3.2. Construction of fillings. We replace S^2 with a disk D^2, and Γ with a trivalent tree with n vertices, whose $n+2$ leaves reach the boundary of the disk at those $n+2$ crossings — we further assume that in a collar neighborhood of ∂D^2, each leaf is a radial interval. A branched double cover S of D^2 over the vertices of the tree is a genus $\frac{1}{2}(n-1)$ surface with one boundary component. To give a hyperelliptic wavefront of S into $D^2 \times \mathbf{R}$, we give a function $f : S \to \mathbf{R}$ analogous to the logarithm function r of Definition 2.1. We replace condition 2.1(5) by the condition that f has no critical points besides the vertices of Γ, and that in a collar neighborhood of D^2 it is linear in the radial coordinate of D^2. These conditions ensure that the

Lagrangian projection — i.e. the map $S \to T^*D^2$ given by $(\pi, \nabla f)$, as in Remark 2.2(3) — is an embedding, with Legendrian boundary at $T^*D^2|_{\partial D^2}$.

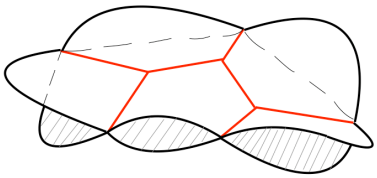

Figure 2.3.1: A tree graph Γ in red, and wavefront for an exact Lagrangian filling of a Legendrian trefoil knot.

2.3.3. Comparison to the Ekholm-Honda-Kálmán construction. Ekholm-Honda-Kálmán observed that a Reeb chord in the Ng projection of the $(2, n)$ torus knot determines a "saddle cobordism" [EHK, §6.5] to the $(2, n-1)$ torus link, with one less Reeb chord in its Ng projection. A total ordering of the Reeb chords gives a cobordism to the $(2, 1)$ torus knot, which is a Legendrian unknot and has a unique disk filling. They observed that many of these cobordisms were Hamiltonian isotopic to each other, making at most $\frac{1}{n+1}\binom{2n}{n}$ Hamiltonian isotopy classes.[2]

A similar factorization is visible in the pictures of trees. A planar tree is dual to an ideal triangulation of the disk with $n+2$ vertices. Those vertices are in natural bijection with the Reeb chords of the spirograph projection. Ordering the vertices determines a labeled, ideal triangulation of the disk with n triangles, by slicing off triangles in order:

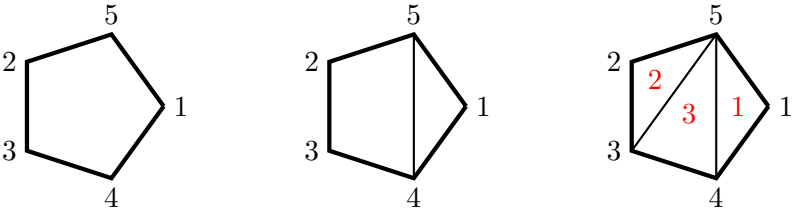

[2]Furthermore, they showed that the $\mathbf{F}_2$ augmentation provided an invariant that showed they make at least $\frac{1}{3}(2^{n+1}-1)$ Hamiltonian isotopy classes. In [STWZ] we used constructible sheaves to give an invariant distinguishing the Catalan-numbers-worth of fillings constructed there.

The labeled triangulation does not depend on the ordering of vertices n, $n+1$, $n+2$. Write a_k for the triangle labeled by k. Then $a_n, \ldots, a_1$ makes a shelling of the triangulated $(n+2)$-gon in the sense of [BM]. In particular, the union $\bigcup_{k=1}^{n} a_k$ makes a triangulated $(n-k+2)$-gon, dual to a planar trivalent tree, and the hyperelliptic wavefront associated to that tree gives a filling of a Legendrian $(2, n-k)$ torus knot or link.

Remark 2.4. The filling we construct only depends on the tree/triangulation, not on the labeling/shelling. We gave an iterative description to highlight the connection to [EHK].

3. Foams and fillings

In some sense we regard the hyperelliptic Legendrians of this paper as two-dimensional analogues of two-strand torus knots. In this section we show that the format of §2.3, where planar trivalent trees give a blueprint for describing the Lagrangian surfaces that fill such knots, can be adapted to describe three-dimensional Lagrangians that fill hyperelliptic Legendrians. The role of trivalent trees is played by foams:[3]

Definition 3.1. Let D^3 be the 3-dimensional ball. A *foam* in D^3 is a stratified subset $\mathbf{F}^2 \subset \mathbf{F}^1 \subset \mathbf{F} \subset D^3$ where

1) As subsets of D^3, $\mathbf{F}^2$ is a finite set of points, $\mathbf{F}^1 - \mathbf{F}^2$ is a finite set of smoothly embedded arcs and $\mathbf{F} - \mathbf{F}^1$ is a finite set of smoothly embedded surfaces.
2) Near each point $\mathbf{F}^2$, we may find smooth coordinates on D^3 identifying $\mathbf{F}$ with the cone over the 1-skeleton of a tetrahedron.
3) Each connected component of $\mathbf{F}^1 - \mathbf{F}^2$ has a neighborhood that looks like the cone on the vertices of a triangle, times an interval.

We furthermore assume that $\mathbf{F}$ is conical in a collar neighborhood N of ∂D^3 — it looks like the preimage of a cubic planar graph under the radial projection $N \to \partial D^3$.

[3]Many related constructions appear in the works described in Section 1.4.1 — see, in particular, Section 5.1.2 of [CCV] and Section 3.1 of [CEHRV] — though not foams *per se* or an emphasis on Legendrian boundaries.

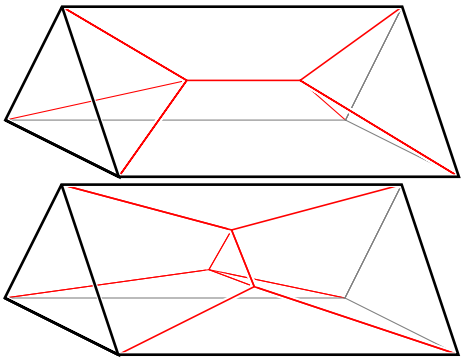

Figure 3.0.1: Two examples of a foam filling a triangular prism graph.

The conditions of the definition are topological analogues of the Plateau conditions for soap films. (The genuine Plateau conditions are metric — soap films have constant mean curvature away from the singularities, where the smooth sheets meet at equilateral angles.) A foam gives a regular cell complex structure on D^3, whose dual complex is a "tetrahedronation" of D^3 — this is similar to the duality between planar trivalent graphs and triangulations. We will say that a foam is *ideal* if every connected component of $D^3 - \mathbf{F}$ contains part of ∂D^3 — then the dual tetrahedronization will have no internal vertices; this is analogous to the cubic planar graph being a cubic planar tree.

3.1. The Harvey-Lawson cone

The building-block example of a foam, is the cone over the 1-skeleton of a regular tetrahedron centered at the origin. Before turning to our general construction let us explain the relationship between this foam and the *Harvey-Lawson cone*. This is the Lagrangian subset $HL \subset \mathbf{C}^3$ given parametrically by

$$(r, e^{is}, e^{it}) \mapsto (re^{is}, re^{it}, re^{-is-it})$$

where $r \geq 0$ and $(s,t) \in S^1 \times S^1$. It is a cone over a two-torus, with a conic singularity at the origin $r = 0$. Coordinatize $\mathbf{C}^3$ as $z = x + iy$, with $z = (z_1, z_2, z_3)$ and likewise for x and y. Then HL is special Lagrangian with respect to the standard Calabi-Yau structure. Define $n = \frac{x}{|x|} \in S^2$ and $f = n \cdot y \in \mathbf{R}$ The restriction of HL to fixed r is Legendrian, with wavefront

projection

$$(s,t) \mapsto e^f n \in \mathbf{R}^3 \setminus \{0\} \cong S^2 \times \mathbf{R}$$

shown in Figure 2.0.1. The Harvey-Lawson cone is therefore a singular Lagrangian filling of the Legendrian surface associated to the tetrahedron. (It is also exact and special.)

We will be interested here in a different projection, to the real three-space:

$$HL \to \mathbf{R}^3_x, \qquad (r,s,t) \mapsto (r\cos(s), r\cos(t), r\cos(-s-t)).$$

This is a double cover branched over the four rays

$$\begin{aligned} &\{(a,a,a) \mid a>0\}, \quad \{(a,-a,-a) \mid a>0\}, \\ &\{(-a,-a,a) \mid a>0\}, \quad \{(-a,a,-a) \mid a>0\} \end{aligned} \tag{3.1.1}$$

Note the generator for the $\mathbf{Z}/2$ Galois group of this covering is given explicitly by $(r,s,t) \mapsto (r,-s,-t)$. The unusually low dimension of the critical value locus — here it has codimension 2 — is related to the fact that both *HL* and the kernel of the projection $\mathbf{C}^3 \to \mathbf{R}^3$ are special Lagrangian of the same phase.

The rays (3.1.1) are generated by the vertices of a regular tetrahedron. The foam whose walls are the sectors through these edges play a role in the *exactness* of *HL*. Indeed being a cone, the set *HL* is contractible and therefore an exact Lagrangian with respect to any primitive α for ω. It is natural to take for α the canonical one-form obtained when identifying $\mathbf{C}^3$ with $T^*\mathbf{R}^3$. Then we compute $\alpha|_{HL} = df$ where

$$f = \frac{1}{4} r^2 (\sin(2s) + \sin(2t) - \sin(2s+2t)) \tag{3.1.2}$$

Note that $f = 0$ precisely when (x_1, x_2, x_3) are in walls of the tetrahedral foam.

Remark 3.2. It is possible to describe *HL* as the graph of a double-valued one-form, indeed as the graph of *df* where

$$f = \pm \frac{1}{2} \sum_{i=1}^{3} \epsilon_i x_i \sqrt{r^2 - x_i^2}$$

where

- $\epsilon_i = \epsilon_i(x_1, x_2, x_3) \in \{-1, 1\}$ is a sign that depends on which chamber of the foam (x_1, x_2, x_3) belongs to. (Thus, let $\sqrt{r^2 - x_i^2}$ always denote the positive square root, and let this ϵ and the $\pm$ in front of the summation sign do the work recording the sign ambiguity of f).
- r is the largest solution to

$$r^3 - (x_1^2 + x_2^2 + x_3^2)r + 2x_1x_2x_3 = 0$$

Note for all $(x_1, x_2, x_3) \in \mathbf{R}^3$ all roots of that cubic equation are real — the largest root is also the unique solution with $r \geq \max(|x_1|, |x_2|, |x_3|)$, and it has multiplicity 2 exactly along the rays (3.1.1).

3.2. Singular exact fillings

Let $\Gamma \subset S^2$ be a cubic planar graph and let $\mathbf{F} \subset D^3$ be a foam whose boundary is Γ. Let $\pi : L(\mathbf{F}) \to D^3$ denote the connected double cover branched over the 1-skeleton of $\mathbf{F}$, and write $\iota_{L(\mathbf{F})}$ for its nontrivial Deck transformation. Then $L(\mathbf{F})$ is a manifold away from the vertices of $\mathbf{F}$ — near each vertex it is diffeomorphic to the Harvey-Lawson cone. The boundary of $L(\mathbf{F})$ is naturally identified with the hyperelliptic surface S branched over the vertices of Γ.

We choose a function $z : L(\mathbf{F}) \to \mathbf{R}$ that is smooth away from the cone points of $L(\mathbf{F})$, that is odd for the Deck transformation of π (i.e. $z \circ \iota = -z$) and that has the following additional properties:

1) In a neighborhood U of each cone point, there are coordinates on $\pi^{-1}(U)$ such that $z|_U$ looks like (3.1.2)
2) The only critical points of z are on $\pi^{-1}(\mathbf{F}^1)$, i.e. on the critical points of π. Furthermore, in a collar neighborhood of ∂D^3, z is quadratic along each ray.

Then (just as in §2.3, but one dimension up) we obtain an exact Lagrangian in T^*D^3 by taking $(\pi, \nabla z)$ whose restriction to $T^*D^3|_{\partial D^3} \cong (T^*S^2) \times \mathbf{R}$ is Legendrian. For example, the function f of Remark 3.2 has all of these properties.

Remark 3.3. Note that item (1) of the conditions on z is much stronger than the analogous item (4) of Definition 2.1. In particular it is only possible when there are coordinates near the vertex of $\mathbf{F}$ that "linearize" the foam. We suspect that it is possible to assume $\mathbf{F}$ is in such a form without affecting

the Hamiltonian isotopy class of the exact Lagrangian $L(\mathbf{F})$, but it would be desirable to relax this condition for other reasons.

