hep-th 1304.5778

arXiv:1304.5778v1 [hep-th] 21 Apr 2013

Topological Strings,
D-Model,
and
Knot Contact Homology

Mina Aganagic[1,2], Tobias Ekholm[3,4], Lenhard Ng[5]
and
Cumrun Vafa[6]

[1]*Center for Theoretical Physics, University of California, Berkeley, USA*
[2]*Department of Mathematics, University of California, Berkeley, USA*
[3]*Department of Mathematics, Uppsala University, Uppsala, Sweden*
[4]*Institut Mittag-Leffler, Djursholm, Sweden*
[5]*Department of Mathematics, Duke University, Durham, USA*
[6]*Jefferson Physical Laboratory, Harvard University, Cambridge, USA*

Abstract

We study the connection between topological strings and contact homology recently proposed in the context of knot invariants. In particular, we establish the proposed relation between the Gromov-Witten disk amplitudes of a Lagrangian associated to a knot and augmentations of its contact homology algebra. This also implies the equality between the Q-deformed A-polynomial and the augmentation polynomial of knot contact homology (in the irreducible case). We also generalize this relation to the case of links and to higher rank representations for knots. The generalization involves a study of the quantum moduli space of special Lagrangian branes with higher Betti numbers probing the Calabi-Yau. This leads to an extension of SYZ, and a new notion of mirror symmetry, involving higher dimensional mirrors. The mirror theory is a topological string, related to D-modules, which we call the "D-model." In the present setting, the mirror manifold is the augmentation variety of the link. Connecting further to contact geometry, we study intersection properties of branches of the augmentation variety guided by the relation to D-modules. This study leads us to propose concrete geometric constructions of Lagrangian fillings for links. We also relate the augmentation variety with the large N limit of the colored HOMFLY, which we conjecture to be related to a Q-deformation of the extension of A-polynomials associated with the link complement.

1 Introduction

The subject of topological strings has been important in both physics and mathematics. From the physical perspective it has been studied from many different viewpoints and plays a prominent role in the understanding of string dualities. From the mathematical perspective, its underpinnings (Gromov-Witten invariants) are well understood and play a central role in symplectic geometry and related subjects. Moreover, in many cases one can directly compute the amplitudes using either physical methods or mathematical techniques and the two agree.

An interesting application of topological strings involves the study of knot and link invariants. In particular it is known that HOMFLY polynomials for links K can be reformulated in terms of topological strings on T^*S^3 where one includes N Lagrangian D-branes wrapping S^3 and some Lagrangian brane L_K intersecting S^3 along the link and asymptotic to the conormal of K at infinity. In this setup, it is possible to shift L_K off of the 0-section $S^3 \subset T^*S^3$, which suggests that the data of the link is imprinted in how L_K intersects the boundary of T^*S^3 at infinity. This ideal boundary is the unit cotangent bundle U^*S^3, which is a contact manifold (topologically $S^3 \times S^2$) with contact 1-form given by the restriction of the action form $p\,dq$. The intersection of L_K with U^*S^3 is a Legendrian torus Λ_K.

Large N transition relates this theory to a corresponding theory in the geometry where S^3 shrinks to zero size and is blown up with blow-up parameter given by Ng_s, where g_s is the string coupling constant. Under the transition, boundary conditions corresponding to strings ending on branes wrapping S^3 close up and disappear, but the transition does not affect the Legendrian torus Λ_K, which is part of the geometry infinitely far away from the tip of the singular cone where S^3 shrinks and gets replaced by S^2.

Contact homology is a theory that was introduced in a setting corresponding exactly to the above mentioned geometry at infinity. More precisely, one studies the geometry of $\mathbb{R} \times \Lambda_K \subset \mathbb{R} \times U^*S^3$, which we can naturally view as the far distance geometry of L_K, using holomorphic disks. In this particular case of conormal tori Λ_K in the unit cotangent bundle of S^3, the resulting theory is known as knot contact homology. From the viewpoint of string theory, this theory involves a Hilbert space of physical open string states corresponding to Reeb chords (flow segments of the Reeb vector field with endpoints on Λ_K) and a BRST-operator deformed by holomorphic disks in $\mathbb{R} \times U^*S^3$ with boundary on $\mathbb{R} \times \Lambda_K$ and with punctures mapping to Reeb chords at infinity. In this language, knot contact homology is the corresponding BRST-cohomology. From a mathematical perspective this structure can be assembled into a differential graded algebra (DGA) which is well studied and in particular can be computed from a braid presentation of a link through a concrete combinatorial description of all relevant moduli spaces of holomorphic disks.

In view of the above discussion it is natural to ask how this DGA is related with the open version of Gromov-Witten theory of L_K. In [1] a partial answer for this was proposed: it was conjectured that the augmentation variety associated to a knot, which is a curve that parametrizes one-dimensional representations of the DGA associated to the knot, is the same as the moduli space of a single Lagrangian L_K corrected by disk instantons. This moduli space is in turn, according to a generalization of SYZ conjecture, related to the mirror curve, which is the locus of singular fibers in the bundle of conics that constitutes the mirror, and as such determines the mirror.[1]

One aim of this paper is to indicate a path to a mathematical proof of this conjecture: in Section 6.3 we show how to get a local parametrization of the augmentation curve in terms of the open Gromov-Witten potential of a Lagrangian filling of the conormal torus of a knot, and that expression agrees with the local parametrization of the mirror curve derived using physical arguments, see Section 2.4. In particular, for knots with irreducible augmentation variety, using the conormal filling we obtain a proof of the conjecture, see Remark 6.7.

---

[1] The A-polynomial curves and their deformations were studied in a context of (generalizing) the volume conjecture in [2], and its relation to topological string was studied in [3, 4, 5, 6, 7]. Mirror symmetry and torus knots were studied in [8] and also in [9].

A further aim is to generalize this from one-component knots to many-component links, and the generalization turns out to involve interesting new ingredients. We consider new Lagrangian branes that have the same asymptotics as L_K but have a topologically different filling. Similar new Lagrangian branes in fact appear already in the case of knots, where different 1-cycles of the Legendrian torus bound and can be shrunk in the different branes. For links with $n > 1$ components, there are further possibilities with Lagrangian branes of various numbers of components. The maximal number of components is n, in which case no two components of the conormal tori of the link at infinity are connected through the brane. However, there are Lagrangian fillings and corresponding branes that connect some of the components at infinity. For example, there is a single component Lagrangian brane which connects them all, see Section 6.6 for a conjectural geometric construction.

Our study of the quantum moduli space of all these Lagrangians leads to a new reformulation of mirror symmetry: the moduli space of branes for an n-component link is n-dimensional and the mirror geometry, instead of being a 1-dimensional curve, will be an n-dimensional variety. It turns out that this variety can be viewed as a Lagrangian variety $V \subset (\mathbb{C}^*)^{2n}$ in a canonical way and we identify it with the augmentation variety from knot contact homology.

An unusual feature of this structure is that we encounter higher-dimensional geometries as the mirror. In fact, it may appear that one cannot formulate topological string in this context since the critical dimension for that theory is effectively 1 for non-compact Calabi-Yau 3-folds but for n-component links we have effective dimension $n > 1$. Nevertheless we propose a string theory mirror even in these cases, based on what we call the topological "D-model". The D-model is the A-model topological string on $(\mathbb{C}^*)^{2n}$ with a Lagrangian D-brane V and a canonical coisotropic brane filling the whole $(\mathbb{C}^*)^{2n}$. In this context the topological string is exact at the level of the annulus, with the analog of higher genus A-model corrections being captured not by higher genera of strings, but by the choice of a flux turned on along the coisotropic brane. The D-model leads to a natural definition of the theory in terms of D-modules (hence the name), while for $n = 1$ (in particular for knots), the D-model is already known to be equivalent to the B-model.

The organization of this paper is as follows. In Section 2 we review the relation between topological strings and Chern-Simons theory, large N transition, and knot invariants. Furthermore, we describe, using a generalized SYZ formulation, how any knot gives a mirror geometry. In Section 3 we introduce new Lagrangians associated to knots with the usual asymptotics at infinity but with different interior topology. We then generalize the discussion from knots to n-component links and show how n-dimensional Lagrangian varieties in flat space $(\mathbb{C}^*)^{2n}$ arise in the the study of the moduli space of Lagrangian branes filling the link conormal. In Section 5, we introduce basic elements of knot contact homology; furthermore, we relate augmentation varieties with the (disk instanton corrected) moduli spaces of Lagrangians associated to knots and links, and study intersection properties of branches of the augmentation variety guided by the D-model. In Section 6, we present a mathematical argument relating knot contact homology for links to disk amplitudes in Gromov-Witten theory, and study geometric constructions of Lagrangian fillings for conormals as well as properties and applications of the resulting Lagrangians. In Section 7 we present some examples. In Section 8, we formulate a conjecture about how to quantize higher dimensional augmentation varieties in terms of the D-model, by relating them to D-modules.