Remark 3.4. In the case of the Harvey-Lawson cone, this construction yields what Nadler denotes by "$\Sigma \subset A$" in the proof of [N2, Thm 3.2] — A is a copy of $\mathbf{R}^3$ and Σ is the hyperelliptic wavefront of Figure 2.0.1.

3.3. Cobordisms

An exact filling of S is an exact Lagrangian cobordism from S to the empty Legendrian. One can ask whether these cobordisms can be factored into simpler cobordisms — as we discussed in §2.3, this is how Ekholm, Honda, and Kálmán discovered their Lagrangians. In this section we indicate a (somewhat) analogous description of the singular exact fillings of §3.2, in particular by rough analogy with the discussion of §2.3.3, but the analogy highlights some differences.

Suppose that Γ has no self-loops or multiple edges, so its planar dual graph $\hat{\Gamma}$ is a triangulation of S^2. Just as a cubic planar tree is dual to an ideal triangulation of a disk, an ideal foam is dual to an ideal "tetahedronation" of the ball D^3 that restricts to $\hat{\Gamma}$ on $\partial D^3 = S^2$. Thus if $\mathbf{F}$ is a foam filling of Γ, let $\hat{\mathbf{F}}$ be the dual tetrahedronation restricting to $\hat{\Gamma}$ on the boundary. If $\mathbf{F}$ is an ideal foam, $\hat{\mathbf{F}}$ is an ideal triangulation: it has no vertices besides those on ∂D^3. For simplicity, we assume $\hat{\mathbf{F}}$ is shellable, but as in the two-dimensional case (see Remark 2.4) this assumption is unnecessary and $L(\mathbf{F})$ depends only on the foam $\mathbf{F}$.

In §2.3.3 we noted that a shelling of the triangles in the dual triangulation of a disk gives an "EHK factorization" of the exact Lagrangian surfaces into elementary cobordisms. Thus, let n be the number of tetrahedra in $\hat{\mathbf{F}}$, equivalently the number of interior vertices in $\mathbf{F}$, and let $a_n, a_{n-1}, \ldots, a_1$ be a shelling of the tetrahedra in $\hat{\mathbf{F}}$. (We use descending indices to better match the notation of §2.3.3; the highest index is the outer shell.) For each k, the union of $a_n, \ldots, a_{n-k}$ is another polyhedron with a tetrahedronation — it is dual to a foam filling $\mathbf{F}(k)$ with boundary a cubic planar graph $\Gamma(k)$. Write $L(\mathbf{F}(k))$ for the singular exact Lagrangian double cover branched over the one-skeleton of $\mathbf{F}(k)$. We will indicate how $L(\mathbf{F}(k+1))$ is obtained from $L(\mathbf{F}(k))$ and a Harvey-Lawson-type Lagrangian HL_{k+1} by gluing $L(\mathbf{F}(k))$ and HL_{k+1} along parts of their boundaries.

For HL_{k+1}, we take the branched double cover of the one-skeleton of the foam associated with a small tetrahedron with one vertex on each face of $a_{n-(k+1)}$. There are two cases to consider: the tetrahedron $a_{n-(k+1)}$ is

incident with $\bigcup_{i=1}^{k} a_{n-i}$ along either exactly one or exactly two of its faces. (Since we assume $\hat{\mathbf{F}}$ has no internal vertices, it is impossible for $a_{n-(k+1)}$ to be indicent with the previous simplices along three of its faces.)

If it is incident along exactly one face, the left-hand figure illustrates $L(\mathbf{F}(k)) \cup HL_{k+1}$ before the gluing, and the right-hand figure illustrates it after the gluing:

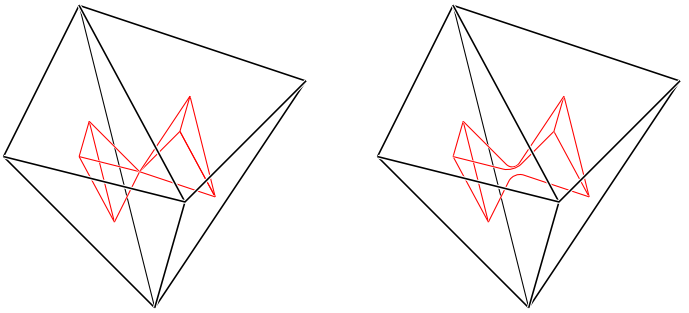

Figure 3.3.1: The black tetrahedra are a_n and a_{n-1}. On the left, one of the red tetrahedra is the boundary of $\mathbf{F}(1)$ and one of them is the boundary of the foam associated with HL_2. On the right, the red figure indicates the boundary of $\mathbf{F}(2)$ obtained by gluing HL_2 to $\mathbf{F}(1)$.

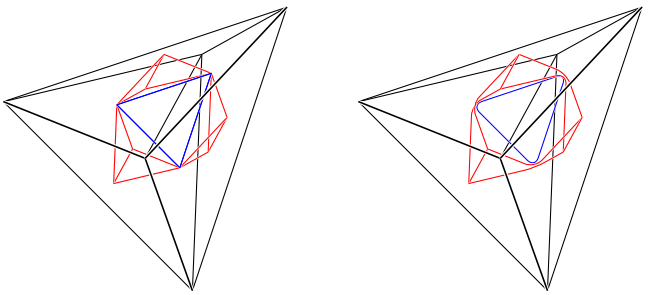

Figure 3.3.2: In case $a_{n-(k+1)}$ shares a face with two tetrahedra of $\hat{\mathbf{F}}(k)$, the red+blue graph that emerges from the cobordism is not the boundary of $\mathbf{F}(k+1)$, but of a disjoint union of $\mathbf{F}(k+1)$ and a spurious (blue) loop.

More precisely, if $a_{n-(k+1)}$ is incident with the previous simplices along exactly one face, then $L(\mathbf{F}(k+1))$ is obtained from $L(\mathbf{F}(k))$ and HL_{k+1} by modifying the disconnected wavefront near the two vertices meeting at the interior face. Specifically, remove from each wavefront a disk neighborhood

of the relevant vertex and replace these disks by gluing in a cylinder, as in Figure 3.3.1.

To understand the case of an interior face of $\mathbf{F}$, i.e. edge of $\hat{\mathbf{F}}$, consider Figure 3.3.2 above.

To define $\mathbf{F}(k+1)$ and obtain from it a singular exact filliing, we construct a cobordism from the boundary of $L(\mathbf{F}(k)) \sqcup HL_{k+1}$ to a new space by modifying the wavefront to remove the spurious blue loop. Consider a neighborhood of the blue loop in the front projection. It is a cylinder. The cobordism takes this cylinder to two disks, just like the two-to-one-sheet hyperboloid cobordism $x^2 + y^2 - z^2 = t,\ t \in [0,1]$. Afterward, the blue loop disappears and the red edges make up a new polytope, or dually define $\mathbf{F}(k+1)$, which in Figure 3.3.2 is a triangular prism.

3.4. The tangles associated to a foam

The one-skeleton $\mathbf{F}^1$ of a foam $\mathbf{F} \subset D^3$ is a kind of singular tangle in the ball, joining the vertices of graph on the boundary. The singularities are at the internal vertices $\mathbf{F}^0$ of the foam, suppose that there are T of these. Then we can associated as many as 3^T smoothings of this singular tangle, as follows. In a neighborhood of each of the T vertices, choose coordinates in which $\mathbf{F}^1$ looks like the rays of (3.1.1), and replace $\mathbf{F}^1$ by one of the smooth hyperbolas

$$\{(x_1, x_2, x_3) \mid x_{i_2} = x_{i_3} \text{ and } x_{i_1}^2 - x_{i_2}x_{i_3} = 1\} \tag{3.4.1}$$

where i_1, i_2, i_3 is a permutation of $\{1,2,3\}$ — the hyperbola is determined just by the value of i_1. In fact, this is precisely the critical locus of the projection to $\mathbf{R}^3_x$ of the Harvey-Lawson smoothing, the solid torus of Example 3.7 below.

We have not investigated what kinds of tangles can appear, the examples we have encountered so far are all quite unlinked.

Remark 3.5. In [CCV] and especially [CEHRV], singular and smooth tangles were considered in a nearly identical context: constructing Lagrangian branes. However, the authors there did not consider fillings of a fixed Legendrian boundary. Our purposes require the consideration of foams. In the work of [CEHRV], all smoothings of tangles were considered. We do not know if such tangles occur within the context of foams.

Example 3.6. Let $\mathbf{F}$ be the foam on the left and $\mathbf{G}$ the foam on the right in Figure 3.0.1. In general, smoothings are naturally indexed by the data of,

at each internal vertex v, a partition of the four edges incident with v into two pairs. Taken up to the dihedral symmetry of the foams, there are three possibilities for $\mathbf{F}$ and six for $\mathbf{G}$. (For $\mathbf{F}$, one of these smoothings is pictured in blue in §4.9.2.) Eight of these nine tangles are abstractly homeomorphic to each other, and in fact to the "trivial tangle" of three parallel strands. The ninth tangle, which appears for $\mathbf{G}$: it is the union of three parallel strands, along with an unknot which is not linked with any strand.

3.5. Nonexact fillings — tangles as caustics

We have constructed singular exact fillings of S with the topology of double covers of the ball, branched over $\mathbf{F}^1$, the edges of a foam. Example 3.7 below suggests that we search for smooth fillings with the topology of a double cover of the ball, branched over a tangle obtained by smoothing $\mathbf{F}^1$ in the sense of §3.4. However the example also suggests that such smoothings will not be exact. Theorem 4.4 gives a strong result making this precise. Let us therefore make some general remarks about nonexact fillings.

Let $(M, \omega = d\alpha)$ be an exact symplectic manifold. In the literature on Lagrangian fillings of Legendrians $\Lambda \subset \partial M$, the condition that the Lagrangian has conical ends is often imposed. For 3-dimensional Lagrangians, this has an undesirable consequence. If L is 3-dimensional and U is a collar neighborhood of the boundary of L, the map $H^1(L, \mathbf{R}) \to H^1(U, \mathbf{R})$ is often injective, so if α is exact on U, it is exact on all of L. This suggests that if we wish to consider non-exact Lagrangians, the condition that L has conical ends is inappropriate. A notion of "asymptotically conical" might be more appropriate, as in [J, Def. 7.1] for special Lagrangians, or [NZ, Def. 5.4.1] for subanalytic Lagrangians.

Example 3.7 (The Harvey-Lawson solid tori). There are three one-parameter families of these, given parametrically by

$$(r, e^{is}, e^{it}) \mapsto (\sqrt{r^2 + \epsilon^2} e^{is}, re^{it}, re^{-is-it})$$

and its permutations. The antiderivative of $\alpha|_{HL_\epsilon}$ is

$$f \text{ of } (3.1.2) + \frac{1}{2}\epsilon^2 \sin(s)\cos(t) + \frac{1}{2}\epsilon^2 s$$

The last term, $\frac{1}{2}\epsilon^2 s$, is not periodic and therefore HL_ϵ is not exact — however it also has no r-dependence. It is an asymptotically conic special Lagrangian, in the sense studied by Joyce (see [J, Example 6.9]). For general tangles,

the existence of asymptotically special Lagrangian smooth embeddings in $\mathbf{C}^3$ is unproven. However, we have the following construction of non-special Lagrangian fillings in $\mathbf{C}^3$, potentially with isolated immersed double points.

If $\mathcal{T}$ is a tangle associated to a foam $\mathbf{F}$, write $\pi : L(\mathcal{T}) \to D^3$ for the double cover branched over $\mathcal{T}$. In the rest of the paper we will refer to any Lagrangian map $L(\mathcal{T}) \to T^*D^3$ whose projection to D^3 is π as a "filling" of S — or a "filling modeled on $\mathcal{T}$" — without heed for asymptotic conditions. Note that even so, the boundary of $L(\mathcal{T})$ is canonically identified with S, since $\mathcal{T} \cap \partial D^3$ is exactly the set of vertices of Γ. We also will not necessarily assume that the map $L(\mathcal{T}) \to T^*D^3$ is an embedding.

Once a filling modeled on $\mathcal{T}$ is given, the pull-back of the canonical one-form on T^*D^3 to $L(\mathcal{T})$ is closed. Conversely, for each closed one-form there is at most one Lagrangian map $L(\mathcal{T}) \to T^*D^3$ whose projection to D^3 is π.