Finally, Appendices A and B contain calculations of a more technical nature.

2 Review

Consider Chern-Simons gauge theory with gauge group $G = SU(N)$ on a closed 3-manifold M at level k, where k is a positive integer. The Chern-Simons partition function is given by the path integral

$$Z_{CS}(S^3) = \int \mathcal{D}A\; e^{\frac{ik}{4\pi} S_{CS}(A)}$$

over the space of connections A with values in the Lie algebra of the gauge group, where

$$S_{CS}(A) = \int_M \mathrm{Tr}\Big(A \wedge dA + \tfrac{2}{3} A \wedge A \wedge A\Big)$$

is the Chern-Simons action of the connection A. The path integral is independent of metric on M and hence gives a topological invariant of 3-manifolds. Submanifolds of dimension 0 in M, i.e. points, carry no topological information but submanifolds of dimension 1, i.e. knots and links, do and in Chern-Simons theory there are corresponding topologically invariant observables. More precisely, we associate a Wilson loop observable in representation R to a knot K by inserting the path ordered exponential

$$\mathrm{Tr}_R\, U(K),$$

where $U(K) = P \exp i \oint_K A$ is the holonomy of A along K, in the path integral. In fact, to define this we also have to choose a framing of the knot, i.e. a non-vanishing vector field in the normal bundle of the knot $K \subset M$. The case of many-component links is similar: a link K is a collection

$$K = K_1 \cup K_2 \cup \ldots K_n$$

of disjoint knots K_j in M. We specify a framing of each knot component of the link and representations

$$R_1,\ R_2, \ldots, R_n$$

coloring $K_1, \ldots, K_n$, respectively. The corresponding link invariant is then the expectation value

$$\langle \mathrm{Tr}_{R_1} U(K_1) \ldots \mathrm{Tr}_{R_n} U(K_n) \rangle = Z_{R_1,\ldots,R_n}(M; K_1, \ldots, K_n)/Z(M)$$

obtained by computing the Chern-Simons path integral with insertion of link observables:

$$Z_{R_1,\ldots,R_n}(K_1, \ldots, K_n) = \int \mathcal{D}A \, e^{\frac{ik}{4\pi} S_{CS}(A)} \ \mathrm{Tr}_{R_1} U(K_1) \ldots \mathrm{Tr}_{R_n} U(K_n),$$

and normalizing it with the path integral in the vacuum.

In [10], Witten explained how to solve the above theory exactly. Any three dimensional topological theory corresponds to a two dimensional rational conformal field theory (CFT). The Hilbert space of the three dimensional theory and operators acting on it can be constructed from conformal blocks of the CFT and from representations of the corresponding modular group. In the Chern-Simons case, the relevant conformal field theory is the $SU(N)_k$ Wess-Zumino-Witten (WZW) model, and one finds that knowledge of the S, T and braiding matrices is all that is needed to solve the theory on any 3-manifold.

In this way, invariants of knots and links in the 3-sphere S^3 that arise from Chern-Simons theory can be explicitly computed. In particular, the polynomial knot invariants considered earlier by Jones correspond to the gauge group $G = SU(2)$ and Wilson lines in the fundamental representation. More generally, for a link $K \subset S^3$ with knot components $K = K_1 \cup \cdots \cup K_n$, the expectation values

$$\langle \mathrm{Tr}_{R_1} U(K_1) \ldots \mathrm{Tr}_{R_n} U(K_n) \rangle = H_{R_1,\ldots,R_n}(K_1, \ldots, K_n)(q, Q)$$

are polynomials in the variables

$$q = e^{\frac{2\pi i}{k+N}}, \qquad Q = q^N$$

with integer coefficients that are independent of both k and N. These polynomials are known as HOMFLY polynomials and were constructed from a mathematical point of view in [11].

2.1 Chern-Simons theory and topological string

Chern-Simons theory on S^3 with gauge group $SU(N)$ is intimately related to the open topological A-model, or Gromov-Witten theory, on

$$X = T^*S^3$$

with N Lagrangian branes on the zero section $S^3 \subset T^*S^3$ as follows. The topological A-model corresponds to counting holomorphic maps with Lagrangian boundary conditions. In T^*S^3, any holomorphic map with boundary on the zero section has vanishing area and is therefore constant. Thus, all maps that contribute to the A-model partition function $Z_{GW}(X)$ are degenerate and it was shown in [12] that their contributions are exactly captured by the Feynman diagrams of $SU(N)$ Chern-Simons theory on S^3. Consequently, the partition functions of the topological A-model on X equals the Chern-Simons partition function on S^3:

$$Z_{GW}(X) = Z_{CS}(S^3),$$

where, as mentioned above, Z_{GW} localizes on the 0-dimensional space of holomorphic maps and is thus given by the (exponentiated) generating function

$$Z_{GW}(X) = \exp\left(\sum_{g,h \geq 0} F_{g,h}\ N^h g_s^{2g-2+h} \right),$$

where $F_{g,h}$ captures the contribution of maps of connected genus g Riemann surfaces to X with h boundary components mapping to $S^3 \subset X$. Here, each boundary component is weighted by a factor N corresponding to the choice of which of the N D-branes wrapping S^3 that it lands on, and the genus counting parameter (or string coupling constant) of the open topological string, g_s, equals the effective value of Chern-Simons coupling constant:

$$g_s = \frac{2\pi i}{k+N}.$$

From the perspective of Chern-Simons perturbation theory, the numbers $F_{g,h}$ arise by organizing the Feynman graphs in the following way: thicken the graphs into ribbon graphs with gauge index labels on the boundary; the number $F_{g,h}$ then keeps track of the contributions from the graphs that give rise to ribbon graphs corresponding to a genus g Riemann surface with h boundary components. We also point out that the parameter q, in terms of which the Chern-Simons knot invariants become polynomial, is given by

$$q = e^{g_s}.$$

Knots and links can be included in the correspondence between Chern-Simons and the topological A-model in the following way [13]. To each knot K in S^3, we associate a Lagrangian L_K in X, which we take to be its conormal in T^*S^3 consisting of all covectors along the knot that annihilate its tangent vector. In particular, intersecting L_K with the zero section, we get the knot itself,

$$L_K \cap S^3 = K.$$

For n-component links $K = K_1 \cup \cdots \cup K_n$ in S^3, we consider the Lagrangian L_K which is the union of the conormals of its components

$$L_K = L_{K_1} \cup \ldots \cup L_{K_n}.$$

We will consider D-branes on L_{K_j} and therefore need to include a sector in the theory that corresponds to worldsheets with boundaries both in the branes on L_K and in the branes on the zero section S^3. We write the partition function of the topological string on X with these branes present as

$$Z_{GW}(X, L_{K_1}, \ldots, L_{K_n})$$

and note that, in addition to depending on g_s and N, it also depends on the moduli of the Lagrangians L_{K_j} which in particular keeps track of the class in $H_1(L_{K_j})$ represented by the boundaries of the worldsheets.

In the case under consideration, L_{K_j} each has the topology of $S^1 \times \mathbb{R}^2$ and we get one modulus x_j for each Lagrangian L_{K_j}. Here, the complex parameter x_j can be written as $x_j = r_j + i \oint_{S^1} A_j$, where the real part r_j can be viewed as coming from the moduli of the deformations of the Lagrangian and the imaginary part is the holonomy of the $U(1)$ gauge field of the brane around the nontrivial cycle $S^1 \times \{0\} \subset L_{K_j}$. In later sections, we will study similar Lagrangians for K of first Betti number $b_1 = r > 1$, in which case the moduli of the Lagrangian is r-dimensional. Giving x_j a nonzero value corresponds to lifting the Lagrangian L_{K_j} off of the S^3.

From the Chern-Simons perspective, assuming there is a single D-brane on L_K for a knot K, computing the partition function $Z_{GW}(X; L_K)$ corresponds to inserting the operator

$$\det(1 - e^{-x}\, U(K))^{-1}. \tag{2.1}$$

This describes the effect of integrating out the bifundamental strings, with one boundary on S^3 and the other on L_K. To relate this to knot invariants, we formally expand the determinant:

$$\det(1 \otimes 1 - e^{-x} \otimes U(K))^{-1} \;= \sum_{S_k} \mathrm{Tr}_{S_k}\, U(K)\, e^{-kx},$$

where the sum ranges over all totally symmetric representations S_k of $SU(N)$ of rank k. Thus, computing in $SU(N)$ Chern-Simons theory on S^3 the following weighted sum of expectation values,

$$\Psi_{K_1,\ldots,K_n}(x_1,\ldots,x_n) = \sum_{k_1,\ldots k_n} \Big\langle \mathrm{Tr}_{S_{k_1}}\, U(K_1) \ldots \mathrm{Tr}_{S_{k_n}}\, U(K_n) \Big\rangle \; e^{-k_1 x_1 - \ldots - k_n x_n},$$

gives the topological string partition function on X with single branes on $L_{K_1}, \ldots, L_{K_n}$ and N branes on S^3. In what follows, we will denote the