To see that a filling modeled on $\mathcal{T}$ is determined by a 1-form on $L(\mathcal{T})$, write x_1, x_2, x_3 for the coordinates on D^3, y_1, y_2, y_3 for the coordinates on T^*D^3, and let s_1, s_2, s_3 be local coordinates on $L(\mathcal{T})$. If $\alpha = \alpha_1 ds_1 + \alpha_2 ds_2 + \alpha_3 ds_3$ is a closed one-form on $L(\mathcal{T})$, then y_1, y_2, y_3 are the solutions to $\sum y_i dx_i = \alpha$, which is equivalent to

$$\begin{pmatrix} \partial x_1/\partial s_1 & \partial x_2/\partial s_1 & \partial x_3/\partial s_1 \\ \partial x_1/\partial s_2 & \partial x_2/\partial s_2 & \partial x_3/\partial s_2 \\ \partial x_1/\partial s_3 & \partial x_2/\partial s_3 & \partial x_3/\partial s_3 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{pmatrix} \tag{3.5.1}$$

The matrix on the left-hand side is the Jacobian of π, which is invertible away from $\mathcal{T}$. For (3.5.1) to have a solution over $\mathcal{T}$ requires that α vanish along $\mathcal{T}$. The existence of such α is addressed by the following:

Proposition 3.8. *Let $\pi : L(\mathcal{T}) \to D^3$ be a branched double cover of a tangle $\mathcal{T} \subset D^3$, and let ι denote the Deck involution. Suppose that $\mathcal{T}$ has no circle components (just strands), and let θ be a closed one-form on $L(\mathcal{T})$. Then there is a function $F : L(\mathcal{T}) \to R$ such that $\alpha := \theta + dF$ is* odd *under ι, i.e. $\iota^*\alpha = -\alpha$. In particular, α vanishes on $\mathcal{T}$. Moreover, any other α' satisfying these properties differs from α by dG, where $G = -G \circ \iota$, i.e. G is odd under ι.*

Before proving the proposition, we remark: it follows that for each class in $H^1(L(\mathcal{T}), \mathbf{R})$, there is a Lagrangian immersion into $\mathbf{R}^6 \cong T^*D^3$ with at worst isolated double points, such that the pull-back of the canonical form of T^*D^3 is in the same cohomology class. Indeed we can select α as in

Proposition 3.8 and construct the immersion as described by (3.5.1). Zeroes of α give double points of the immersion.

Proof. Let us write L for $L(\mathcal{T})$. If $Z^1(L,\mathcal{T})$ is the space of closed 1-forms that vanish on $\mathcal{T}$, and $B^1(L,\mathcal{T})$ is the subspace of 1-forms like dG, where G vanishes on $\mathcal{T}$, then we have a short exact sequence of ι-modules

$$0 \to B^1(L,\mathcal{T}) \to Z^1(L,\mathcal{T}) \to H^1(L,\mathcal{T}) \to 0$$

We will prove the Proposition by proving that the tautological map $H^1(L,\mathcal{T}) \to H^1(L)$ is an isomorphism when restricted to the (-1)-eigenspace of $H^1(L,\mathcal{T})$.

We have assumed $\mathcal{T}$ has no circle components, so $H^1(\mathcal{T}) = 0$. The long exact sequence of the pair $(L,\mathcal{T})$ therefore induces a short exact sequence

$$H^0(L) \to H^0(T) \to H^1(L,\mathcal{T}) \to H^1(L) \to 0 \tag{3.5.2}$$

(3.5.2) is a short exact sequence of ι-modules. Both $H^0(\mathcal{T})/H^0(L)$ and $H^1(L)$, have dimension g over $\mathbf{R}$ — ι acts trivially on $H^0(\mathcal{T})/H^0(L)$ and (since $L/\iota = D^3$ and $H^1(D^3) = 0$) by -1 on $H^1(L)$. Thus the (-1)-eigenspace of the ι-action on $H^1(L,\mathcal{T})$ is identified with $H^1(L)$. □

Remark 3.9. It is tempting to speculate that by choosing G in the Proposition correctly, one can make the map $L(\mathcal{T}) \to T^*D^3 \cong \mathbf{R}^6$ a special Lagrangian immersion, or even embedding. This amounts to showing the existence of a solution to a PDE of Monge-Ampère type for the single function G.

4. Constructible sheaves

To this point, we have defined a Legendrian surface S from a cubic, planar graph Γ, and singular, exact Lagrangian fillings from foams, as well as their smoothings via tangles. In this section, we study a category of sheaves defined by S, which by [N, NZ] is a Fukaya category. We give a concrete description of this category and its moduli space $\mathcal{M}$ of simple objects, then use this moduli space to give many examples of non-isotopic Legendrians with the same genus and classical invariants, and to prove non-fillability for Legendrians constructed from simple cubic planar graphs. We show that $\mathcal{M}$ embeds into a period domain as a Lagrangian submanifold, and define a generalized notion of "phase" and "framing" (related to the tangle description of the filling), to define conjectural open Gromov-Witten invariants à la [AV, AKV].

Fix a commutative ground ring k and write $Sh(\mathbf{R}^3)$ for the k-linear dg-derived category of constructible sheaves of k-modules on $\mathbf{R}^3$. In this section, we study full subcategories of $Sh(\mathbf{R}^3)$ with singular support determined by a hyperelliptic wavefront (Definition 2.1) — i.e. with singular support in either the Legendrian lift or the extended Legendrian lift (see Section 2.1) of a hyperelliptic wavefront. Thus fix a graph $\Gamma \subset S^2$, let $\Psi \subset \mathbf{R}^3$ be the image of the associated hyperelliptic wavefront, and let $S \subset S^+ \subset T^\infty\mathbf{R}^3$ be the associated Legendrians. We define

$$\mathcal{C}(S) \subset \mathcal{C}(S^+) \subset Sh(\mathbf{R}^3)$$

to be the full subcategories of sheaves that are compactly supported, and that have singular support in S or S^+. That is, $\mathcal{C}(S) := Sh(\mathbf{R}^3, S)$ and $\mathcal{C}(S^+) := Sh(\mathbf{R}^3, S^+)$.

4.1. Regular cell decomposition

We let g be as in (2.0.1), and write Ψ for the image of the hyperelliptic wavefront. The filtration

$$(\text{vertices of } \Gamma) \subset \Gamma \subset \Psi \subset \mathbf{R}^3 \tag{4.1.1}$$

gives a Whitney stratification of $\mathbf{R}^3$, with $2g+2$ strata of dimension zero, $3g+3$ strata of dimension one, $2g+6$ strata of dimension two and $g+5$ top-dimensional strata. The edges and vertices of Γ are also edges and vertices of (4.1.1), but let us give some vocabulary for the two- and three-dimensional strata.

For the top-dimensional strata, there is unique region containing the origin and a unique region incident with the point at ∞, that we call the "inner" and "outer" region respectively. The remaining open strata are in one-to-one correspondence with the $g+3$ faces of Γ, we call these regions "pillows." Each two-dimensional stratum is in the boundary of a unique pillow — we call them "sheets." Each pillow has exactly two sheets at its boundary, the "inner sheet" which is incident with the inner region and the "outer sheet" incident with the outer region.

If we omit the outer region (which has the topology of $S^2 \times \mathbf{R}$), the stratification is a regular cell complex. We therefore have an equivalence (see e.g. [KS, §8.1] [STZ2, Prop. 3.9]) between

- the derived category of sheaves that are constructible for the stratification and acyclic in the outer region

- the derived category of functors from the partially ordered set of strata (not including the outer region) to k-modules.

We will use this equivalence freely in what follows, describing a sheaf by a strictly commutative diagram of chain complexes in the shape of the poset of strata. It is convenient to include the outer region in these diagrams, but it is always to be labeled by the zero complex.

We will describe conditions on such diagrams to belong to $\mathcal{C}(S^+)$ and $\mathcal{C}(S)$. These conditions are local, so the cases to be considered are a neighborhood of a vertex, and edge, and an inner or outer sheet. In the end we find that an object of $\mathcal{C}(S^+)$ or of $\mathcal{C}(S)$ has a very simple description, which we summarize in §4.2.

4.1.1. Local study at one- and two-dimensional strata. In codimension two or less, a wavefront hypersurface (of a manifold of any dimension n) looks like a product of a smooth hypersurface in $\mathbf{R}^{n-2}$ with the front diagram of a Legendrian knot in $\mathbf{R}^3$ — we have studied these in detail in [STZ2]. We recall the local descriptions for a front without cusps (thus without any genuine front singularities) here — it is a special case of [STZ2, Thm. 3.12].

Suppose Σ is a sheet incident with the two regions R_1 and R_2, with R_1 farther from the origin than R_2. Then a sheaf F is given near Σ by a diagram

$$F(R_1) \leftarrow F(\Sigma) \rightarrow F(R_2)$$

If Σ is an outer sheet, then R_1 is the outer region and we require $F(R_1)$ to be the zero complex. If the singular support is to lie in S^+ (or S, away from the edges there is no difference), the map $F(\Sigma) \rightarrow F(R_1)$ must be an isomorphism.

If i is an edge of Γ, a sheaf F is given near i by a diagram

(4.1.2)

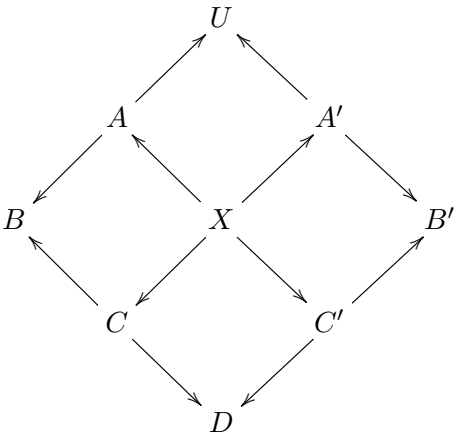

where U is the complex labeling the outer region (it must be zero), B and B' label the pillows, D labels the inner region, A and A' label upper sheets, C and C' label lower sheets, and X labels the edge. The sheaf belongs to $\mathcal{C}(S^+)$ if and only if the six maps $X \to A$, $X \to A'$, $A \to U$, $A' \to U$, $C \to B$, $C' \to B'$ are quasi-isomorphisms. Note in particular that this requires that X, A, and A' are acyclic. It belongs to $\mathcal{C}(S)$ if and only if the square at the bottom is exact (i.e. it realizes X as the homotopy fiber product of the maps $C, C' \to D$).

4.1.2. Local study at a vertex. The natural stratification of the wavefront given by (2.0.2) has one vertex, three edges, six 2-dimensional strata (wedge-shaped sheets), and five 3-dimensional strata (three pillows in between the sheets, an outer region and an inner region). Their pattern of incidences is recorded in (4.1.3) — the vertex is denoted Z, the edges are denoted X, the sheets are the A's and C's, the three pillows are Bs, and the big open regions are U and D.

The sheaf corresponding to (4.1.3) belongs to $\mathcal{C}(S^+)$ if and only if the following conditions hold (where X means any X_i, etc.):

(4.1.3)

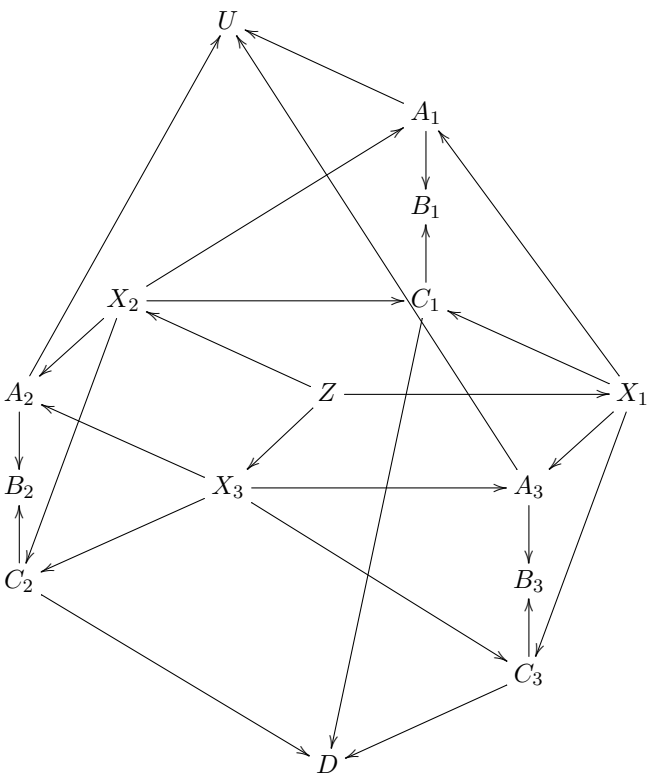

1) $U = 0$,

2) all the maps from $A \to U$, $X \to A$, and $C \to B$ are quasi-isomorphisms,

3) The total complex of

(4.1.4)
$$Z \to \big(X_1 \oplus X_2 \oplus X_3\big) \to \begin{pmatrix} A_1 \oplus A_2 \oplus A_3 \oplus \\ C_1 \oplus C_2 \oplus C_3 \end{pmatrix} \to \begin{pmatrix} B_1 \oplus B_2 \oplus B_3 \oplus \\ D \end{pmatrix}$$

is acyclic.

It belongs to $\mathcal{C}(S)$ if and only if it obeys the further condition

(4) Each commutative square between X and D is exact, i.e. the total complex of
$$X_i \to C_i \oplus C_{i-1} \to D$$
is acyclic.

Indeed the three commutative 3×3 squares in (4.1.3)

$$\begin{array}{ccc} U & A_1 & B_1 \\ A_3 & X_1 & C_1 \\ B_3 & C_3 & D \end{array} \qquad \begin{array}{ccc} U & A_2 & B_2 \\ A_1 & X_2 & C_2 \\ B_1 & C_1 & D \end{array} \qquad \begin{array}{ccc} U & A_3 & B_3 \\ A_2 & X_3 & C_3 \\ B_2 & C_2 & D \end{array}$$

are each of the form considered in (4.1.2), which establishes (1) and (2). The equation (4.1.4) gives the microstalk of the sheaf at the singular point in the vertical codirection.

4.2. Concrete description of $\mathcal{C}(S)$ and $\mathcal{C}(S^+)$

Consider the star-shaped quiver with a single sink (let us call it o) and $g+3$ sources indexed by the faces a of Γ. The derived category of this quiver is equivalent to the category of constructible sheaves on the union of the inner region and the inner sheets (labeled by Ds and Cs in the local descriptions §4.1.1–4.1.2). It follows from the discussion in §4.1 that the restriction functor from $\mathcal{C}(S^+)$ to this union is an equivalence.

In this quiver description, $\mathcal{C}(S) \subset \mathcal{C}(S^+)$ is a full subcategory, with F belonging to $\mathcal{C}(S)$ if and only if whenever a and b are faces of Γ separated

by an edge,

(4.2.1) the map $F(a) \oplus F(b) \to F(o)$ is an isomorphism.

(In the analogous story for Legendrian curves, this is the "crossing condition" of [STZ2].)

Remark 4.1. Note in particular that the quasi-equivalence class of the dg category $\mathcal{C}(S^+)$, since it only depends on the genus, g, does not depend on the graph Γ. A more precise form of this statement can be obtained from the picture of §2.2: if Γ' is obtained from Γ by a sequence of edge moves, then $(S')^+$ will be obtained from $(S)^+$ by a sequence of disk moves, which induce equivalences of categories as in [STW, §2.3]. In particular, the moduli spaces of objects in $\mathcal{C}(S)$ and $\mathcal{C}(S')$ will be related to each other by a sequence of cluster transformations.

4.3. Moduli of microlocal rank one objects

Note (4.2.1) has the consequence that, if $F \in \mathcal{C}(S)$ and $F(o)$ is a vector space concentrated in degree 0, then $F(o)$ must be even-dimensional, with the spaces labeling the pillows being concentrated in degree 0 and having half the dimension. We call that dimension the *microlocal rank* of the object.

Thus if F has microlocal rank one, $F(o)$ is a 2-dimensional vector space. Let us say that a framed object is one equipped with an isomorphism $F(o) \cong k^2$. A framing rigidifies an object: there are no automorphisms of an object (nor even self-homotopies of the identity automorphism) that preserve the framing except for the identity — there is a fine moduli space of framed sheaves of microlocal rank one. It can be described concretely as an open subset of $(\mathbf{P}^1)^{(g+3)}$, where the factors of $\mathbf{P}^1$ are indexed by the faces of Γ.

We define an affine open subset $\mathcal{M}_{fr}(\Gamma) \subset (\mathbf{P}^1)^{g+3}$ as follows. A point in $\mathcal{M}_{fr}(\Gamma)$ is a collection of $z_a \in \mathbf{P}^1$, one for each face a of Γ, subject to the condition that $z_a \neq z_b$ whenever a and b share an edge of Γ.

PGL_2 acts diagonally on $\mathcal{M}_{fr}(\Gamma)$. Define $\mathcal{M}(\Gamma)$ as the quotient $\mathcal{M}(\Gamma) := \mathcal{M}_{fr}(\Gamma)/\mathrm{PGL}_2$.

4.4. The chromatic polynomial as a Legendrian invariant

The definition of $\mathcal{M}_{fr}(\Gamma)$ makes sense over any commutative ring, so we consider it over the finite fields $\mathbf{F}_q$ with q a prime power. The number of $\mathbf{F}_q$-points of $\mathcal{M}_{fr}(\Gamma)$ is equal to the number of $(q+1)$-colorings of the map

defined by Γ, or equivalently the number of $(q+1)$-colorings of the dual planar graph $\hat{\Gamma}$ — these are the values at $q+1$ of the *chromatic polynomial* of $\hat{\Gamma}$.

Theorem 4.2. *Suppose Γ and Γ' are cubic planar graphs with the same number of vertices, so the respective surfaces S and S' have the same genus and the same classical invariants (cf. Prop. 2.3). If $\hat{\Gamma}$ and $\hat{\Gamma}'$ do not have the same chromatic polynomial, then S and S' are not Legendrian isotopic.*

Proof. We prove the contrapositive. Whenever S is Legendrian isotopic to S', the GKS-equivalence [GKS] gives an isomorphism between $\mathcal{M}(\Gamma)$ and $\mathcal{M}(\Gamma')$. In particular, these moduli spaces have the same number of points over different fields. By considering $\mathbf{F}_q$, this means the chromatic polynomials are equal at all prime powers, and thus equal. □

Remark 4.3. 1) It is natural to ask whether a Legendrian isotopy between S and S' implies that Γ and Γ' are equivalent as planar graphs. Counterexamples have recently been obtained by Roger Casals and Emmy Murphy. In particular Casals has constructed an infinite family of examples, the simplest of which are the following:

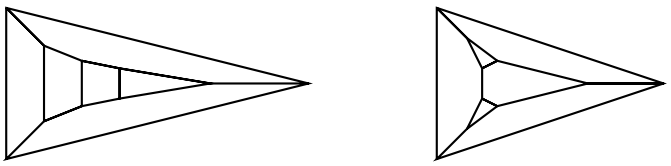

In general Casals's examples are obtained from the blow-up process of §5.2.

2) Dimitroglou Rizell [D] has constructed for each integer g a family of $g+1$ Legendrian embeddings of a genus-g surface into S^5, each with $g+1$ Reeb chords, no two of which are Legendrian isotopic. The hyperelliptic Legendrians have $g+3$ Reeb chords apiece. We do not know whether there is a literature on the number of different chromatic polynomials of planar graphs, but we suspect that Theorem 4.2 shows that the number of pairwise distinct Legendrian surfaces in our family grows at least exponentially in g. (Since this paper first appeared as a preprint, Agol and Krushkal produced a proof in [AK, Thm. 1.3]).

4.5. Exact fillings of hyperelliptic Legendrians

We have discussed a family of singular exact fillings and nonexact fillings in Section 3. Here we can prove the following:

Theorem 4.4. *Let $S \subset T^\infty \mathbf{R}^3$ be the genus-g Legendrian surface defined by a simple, cubic planar graph Γ. Then S has no smooth oriented graded exact Lagrangian fillings in $\mathbf{R}^6$.*

We have given the proof already in the introduction. Note that it is not strictly necessary for Γ to be simple — the proof works so long as the graph dual to Γ has no multiple edges, or even if the number of such edges counted without multiplicity is at least $2g + 4$. We believe the "graded" condition can be removed, by using ungraded Floer groups in the construction of [NZ] — the coefficients of such ungraded groups must have characteristic 2, but we can still appeal to the same properties of the chromatic polynomial by counting points over the fields $\mathbf{F}_q$ for q a power of 2.

Theorem 4.4 shows there are no smooth exact Lagrangians which define sheaf objects of $\mathcal{C}(\Gamma)$. Nevertheless, the following proposition shows that the objects still behave cohomologically as though they were genus-g handlebodies.

Proposition 4.5. *Suppose $F \in \mathcal{C}(\Gamma) \subset Sh(\mathbf{R}^3)$ has microlocal rank one. Let us work over a field. Then*

$$\dim \mathrm{Ext}^i(F, F) = \begin{cases} 1 & \text{if } i = 0 \\ g & \text{if } i = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Proof. By §4.2, $\mathcal{C}(\Gamma)$ is a full subcategory of the derived category of a simple quiver, with one sink and $g + 3$ sources. The objects of microlocal rank one are concentrated in homological degree zero, which implies that the Ext groups vanish for $i \notin \{0, 1\}$. As representations of the quiver, the objects of microlocal rank one have dimension 1 along every source and dimension 2 along every sink; furthermore each map from a source to a sink is an inclusion, and at least three of these maps have distinct images — this implies that any endomorphism $F \to F$ must be a scalar.

The Ext^1 calculation follows from the identification of $\mathrm{Ext}^1(F, F)$ with the tangent space to F in $\mathcal{M}$. More concretely, a class in $\mathrm{Ext}^1(F, F)$ can be

represented by a short exact sequence

$$0 \to F \to E \to F \to 0$$

If the stalk of F in the middle region is a two-dimensional vector space V and the stalks in the other regions are lines $L_i \subset V$, then we may assume that the stalk of E in the middle region in $V \oplus V$, the stalk in the other regions is $L_i \oplus L_i$, and that the inclusion maps $L_i \oplus L_i \to V \oplus V$ each have the form

$$\begin{pmatrix} \mathrm{inc}_{L_i \to V} & \phi_i \\ 0 & \mathrm{inc}_{L_i \to V} \end{pmatrix}$$

where the ϕ_i are arbitrary linear maps $L_i \to V$. The data $\{\phi_i\}$ and $\{\phi_i'\}$ represent equivalent extensions if there is a commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & F & \longrightarrow & E & \longrightarrow & F & \longrightarrow & 0 \\ \downarrow & & \downarrow{=} & & \downarrow & & \downarrow{=} & & \downarrow \\ 0 & \longrightarrow & F & \longrightarrow & E' & \longrightarrow & F & \longrightarrow & 0 \end{array}$$

or equivalently, if there is a map $\psi_V : V \to V$ intertwining ϕ_i and ϕ_i'. In other words, $\mathrm{Ext}^1(F, F)$ is the g-dimensional cokernel of the map $\mathrm{End}(V) \to k^{g+3}$. □

4.6. Period domain

Let $\mathcal{P} := H^1(S; \mathbb{G}_\mathrm{m})$ be the "period domain" of the surface S — the name is explained in the next section. A basis for $H^1(S, \mathbf{Z})$ gives an identification $\mathcal{P} \cong (\mathbb{G}_\mathrm{m})^{2g}$. Here we give an alternative description of $\mathcal{P}$ as a subtorus of $(\mathbb{G}_\mathrm{m})^e = (\mathbb{G}_\mathrm{m})^{(3g+3)}$.

Each edge i of Γ determines a loop on S, in the following way. Let $U_i \subset S^2$ be a neighborhood of i:

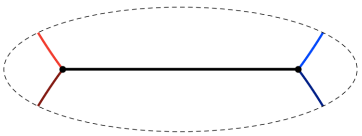

The preimage of U_i in S is an annulus:

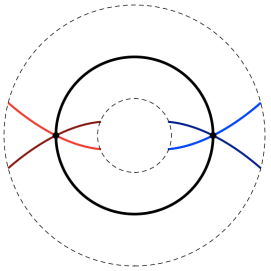

The preimage of a loop around the edge is disconnected, but the two oriented components represent the same homology class in $H_1(S)$, which we denote by γ_i. The γ_i generate $H_1(S)$, with the relations

$$\sum_{i \in \overline{f}} \gamma_i = 0 \tag{4.6.1}$$

where $\overline{f}$ a closed face, i.e. the closure of a connected component of $S^2 \setminus \Gamma$. The intersection form $H_1(S) \otimes H_1(S) \to H_0(S) = \mathbf{Z}$ can be computed from the ribbon structure (i.e. the cyclic ordering on the half-edges incident with a given vertex) on Γ, as follows:

$$\langle e_1, e_2 \rangle = \begin{cases} 0 & \text{if } e_1 \text{ and } e_2 \text{ are not incident,} \\ 1 & \text{if } e_1 \text{ and } e_2 \text{ are incident like} \\ -1 & \text{if } e_1 \text{ and } e_2 \text{ are incident like} \end{cases} \tag{4.6.2}$$

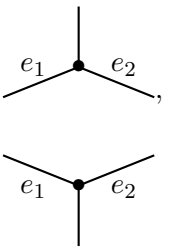

We regard $\mathcal{P}$ as an algebraic torus, whose character lattice is $H_1(S, \mathbf{Z})$. The intersection form on the character lattice is an element of

$$\mathrm{Hom}\left(\bigwedge^2 H_1(S, \mathbf{Z}), \mathbf{Z}\right),$$

but being nondegenerate it induces an element in $\bigwedge^2_{\mathbf{Z}}(H_1(S, \mathbf{Z}))$, which in turn induces a translation-invariant algebraic symplectic form on $\mathcal{P}$.

4.7. Period map

We parametrize $\mathcal{M}_{fr}$ in Definition 4.3 as an open subset of $(\mathbf{P}^1)^{(g+3)}$ and $\mathcal{M}$ as the quotient $\mathcal{M}_{fr}/\mathrm{PGL}_2$. In §4.6, we defined an algebraic torus $\mathcal{P}$, with a distinguished character $x_e : \mathcal{P} \to \mathbb{G}_{\mathrm{m}}$ for each edge of Γ. We now define a map $\mathcal{M} \to \mathcal{P}$, by the formula

$$x_i(z \in \mathcal{M}_{fr}) = -\frac{z_b - z_a}{z_c - z_b} \cdot \frac{z_d - z_c}{z_a - z_d} \tag{4.7.1}$$

when a, b, c, d are the faces surrounding an edge i in the following pattern:

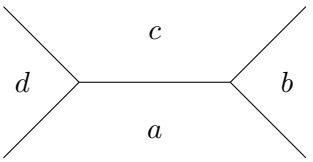

One easily verifies the relations $\prod_{e \in \overline{f}} x_e = 1$. Shen has pointed out to us some additional relations which cut $\mathcal{M}$ out of $\mathcal{P}$ as a complete intersection. Meanwhile there are several ways to see the following:

Proposition 4.6. *$\varphi : \mathcal{M} \to \mathcal{P}$ is a Lagrangian embedding.*

Proof. As mentioned in Section 1.4, an identical moduli space and Lagrangian embedding was considered in the work of Dimofte-Gabella-Goncharov — see, e.g., Section 2.3 and Theorem 4.2 of [DGG1]. □

4.8. Fillings and framings of the period domain

Let $\mathcal{T}$ be a tangle obtained by smoothing the 1-skeleton of a foam, as in §3.4, and let $L := L(\mathcal{T})$ be the Lagrangian 3-manifold associated to $\mathcal{T}$ in §3.5. In this section we will assume that $\mathcal{T}$ has no circle components. Then the identification of S with the boundary of L induces a projection $\pi : H_1(S, \mathbf{Z}) \to H_1(L, \mathbf{Z})$ that we will call the *phase* associated to $\mathcal{T}$. It is a surjection, and the kernel is identified by Poincaré duality with $H^1(L, \mathbf{Z})$. The phase therefore determines a short exact sequence

$$0 \longrightarrow H^1(L, \mathbf{Z}) \longrightarrow H_1(S, \mathbf{Z}) \xrightarrow{\pi} H_1(L, \mathbf{Z}) \longrightarrow 0 \tag{4.8.1}$$

The kernel is isotropic with respect to the intersection form on S. We define an *OGW framing*, $\mathfrak{f} : H_1(L, \mathbf{Z}) \to H_1(S, \mathbf{Z})$, to be a splitting of this short

exact sequence with the same property, i.e. a map
(4.8.2)
$$\mathfrak{f} : H_1(L, \mathbf{Z}) \to H_1(S, \mathbf{Z}) \quad \text{with } (\pi \circ \mathfrak{f})(x) = x \text{ and whose image is isotropic}$$

Such an $\mathfrak{f}$ gives a decomposition $H_1(S, \mathbf{Z}) \cong H_1(L, \mathbf{Z}) \oplus H^1(L, \mathbf{Z})$ into dual isotropic subspaces. We define $T_L := H^1(L, \mathbb{G}_{\mathrm{m}})$; it is a Lagrangian subspace of the period domain $\mathcal{P}$ of §4.6. By applying $\mathrm{Hom}(-, \mathbb{G}_{\mathrm{m}})$ to the splitting (4.8.2), we get a symplectic covering map (essentially equivalent to the data of a framing $\mathfrak{f}$, so we reuse the notation)

$$\mathfrak{f} : T^*(T_L) \to \mathcal{P}.$$

When we pull $\mathcal{M}$ back along $\mathfrak{f}$, it looks like the graph of a closed one-form whose antiderivative is a multiple-valued holomorphic function W on T_L:

$$\mathcal{M} = \mathrm{Graph}(dW), \qquad W : T_L \to \mathbf{C}.$$

In the examples we have checked, it is an integral linear combination of dilogarithms in natural coordinates defined by the phase and framing.

> *We conjecture that W is the generating function for the genus-zero open Gromov-Witten invariants of L in $\mathbf{R}^6$*

by analogy with the formulas of [AV, AKV] for AV branes, proven by Katz and Liu [KL].

Remark 4.7. The open Gromov-Witten invariants have not yet been defined in this generality, and the role that our framings should play in the theory is still somewhat mysterious — but we understand from Jake Solomon that our treatment of framings is in line with expectations. We further expect that, perhaps after choosing asymptotically radial tangles, the associated Lagrangians will satisfy the anticipated requisite bounded-geometry requirements (see [GS]) to ensure that the moduli spaces of disks are compact. Solomon-Tukachinsky [ST2, Section 1.2.6] have a project to study open Gromov-Witten theory in this setting.

4.9. Examples of OGW framings

The Lagrangians $L(\mathcal{T})$ of §3.5 are determined topologically by a tangle in the 3-ball whose endpoints are on the vertices of Γ. In the diagrams below, we sketch in blue arcs the radial projection of the tangle onto S^2. The brown

line segments indicate generators of $H_1(L, \mathbf{Z})$ up to sign — more precisely, the inverse image of each brown line under the double cover $L \to D^3$ is a circle, whose orientation we do not specify.

4.9.1. Tetrahedron. If Γ is the tetrahedron graph, $H_1(S, \mathbf{Z})$ is generated by e_1, e_2, e_3 subject to the relation $e_1 + e_2 + e_3 = 0$. Each e_i labels a pair of opposite edges, equivalent under the relations, as in the following diagram.

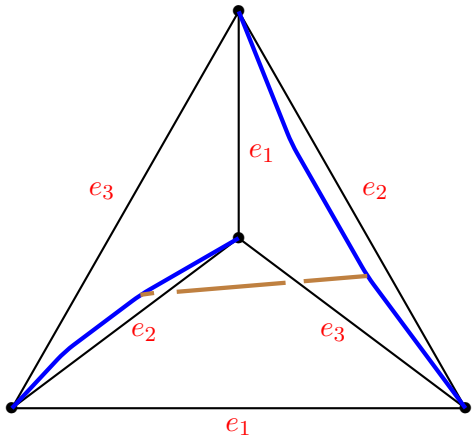

The intersection form (4.6.2) is given by $\omega(e_1, e_2) = \omega(e_2, e_3) = \omega(e_3, e_1) = 1$. The blue lines indicate a tangle in the interior of the tetrahedron, determining a phase. In the phase pictured, the loop over e_2 (as an element of $H_1(S, \mathbf{Z})$) maps to zero in $H_1(L, \mathbf{Z})$. More generally, each of the three generators e_i determines a phase, and $H_1(L, \mathbf{Z})$ is the quotient of $H_1(S, \mathbf{Z})$ by e_i.

Here is a visualization of the map of lattices $H_1(S, \mathbf{Z}) \to H_1(L, \mathbf{Z})$:

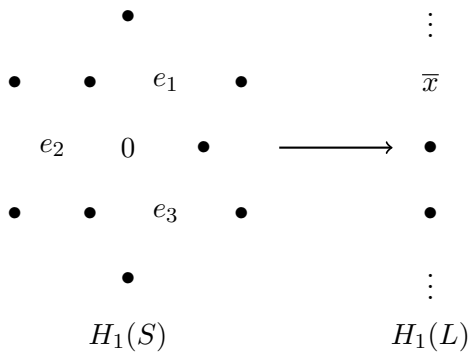

The kernel of this map is naturally identified with $H^1(L, \mathbf{Z})$ — in particular $H^1(L, \mathbf{Z})$ has a canonical generator, the edge e_2 in the picture. This gives us a preferred orientation for the loop in L that projects to the brown line segment, i.e. a preferred generator for $H_1(L, \mathbf{Z}) = \mathrm{Hom}(H^1(L, \mathbf{Z}), \mathbf{Z})$. In other words we choose the generator $\overline{x}$ for which

$$\omega(x, e_2) > 0 \tag{4.9.1}$$

holds for any representative $x \in H_1(S, \mathbf{Z})$ projecting to $\overline{x}$. The combinatorics of the situation gives us a distinguished choice of x, namely $x = e_1$. (This is not typical for more general graphs.)

When we think of $H^1(L, \mathbf{Z})$ as functions (homomorphisms) on $H_1(L, \mathbb{G}_{\mathrm{m}})$, we will write e^v instead of e_2 for the canonical generator. Similarly when considering it as a function on $H^1(S, \mathbb{G}_{\mathrm{m}})$, we write e^u instead of e_1 for the lift of distinguished lift of the canonical generator of $H_1(L, \mathbf{Z})$.

Any splitting of (4.8.1) is Lagrangian, so provides an OGW framing. The map (4.8.2) carries the canonical generator of $H_1(L, \mathbf{Z})$ to $e_1 + pe_2$; we write this as e^{u+pv} when we think of it as a function on $\mathcal{P}$, with $p = 0$ giving the distinguished splitting. Then v and e^{u+pv} define coordinates on T^*T_L.

4.9.2. Triangular prism. For the triangular prism, the lattice $H_1(S, \mathbf{Z})$ has rank 4.

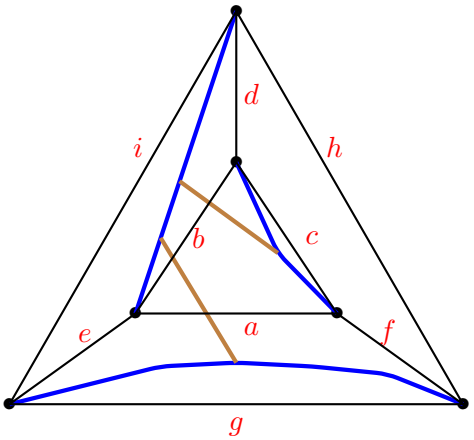

Inside of $H_1(S, \mathbf{Z})$ we can find the product of two triangular lattices — one where $a + b + c = 0$ and one where $g + h + i = 0$. In these coordinates, the relations (4.6.1) imply

$$d = a + g \qquad e = c + h \qquad f = b + i$$

The vectors b, c, h, i make a basis for $H_1(S, \mathbf{Z})$, and the symplectic form on $\mathcal{P}$ is given by $db \wedge dc + dh \wedge di$.

Let L be the branched double cover of the blue tangle in the diagram above. We will see that OGW framings of L are naturally indexed by 2×2 symmetric matrices with integer entries. If we put $u_1 = \log(b), u_2 = \log(h)$ and $v_1 = \log(c), v_2 = \log(i)$, then each of these framings gives an identification of the universal cover of $\mathcal{P}$ with $T^*\mathbf{C}^2$, which always has v_1, v_2 for momentum coordinates and $u_s + M_s^t v_t$, $s, t = 1, 2$, i.e.

$$u_1 + M_1^1 v_1 + M_1^2 v_2 \qquad u_2 + M_2^1 v_1 + M_2^2 v_2$$

for position coordinates.

The two brown arcs lift to two loops in L, which is a handlebody of genus two. Choosing an orientation for each of those loops gives a basis for $H_1(L, \mathbf{Z})$. With respect to this basis for $H_1(L, \mathbf{Z})$, and the (b, c, h, i) basis of $H_1(S, \mathbf{Z})$, the projection $H_1(S, \mathbf{Z}) \to H_1(L, \mathbf{Z})$ is the matrix

$$\begin{pmatrix} \epsilon_b & 0 & 0 & 0 \\ 0 & 0 & \epsilon_h & 0 \end{pmatrix} \tag{4.9.2}$$

where ϵ_b and ϵ_h are arbitrary signs. Let us put $\epsilon_b = \epsilon_h = 1$.

A splitting of the map $H_1(S, \mathbf{Z}) \to H_1(L, \mathbf{Z})$ is given by a 4×2 integer matrix whose product with (4.9.2) is the 2×2 identity matrix, it's general form is

$$\begin{pmatrix} 1 & 0 \\ \alpha & \beta \\ 0 & 1 \\ \gamma & \delta \end{pmatrix}$$

The splitting is Lagrangian if the symplectic pairing of the two columns is zero, i.e. if $\beta = \gamma$. The framing matrix is therefore $M = \left(\begin{smallmatrix} \alpha & \beta \\ \beta & \delta \end{smallmatrix}\right)$. The open Gromov-Witten invariants for this example will be discussed in Section 5.3.

5. Computations, examples

In this section, we choose some graphs Γ and compute superpotentials $W : T_L \to \mathbf{C}$. We also derive a blow-up formula relating moduli spaces a graph and its blow-up, but we begin immediately below with the fundamental example: the tetrahedron.

5.1. Tetrahedron

In the case of a tetrahedron, we recover the results of [AV, AKV, N2].

We begin by choosing an OGW framing, so we continue with the choice from Section 4.9.1. The figure below includes the face coordinates z_i for the framed moduli space.

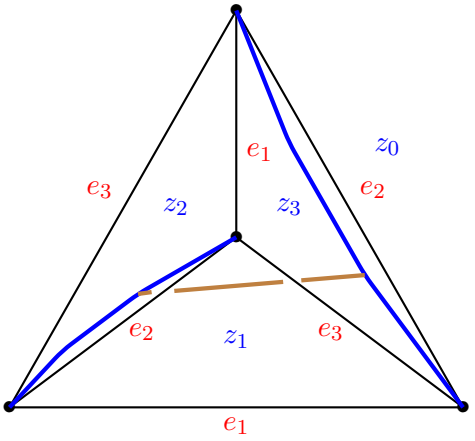

A point of $\mathcal{M}$ is a PGL_2-orbit of quadruples (z_0, z_1, z_2, z_3) which are pairwise distinct. To compute the image of the period map (4.7.1), we may assume $z_0 = 0, z_1 = 1, z_2 = \infty$, and $z_3 = z$. Then we find that $\mathcal{M}$ is parametrized by

$$(x_1, x_2, x_3) = \left(\frac{z}{1-z}, z-1, \frac{-1}{z}\right).$$

The period domain $\mathcal{P} \subset (\mathbf{C}^*)^{\times 3}$ is defined by the face relation $x_1x_2x_3 = 1$, and $\mathcal{M}$ is cut out of $\mathcal{P}$ by the further equation $1 + x_2 + x_1x_2 = 0$. If we identify $\mathcal{P}$ with $(\mathbf{C}^*)^2 = \{(x, y)\}$ using coordinates $x = -x_3$ and $y = -x_2x_3$, then $\mathcal{M}$ is the pair of pants

$$x + y = 1, \tag{5.1.1}$$

We take the OGW framing from Section 4.9.1. We regard e^v and e^{u-pv} as $\mathbf{C}^*$-valued functions on $\mathcal{P}$, with e^v cutting out T_L. Recall that $e^v = e_2 = x_2$ while $e^u = e_1 = x_1$. The canonical 1-form in this framing is given by $vd(u - pv)$.

Set

$$V = -e^{-v}, \qquad U = -(-1)^p e^{u+pv}.$$

The choice of sign in front is a kind of mirror map — see Section 6.1 of [AKV]. It is the ϵ_f of (1.2.1). Write the defining equation as $1 + \frac{1}{x_2} + x_1 = 0$. Then U and V obey

$$UV^p + V = 1.$$

The one-form is taken to be $-\log V d\log U$, a shift of $i\pi d\log U$ from $v d\log U$ due to the mirror map, for which we have no good explanation. (Similar sign choices and shifts will be presented without comment in later examples.)

Let us consider the canonical ($p = 0$) framing. Then since we can solve for $V = 1 - U$, we have for the one-form

$$-\log(1-U)U^{-1}dU = \sum_{n=0}^{\infty} \frac{1}{n} U^{n-1} dU = d\left(\sum_{n=0}^{\infty} \frac{1}{n^2} U^n \right) = d\mathrm{Li}_2(U).$$

The conjectural open Gromov-Witten generating function is therefore $W^{(p=0)} = \mathrm{Li}_2(U)$. It obeys the integrality condition, Equation 5.4.1.

For a general OGW framing labeled by p, the corresponding open Gromov-Witten invariants have generating function $W^{(p)}$ precisely as in the work of Aganagic-Klemm-Vafa — see Section 6.1 of [AKV].

Many examples and moduli spaces can be related to the tetrahedron case by means of a blow-up formula, which we now describe.

5.2. Blow-up formula

Let Γ be a simple cubic planar graph, S the corresponding Legendrian, $\mathcal{M}$ the moduli space and $\mathcal{P}$ the period domain. Let $v \in \Gamma$ be a vertex. The blow-up of Γ at v is a new graph Γ' constructed from Γ by a local modification: v is replaced by a small "exceptional" triangle with vertices connected to the edges incident to v (see picture below). Let Γ' be the blow-up of Γ at v, and define S', $\mathcal{M}'$ and $\mathcal{P}'$ respectively.

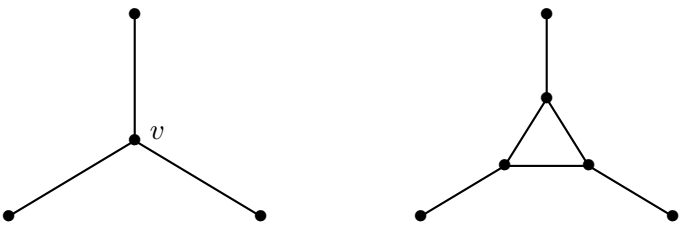

Blowing up at a vertex v

Proposition 5.1. *$\mathcal{P}'$ has a symplectic decomposition $\mathcal{P}' \cong \mathcal{P} \times (\mathbf{C}^*)^2$ with respect to which $\mathcal{M}' \subset \mathcal{P}'$ is a Lagrangian product $\mathcal{M} \times H$, where $H \subset (\mathbf{C}^*)^2$ is the pair of pants $x + y = 1$.*

Proof. Let E be the set of edges of Γ and E' the set of edges of Γ'. Note $E \subset E'$, where edges of Γ incident to v are mapped to their obvious counterparts (proper transforms) in Γ'. Let $N := \mathbf{Z}\langle E\rangle$ and $N' := \mathbf{Z}\langle E'\rangle$ be the respective edge lattices, each endowed with their induced intersection forms A and A' from Equation (4.6.2). We define an inclusion $i : N \to N'$ as follows. For an edge $e \in E$ not incident to v, $i(e) = e$. For e incident to v put $i(e) = e - e'$, where e' is the unique exceptional edge *not* adjacent to e. The inclusion i induces a map of the same name, $i : N \to N'$. Put N_0 for the lattice generated by the three exceptional edges. Then a simple case-by-case check shows that

$$N' \cong i(N) \oplus N_0 \text{ is an orthogonal decomposition with respect to } A', \text{ and } A'|_{i(N)} \cong A$$

Now i^* induces a map $j : Hom(E', \mathbf{C}^*) \to Hom(E, \mathbf{C}^*)$. Let ω' and ω be the two-forms on these spaces defined by A' and A. It follows that these forms are related by $\omega' = j^*\omega + \omega_0$ where ω_0 is $\sum_{i=1}^{3} \frac{dx_i}{x_i} \wedge \frac{dx_j}{x_j}$, the sum taken over the three exceptional divisors. These x_i obey $\prod_{i=1}^{3} x_i = 1$ and as in Equation 5.1.1 parametrize a pair of pants. That these pants split as a Cartesian factor follows from the observation (by direct calculation) that for $e \in E$ we have $x_e = -x_{i(e)}$. □

Remark 5.2 (Tetrahedron, revisited). The tetrahedron graph is the blow-up of the "Θ" graph, the unique (non-simple) planar graph with two vertices and three edges. The Θ graph has zero symplectic form and moduli space equal to a point, so Proposition 5.1 establishes that the moduli space for the tetrahedron, which corresponds to the Aganagic-Vafa brane, is a pair of pants — as we have already seen in Section 5.1 above.

5.3. Triangular prism

The triangular prism is the blow-up of the tetrahedron graph at any vertex. We label the edges as in Section 4.9.2.

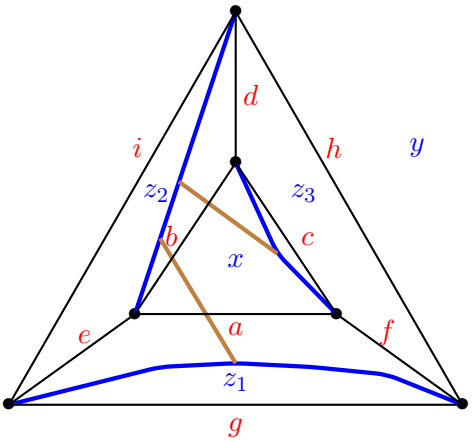

Figure 5.3.1: The blue arcs denote the tangle. The brown arcs generate $H_1(L, \mathbf{Z})$.

Using the blow-up procedure, we confirm that the intersection form in the basis

$$\{a, b, c, g, h, i, d - a - g, e - c - h, f - b - i\}$$

is

$$H \oplus H \oplus 0, \qquad H = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$$

Let's write down the period maps.

$$\begin{array}{lll} x_a = -\frac{x-z_2}{z_2-z_1} \cdot \frac{z_1-z_3}{z_3-x} & x_d = -\frac{z_2-x}{x-z_3} \cdot \frac{z_3-y}{y-z_2} & x_g = -\frac{z_1-z_2}{z_2-y} \cdot \frac{y-z_3}{z_3-z_1} \\ x_b = -\frac{x-z_3}{z_3-z_2} \cdot \frac{z_2-z_1}{z_1-x} & x_e = -\frac{z_1-x}{x-z_2} \cdot \frac{z_2-y}{y-z_1} & x_h = -\frac{z_3-z_1}{z_1-y} \cdot \frac{y-z_2}{z_2-z_3} \\ x_c = -\frac{x-z_1}{z_1-z_3} \cdot \frac{z_3-z_2}{z_2-x} & x_f = -\frac{z_1-y}{y-z_3} \cdot \frac{z_3-x}{x-z_1} & x_i = -\frac{z_2-z_3}{z_3-y} \cdot \frac{y-z_1}{z_1-z_2} \end{array}$$

With the blow-up basis as our guide, we note the following relations.

$$\begin{gathered} x_a x_b x_c = 1, \qquad 1 + x_c + x_b x_c = 0; \\ x_g x_h x_i = 1, \qquad 1 + x_h + x_g x_h = 0; \\ \frac{x_d}{x_a x_g} = -1, \qquad \frac{x_e}{x_c x_h} = -1, \qquad \frac{x_f}{x_b x_i} = -1. \end{gathered}$$

We see that the intersection form is nondegenerate on $H_1(S, \mathbf{Z})$ generated by $\{b, c, g, h\}$ and the image of the period map is the Cartesian product in $(\mathbf{C}^*)^4$ of two pairs of pants.

Let us continue our analysis by picking up from Section 4.9.2, where we considered an OGW phase and family of framings described by the framing matrix $M = \left(\begin{smallmatrix}\alpha & \beta \\ \beta & \delta\end{smallmatrix}\right)$. We begin by considering "zero framing," $M = 0$. Proceeding by analogy with Section 5.1, define the corresponding coordinates as follows. Put

$$U_1 = -x_b = -e^{u_1}, \quad V_1 = -\frac{1}{x_c} = -e^{-v_1};$$
$$U_2 = -x_h = -e^{u_2}, \quad V_2 = -\frac{1}{x_i} = -e^{-v_2}.$$

Then

$$U_1 + V_1 = 1, \qquad U_2 + V_2 = 1. \tag{5.3.1}$$

The symplectic form is $\sum_i d\log(V_i) \wedge d\log(U_i)$. The moduli space is Lagrangian and can be written as the graph of $dW(U_1, U_2)$. That is, we can solve the equation $-\log(V_i) = \partial_{\log(U_i)} W$ for W. To do so, solve the defining equation for $-\log(V_i) = -\log(1 - U_i)$, which gives

$$W = \mathrm{Li}_2(U_1) + \mathrm{Li}_2(U_2).$$

We now study how W changes for different framings. Make the change of coordinates $U_1 \to U_1 V_1^{\alpha} V_2^{\beta} (-1)^{\alpha}$, $U_2 \to U_2 V_1^{\beta} V_2^{\delta}$, with V_i unchanged, so that Equation 5.3.1 now reads

$$U_1 V_1^{\alpha} V_2^{\beta} (-1)^{\alpha} + V_1 = 1, \qquad U_2 V_1^{\beta} V_2^{\delta} (-1)^{\delta} + V_2 = 1.$$

We then try to write $-\log(V_i) = \partial_{\log(U_i)} W(U)$. The equations define a new function $W(U)$ for each choice of M. If M is diagonal, $M = \mathrm{diag}(\alpha, \delta)$, then the equations above decouple and $W = W^{(\alpha)}(x_1) + W^{(\delta)}(x_2)$, where $W^{(p)}$ is the framing-p superpotential as in Section 5.1.

Let us investigate some non-diagonal framings, M. Consider the family $M = \left(\begin{smallmatrix}0 & p \\ p & 0\end{smallmatrix}\right)$. When $p = 1$ we can solve the equations for the V_i. The equations are

$$U_1 V_2 + V_1 = 1, \qquad U_2 V_1 + V_2 = 1.$$

We can solve

$$-\log V_1 = -\log(1 - U_1) + \log(1 - U_1 U_2),$$
$$-\log(V_2) = -\log(1 - U_2) + \log(1 - U_1 U_2).$$

Putting $-\log(V_i) = \partial_{\log(U_i)} W(U)$ gives

$$W = \mathrm{Li}_2(U_1) + \mathrm{Li}_2(U_2) - \mathrm{Li}_2(U_1 U_2),$$

an integral linear combination of dilagarithms with arguments labeled by $H_1(L, \mathbf{Z})$, as expected — see Equation (5.4.1). When $p = -1$, the equations for the y_i are quadratic, and an exact solution for W seems out of reach. We can instead develop a power series solution $W = \sum_{a,b} K_{a,b} U_1^a U_2^b$, solve for the conjectural open Gromov-Witten invariants $K_{a,b}$. We then want to check that they define integer conjectural BPS numbers after accounting for d-fold covers with the $1/d^2$ multiple-cover formula of Ooguri-Vafa — equivalently, check that to a specified order W has the form

$$W = \sum_{d=(d_1,d_2)} a(d) \, \mathrm{Li}_2(U_1^{d_1} U_2^{d_2})$$

with $a(d)$ integers, where d_1 and d_2 refer to homology classes corresponding to the upper and lower brown arcs of Figure 5.3.1, respectively (and we have suppressed the dependence on framing in the notation $a(d)$). We find for $a(d)$ the following numbers.

$d_1 \setminus d_2$	0	1	2	3	4	5	6	7	8	9
0	0	1	0	0	0	0	0	0	0	0
1	1	1	1	1	1	1	1	1	1	
2	0	1	2	4	6	9	12	16		
3	0	1	4	11	25	49	87			
4	0	1	6	25	76	196				
5	0	1	9	49	196					
6	0	1	12	87						
7	0	1	16							
8	0	1								
9	0									

We have written a computer code to implement this procedure, and integrality has been verified in all of the hundreds of examples checked.

5.4. The cube

The 1-skeleton of a cube is not obtained from the "blow-up" construction, and the moduli space is not a product of pairs of pants, so presents an interesting new test of our methods.

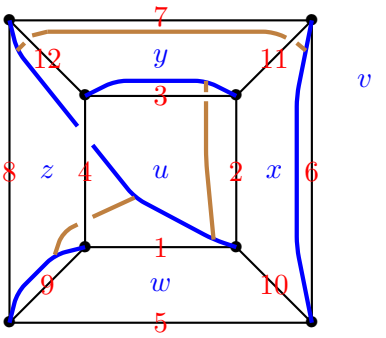

We choose the following basis for $H^1(S, \mathbf{Z})$

$$\{-(e_1 + e_3), e_9; -e_2, e_3; -e_7, e_6\},$$

in which the symplectic form is standard. The phase is evident from the blue tangle, which determines the kernel of the map $H_1(S, \mathbf{Z}) \to H_1(L, \mathbf{Z})$ to be $e_9 =: e^{v_1}, e_3 =: e^{v_2}, e_6 =: e^{v_3}$. In choosing this basis we have also selected a lift of the brown generators of $H_1(L, \mathbf{Z})$, namely $e_1 + e_3 =: e^{u_1}, e_2 =: e^{u_2}, e_7 =: e^{u_3}$. We must try to express the v_i in terms of the u_i.

To do so, calculate the monodromy from Equation (4.7.1):

$$\begin{array}{lll}
x_1 = -\frac{u-z}{z-w} \cdot \frac{w-x}{x-u} & x_2 = -\frac{u-w}{w-x} \cdot \frac{x-y}{y-u} & x_3 = -\frac{u-x}{x-y} \cdot \frac{y-z}{z-u} \\
x_4 = -\frac{u-y}{y-z} \cdot \frac{z-w}{w-u} & x_5 = -\frac{v-x}{x-w} \cdot \frac{w-z}{z-v} & x_6 = -\frac{v-y}{y-x} \cdot \frac{x-w}{w-v} \\
x_7 = -\frac{v-z}{z-y} \cdot \frac{y-x}{x-v} & x_8 = -\frac{v-w}{w-z} \cdot \frac{z-y}{y-v} & x_9 = -\frac{w-u}{u-z} \cdot \frac{z-v}{v-w} \\
x_{10} = -\frac{x-u}{u-w} \cdot \frac{w-v}{v-x} & x_{11} = -\frac{y-u}{u-x} \cdot \frac{x-v}{v-y} & x_{12} = -\frac{z-u}{u-y} \cdot \frac{y-v}{v-z}
\end{array}$$

We have the face relations:

$$\begin{array}{ll}
x_1x_2x_3x_4 = 1 & x_5x_6x_7x_8 = 1 \\
x_1x_9x_5x_{10} = 1 & x_2x_{10}x_6x_{11} = 1 \\
x_3x_{11}x_7x_{12} = 1 & x_4x_{12}x_8x_9 = 1
\end{array}$$

We can use the last 3 to eliminate x_{12}, x_{11}, x_{10} in terms of the others, and use the first two to eliminate x_4, x_5.

$$x_4 = \frac{1}{x_1x_2x_3}, \qquad x_5 = \frac{1}{x_6x_7x_6},$$
$$x_{12} = \frac{x_1x_2x_3}{x_8x_9}, \qquad x_{11} = \frac{x_8x_9}{x_1x_2x_3^2x_7}, \qquad x_{10} = \frac{x_1x_3^2x_7}{x_6x_8x_9}.$$

The last relation gives

$$x_8 = \pm\frac{x_1 x_3}{x_6}$$

So the codomain of the period map is coordinatized by

$$\begin{array}{lll} x_9 = -\frac{w-u}{u-z}\cdot\frac{z-v}{v-w} & x_1 = -\frac{u-z}{z-w}\cdot\frac{w-x}{x-u} & x_2 = -\frac{u-w}{w-x}\cdot\frac{x-y}{y-u} \\ x_3 = -\frac{u-x}{x-y}\cdot\frac{y-z}{z-u} & x_6 = -\frac{v-y}{y-x}\cdot\frac{x-w}{w-v} & x_7 = -\frac{v-z}{z-y}\cdot\frac{y-x}{x-v} \end{array}$$

and we can describe the image, the chromatic Lagrangian, by finding relations among combinations of these variables.

Here's one:

$$x_6\left(1+\frac{1}{x_1 x_3 x_7}\right)+\frac{1}{x_7}+1=0$$

We can eliminate $x_6 = -\frac{x_1x_3x_7+x_1x_3}{x_1x_3x_7+1}$.

In looking for relations, we will set $w=0, x=1, u=\infty$. Then

$$\begin{array}{lll} x_9 = -\dfrac{v-z}{v}, & x_1 = -\dfrac{1}{z}, & x_2 = y-1, \\ x_3 = -\dfrac{y-z}{y-1}, & x_6 = -\dfrac{y-v}{y-1}\cdot\dfrac{1}{v}, & x_7 = -\dfrac{v-z}{z-y}\cdot\dfrac{y-1}{1-v}. \end{array}$$

Corresponding to our symplectic basis, define

$$\begin{array}{lll} a_1 = (x_1x_3)^{-1}, & b_1 = x_9, & a_2 = (x_2)^{-1}, \\ b_2 = x_3, & a_3 = x_7^{-1}, & b_3 = x_6. \end{array}$$

Then

$$\omega = \sum_{i=1}^{3}\frac{1}{a_i b_i}da_i\wedge db_i = \sum_{i=1}^{3}\frac{1}{U_iV_i}dU_i\wedge dV_i$$

where we have defined mirror coordinates U_i and V_i by the following choice of signs:

$$U_i = -a_i^{-1}, \qquad V_i = -b_i^{-1}.$$

We can solve for y, z, v in terms of U_i and use the results to express the V_i as functions of the U_i.

$$y = 1+U_2, \qquad z = \frac{1+U_2}{1-U_1U_2}, \qquad v = \frac{-(U_1U_2+U_1)U_3+U_2+1}{-(U_1U_2+U_1)U_3-U_1U_2+1}$$

from which we get

$$b_1 = \frac{U_1U_2U_3(U_1+1)}{(U_1U_2-1)(U_1U_3-1)}, \quad b_2 = \frac{U_1(U_2+1)}{U_1U_2-1}, \quad b_3 = \frac{U_1(U_3+1)}{U_1U_3-1}.$$

(As an example, the expression for b_3 is equivalent to the relation found above involving $x_6 = b_3$.) Lifting to the universal cover, where all Lagrangians are exact, we define

$$u_i = -\log U_i, \qquad v_i = -\log V_i$$

so

$$\omega = \sum_i du_i \wedge dv_i = -d \sum v_i du_i.$$

Then by above, the v_i are functions of the u_i and Lagrangianicity can be verified:

$$\partial_{u_j} v_i = \partial_{u_i} v_j.$$

We now ask Mathematica determine a function $W(U)$ such that $v_i = \partial_{\log U_i} W$. We find the solution

$$W = \mathrm{Li}_2(U_1) + \mathrm{Li}_2(U_2) + \mathrm{Li}_2(U_3) - \mathrm{Li}_2(U_1 U_2) - \mathrm{Li}_2(U_1 U_3).$$

Since W is expressed purely in terms of dilogarithms of monomials in the U_i, it satisfies the open Gromov-Witten integrality constraint, Equation (5.4.1). The result predicts the existence of unique holomorphic disks in $\mathbf{C}^3$ bounding the smooth, non-exact Lagrangian L in various homology classes labeled by the u_i. In particular, we expect the following BPS numbers for this OGW framing:

$$\boxed{a(1,0,0) = a(0,1,0) = a(0,0,1) = 1, \qquad a(1,1,0) = a(1,0,1) = -1}$$

Appendix: Physical contexts

There is a wide array of physical set-ups where the mathematics of the present paper applies. Here is a partial list.

- **Type-IIA on Noncompact Calabi-Yau.** Type-IIA string theory on a noncompact Calabi-Yau manifold X has an effective four-dimensional supergravity theory with $\mathcal{N} = 2$ supersymmetry and $b_2(X) + 1$ abelian gauge fields arising from the Ramond-Ramond sector. They are part of chiral superfields. (In our set-up, the Calabi-Yau manifold is $\mathbf{C}^3$, so there is more supersymmetry, but it may be a good idea to think of $\mathbf{C}^3$ as a special case of the more general set-up.) Different couplings

of these gauge fields are described by topological string amplitudes at genus g.

A D4-brane whose (five-dimensional) world volume fills a two-plane in spactime cross a supersymmetric Lagrangian three-cycle $L \subset X$ creates a BPS domain wall. These domain walls have a net effect on the 4d physics. The corresponding term in the 4d action, which includes a delta function supported on the domain wall, includes the contribution of D2-D0 branes ending on the D4-brane. These can be computed by open topological string amplitudes: open Gromov-Witten invariants counting holomorphic maps from a disk with boundary lying in L. As we recall below, Ooguri-Vafa derived integrality results for open Gromov-Witten theory by comparison of this set-up with M-Theory, analagous to the Gopakumar-Vafa formula in the closed case.

References: [AGNT, OV]

- **Type II-A, Part 2**

  One could instead consider a D6-brane wrapping all four dimensions of spacetime cross $L \subset X$. In this case, the open Gromov-Witten invariants contribute couplings involving chiral fields, $b_1(L)$ in number, the $b_2(X) + 1$ closed chiral fields discussed above, and the graviphoton multiplet. One such term is the superpotential, and it is computed from disk invariants [OV].

- **M-Theory on a G_2-holonomy manifold.**

  The M-Theory equivalent of IIA on a Calabi-Yau manifold X is M-Theory on $Y = X \times S^1$, where the radius of the circle is related to the string coupling constant. To model the D4-brane set-up, one takes an M5-brane wrapping $\mathbf{R}^2 \times L \times S^1$. Contributions from bound D2-D0 branes are encoded by $M2$-branes with boundary on the M5-brane. The non-perturbative M-theory set-up allows one to calculate these contributions without resorting to the perturbative methods of Gromov-Witten theory. Using this perspective, Ooguri-Vafa determine strict integrality conjectures for open Gromov-Witten theory, analogous to Gopakumar-Vafa formulas in the closed case. As a result, the superpotential W of the four-dimensional theory must have an expression as follows. First choose a (non-canonical) splitting $H_2(X, L; \mathbf{Z}) \cong H_2(X; \mathbf{Z}) \oplus H_1(L, \mathbf{Z})$ along with bases $\{C_j\}$ for $H_2(X; \mathbf{Z})$ and $\{\gamma_i\}$ for $H_1(L; \mathbf{Z})$. Then introduce corresponding coordinates $q_j, x_i \in \mathbf{C}^*$; and for $\beta = \sum_j \beta_j C_j \in H_2(X; \mathbf{Z})$ and $\gamma = \sum_i d_i \gamma_i \in H_1(L; \mathbf{Z})$, define

$q^\beta := \prod_j q_j^{\beta_j}$ and $x^\gamma := \prod_i x_i^{d_i}$. Then

$$W(t, x_i) = \sum_{n=1}^{\infty} \sum_{\beta \in H_2(X)} \sum_{\gamma \in H_1(L)} \sum_{s \in \frac{1}{2}\mathbf{Z}_+} n_{\beta,\gamma,s} \frac{1}{n^2} q^{n\beta} x^{n\gamma} u^s, \qquad n_{\beta,\gamma,s} \in \mathbf{Z}$$

A special case arises when, as for our examples, $H_2(X) = 0$ so the q term disappears (or one can take the $q \to 1$ limit) and we ignore the spin dependence by setting $u \to 1$ as well. We then have

$$W(x) = \sum_{\gamma \in H_1(L;\mathbf{Z})} c_\gamma \mathrm{Li}_2(x^\gamma), \qquad c_\gamma \in \mathbf{Z}, \tag{5.4.1}$$

where $\mathrm{Li}_2(x)$ is the dilogarithm function $\sum_{n \geq 1} \frac{x^n}{n^2}$. The integrality constraint to which we subject our conjectural superpotentials is that W can be written in this form with the c_γ integers.

- **M-Theory, Part 2**

  The M-Theory equivalent of IIA on X with a D6-brane wrapping four-dimensional spacetime cross a Lagrangian $L \subset X$ is entirely geometrical. There are no branes, but rather, the compactification manifold is a G_2-holonomy seven-fold Y, and should take the form of a singular S^1-fibration, where the circle fiber degenerates to a point over the locus L. A local model should be the Taub-NUT geometry in the four dimensions transverse to (the lift of) L in Y. The effective theory on the 4d spacetime has $\mathcal{N} = 1$ supersymmetry. There are as many chiral superfields as there are L^2-harmonic three-forms. (This number equals $b_4(Y) = b_3(Y)$ if Y is compact, but that is not the case here.) This picture is largely conjectural, but for the case of the (singular) Harvey-Lawson brane it is rigorous, as shown by Atiyah and Witten.

  In this set-up, the superpotential for the chiral superfields of the four-dimensional theory is generated by M2-branes. There are no M5-branes present, so the M2-branes are closed. As explained at the end of the Section 3.1 of [AKV], for example, the superpotential term receives contributions from M2-instantons which are homology three-spheres. The three sphere geometry comes from a disk cross the M-theory circle, with the fibers over the boundary circle of the disk identified, since the M-theory circle collapses there.

  These ideas are explored in the following works: [AKV, AMV, AV3, AW].

- **Dimensional reduction of six-dimensional theories**

As discussed in Section 1.4, the prior works [CCV, CNV, CEHRV] as well as [DGG1, DGG2, DGH] employ overlapping mathematical machinery to the present work.

The works of Section 1.4.1 consider the effective theory of M5-branes wrapping spacetime crossed with a branched double cover of $\mathbf{R}^3$. The authors arrive at dual three-dimensional theories by considering different Seifert surfaces bounding the branch locus, a tangle. The effective three-dimensional theory is an abelian $\mathcal{N} = 2$ Chern-Simons theory coupled to chiral matter, whose partition function can be computed by supersymmetric localization. The authors also give an interpretation in terms of a quantum-mechanical wavefunction that bears a strong resemblance to the superpotential W computed here. A main point of these works is that different Seifert surfaces generate equivalent Lagrangian descriptions of a single quantum theory. Invariance of the three-dimensional theory amounts to quantum dilogarithm identities, giving a physical explanation of the Kontsevich-Soibelman wall-crossing formulas.

The works of Section 1.4.2 propose various dualities by considering dimensional reductions of supersymmetric theories on brane worldvolumes. For example, one can consider a stack of n M5-branes wrapping $\mathbf{R}^2 \times S^1 \times L$, where L is a (typically noncompact) three-dimensional Lagrangian in a (typically noncompact) Calabi-Yau three-fold X. Boundary conditions on L are fixed: some prototypes are when L is a knot complement for a hyperbolic knot $K \subset S^3$ or a Harvey-Lawson/Aganagic-Vafa brane in a toric Calabi-Yau threefold. By relating the six-dimensional theory on the brane to the reduction to *either* $\mathbf{R}^2 \times S^1$ or to the three-fold L, a duality is proposed in [DGH, DGG2] between a three-dimensional $\mathcal{N} = 2$ supersymmetric theory on $\mathbf{R}^2 \times S^1$ and an $SL(n, \mathbf{C})$ Chern-Simons (CS) theory on L. The CS partition function has an expression in terms of dilogarithms arising as the hyperbolic volumes of ideal tetrahedra. The partition function of the $\mathcal{N} = 2$ theory is an index, counting BPS states, and involves dilogarithms, as we relate below.

[The earlier work of [DGH] considered a further reduction on S^1 to a two-dimensional $\mathcal{N} = (2, 2)$ supersymmetric theory on $\mathbf{R}^2$. In the reduction to two dimensions, the partition function is expressed as a sum over vortex instantons on $\mathbf{C}$, and the vortex center locations are described by symmetric polynomials. The fixed-point theorem relates the character over this moduli space to monomial contributions, leading to the dilogarithm expressions above.]

As described in [DGH], all this can be seen from the spacetime perspective by decoupling BPS states from the bulk by setting the Kähler parameters of the Calabi-Yau to large volume, similar to the $H_2(X, \mathbf{Z}) = 0$ case discussed above.

References

[AK] I. Agol and S. Krushkal, *Structure of the flow and Yamada polynomials of cubic graphs*, arXiv:1801.00502.

[AGNT] I. Antoniadis, E. Gava, K. S. Narain, and T. R. Taylor, *Topological amplitudes in string theory,* Nucl. Phys. **B413** (1994), 162–184.

[AMV] M. Atiyah, J. Maldacena, and C. Vafa, *An M-theory flop as a large $\mathcal{N}$ duality*, arXiv:hep-th/0011256.

[AV] M. Aganagic and C. Vafa, *Mirror symmetry, D-branes and counting holomorphic discs*, arXiv:hep-th/0012041.

[AV2] M Aganagic and C. Vafa, *Large N duality, mirror symmetry, and a Q-deformed A-polynomial for knots*, arXiv:1204.4709.

[AV3] M. Aganagic and C. Vafa, *Mirror symmetry and a G_2 flop*, arXiv:hep-th/0105225.

[AW] M. Atiyah and E. Witten, *M-theory dynamics on a manifold of G_2 holonomy*, arXiv:hep-th/0107177.

[AKV] M. Aganagic, A. Klemm, and C. Vafa, *Disk instantons, mirror symmetry and the duality web*, Z. Naturforsch. **A57** (2002) 1–28; arXiv:hep-th/0105045.

[B] G. Birkhoff, *A determinant formula for the number of ways of coloring a map*, Ann. Math. (1912), 42–46.

[BD] C. Brav and T. Dyckerhoff, *Relative Calabi-Yau structures*, arXiv:1606.00619.

[BM] H. Brugesser and P. Mani, *Shellable decompositions of cells and spheres*, Math. Scand. **29** (1971), 197–205.

[CCV] S. Ceccoti, C. Córdova, and C. Vafa, *Braids, walls, and mirrors*, arXiv:1110.2115.

[CNV] S. Ceccoti, A. Neitzke, and C. Vafa, *R-twisting and 4d/2d correspondences*, arXiv:1006.3435.

[CEHRV] C. Córdova, S. Espahbodi, B. Haghighat, A. Rastogi, and C. Vafa, *Tangles, generalized Reidemeister moves, and thre-dimensional mirror symmetry*, arXiv:1211.3730.

[D] G. Dimitroglou Rizell, *Knotted Legendrian surfaces with few Reeb chords*, Algebraic and Geometric Topology **11** (2011).

[DGG1] T. Dimofte, M. Gabella, and A. Goncharov, *K-deompositions and 3d gauge theories*, arXiv:1301.0192.

[DGG2] T. Dimofte, D. Gaiotto, and S. Gukov, *Gauge theories labelled by three-manifolds*, Commun. Math. Phys. **325** (2014), 367–417; arXiv:1108.4389.

[DGH] T. Dimofte, S. Gukov, and L. Hollands, *Vortex counting and Lagrangian* 3*-manifolds*, Lett. Math. Phys. **98** (2011), 225–287; arXiv:1006.0977.

[E] T. Ekholm, *Morse flow trees and Legendrian contact homology in* 1*-jet spaces,* Geom. Topol. **11** (2007), 1083–1224.

[EHK] T. Ekholm, K. Honda, and T. Kálmán, *Legendrian knots and exact Lagrangian Cobordisms*, arXiv:1212.1519.

[GS] Y. Groman and J. Solomon, *A reverse isoperimetric inequality for* J*-holomorphic curves*, Geom. Funct. Anal. **24** (2014), 1448–1515.

[GKS] S. Guillermou, M. Kashiwara, and P. Schapira, *Sheaf quantization of Hamiltonian isotopies and applications to nondisplaceability problems*, Duke Math J. **161** (2012), 201–245.

[J] D. Joyce, *Special Lagrangian submanifolds with isolated conical singularities V: Survey and applications*, Jour. Diff. Geom. **63** (2003), 279–347.

[KL] S. Katz and C.-C. M. Liu, *Enumerative geometry of stable maps with Lagrangian boundary conditions and multiple covers of the disc*, Adv. Theor. Math. Phys. **5** (2000), 1–49.

[KS] M. Kashiwara and P. Schapira, *Sheaves on manifolds*, Grundlehren der Mathematischen Wissenschafte **292**, (Springer-Verlag, 1990).

[N] D. Nadler, *Microlocal branes are constructible sheaves*, Selecta Math. (N.S.) **15** (2009), 563–619.

[N2] D. Nadler, *A combinatorial calculation of the Landau-Ginzburg model* $M = \mathbb{C}^3$, $W = z_1 z_2 z_3$, arXiv:1507.08735.

[NZ] D. Nadler and E. Zaslow, *Constructible sheaves and the Fukaya category*, J. Amer. Math. Soc. **22** (2009), 233–286.

[OV] H. Ooguri and C. Vafa, *Knot invariants and topological strings*, Nucl. Phys. **B577** (2000), 419–438.

[P] Y. Pan, *Exact Lagrangian fillings of Legendrian* $(2, n)$ *torus knots*, arXiv:1607.03167.

[RS] D. Rutherford and M. Sullivan, *Cellular Legendrian contact homology for surfaces, Pari I*, arXiv:1608.02984; and *Cellular Legendrian contact homology for surfaces, Pari II*, arXiv:1608:03011.

[SS] S. Scherotzke and N. Sibilla, *The nonequivariant coherent constructible correspondence and tilting*, Selecta Math, to appear; arXiv:1402.3360.

[ST1] V. Shende and A. Takeda, *Symplectic structures from topological Fukaya categories*, arXiv:1605.02721.

[ST2] J. Solomon and S. Tukachinsky, *Point-like bounding chains in open Gromov-Witten theory*, arXiv:1608.02495.

[STWZ] V. Shende, D. Treumann, H. Williams, and E. Zaslow, *Cluster varieties from Legendrian knots*, arXiv:1512.08942.

[STW] V. Shende, D. Treumann, and H. Williams, *On the combinatorics of exact Lagrangian surfaces*, arXiv:1603.07449.

[STZ1] L. Shen, D. Treumann, and E. Zaslow, *Framing duality*, in preparation.

[STZ2] V. Shende, D. Treumann, and E. Zaslow, *Legendrian knots and constructible sheaves*, arXiv:1402.0490.

[Y] M.-L. Yau, *Surgery and invariants of Lagrangian surfaces*, arXiv:1306.5304.

Department of Mathematics, Boston College
Chestnut Hill, MA 02467-3806, USA
E-mail address: treumann@bc.edu

Department of Mathematics, Northwestern University
Evanston, IL 60208-2730, USA
E-mail address: zaslow@math.northwestern.edu