HOMFLY polynomials

$$\left\langle \mathrm{Tr}_{S_{k_1}} U(K_1) \ldots \mathrm{Tr}_{S_{k_n}} U(K_n) \right\rangle = H_{S_{k_1} \ldots S_{k_n}}(K_1, \ldots, K_n)$$

simply by

$$H_{S_{k_1} \ldots S_{k_n}}(K_1, \ldots, K_n) = H_{k_1 \ldots k_n}(K_1, \ldots, K_n).$$

Since we isolated the part of the topological string amplitude on X with some boundary component on the Lagrangian branes on $L_{K_1} \cup \ldots \cup L_{K_n}$, we get the following equation explicitly relating HOMFLY to topological string partition functions:

$$\Psi_{K_1 \ldots K_n}(x_1, \ldots, x_n) = Z_{GW}(X, L_{K_1}, \ldots, L_{K_n})/Z_{GW}(X).$$

2.2 Higher representations and multiple branes

In Section 2.1, we considered knots and links with a single brane on the conormal of each component. It is natural to ask what happens if we instead insert several branes on the conormals. As it turns out, this reduces to a special case of single brane insertions. We explain this in the case of a single component knot K; the case of many component links is then an immediate generalization.

Consider $X = T^*S^3$ with $n > 1$ branes on the Lagrangian conormal L_K of a knot K. Let V be the $n \times n$ matrix of holonomies on the branes, with eigenvalues $(e^{x_1}, \ldots, e^{x_n})$. The topological string partition function of the n branes

$$\Psi_K^{[n]}(x_1, \ldots, x_n).$$

has a contribution from worldsheets with at least one boundary component on one of the branes wrapping S^3 that can be computed by inserting

$$\det(1 \otimes 1 - V^{-1} \otimes U(K))^{-1}$$

in the Chern-Simons path integral, and which describes the effect of integrating out the bifundamental strings between L_K and the S^3, generalizing (2.1) to the case of more than one brane on L_K.

Consider instead n distinct Lagrangians that are copies of L_K separated from it by moduli corresponding to $\mathrm{Re}(x_i)$. More precisely, we take the conormal of the n-component link $\tilde{K} = K_\epsilon \cup \ldots \cup K_{n\epsilon}$ where $K_{j\epsilon}$ is the knot obtained by shifting K a distance $j\epsilon$, where $\epsilon > 0$ is very small, along the framing vector field of K used to define the quantum invariants (i.e. the expectation values $\langle \mathrm{Tr}_R\, U(K)\rangle$). Note that the topological string partition function for a single brane on each $L_{K_{j\epsilon}}$ has exactly the same contribution, as follows e.g. from the simple mathematical fact that

$$\det(1 \otimes 1 - V^{-1} \otimes U)^{-1} = \prod_i \det(1 - e^{-x_i} U)^{-1},$$

which then holds inside the expectation values as well. In fact the system with n branes on L_K is physically indistinguishable from the system with single branes on all components of the conormal $L_{\tilde{K}}$. The expression on the left is more naturally associated to the former, whereas the one on the right is more naturally associated to the latter.

Treating $\tilde{K}$ as a general link as considered in Section 2.1, disregarding the effects of the branes being very close, we get a corresponding string partition function

$$\Psi_{\tilde{K}}(x_1, \ldots, x_n)$$

given by the insertion in the Chern-Simons path integral described above. However, when computing the string partition for n coincident branes on L_K, it is natural to take these effects into account: contributions that are not captured by Chern-Simons theory correspond to worldsheets with no boundary component on S^3. For parallel branes these come from short strings connecting different branes; such strings only contribute nontrivially to the annulus diagram of the topological string with amplitude $\sum_{k>0} \frac{1}{n} e^{kx_i} e^{-kx_j}$

corresponding to strings connecting L_{K_i} to L_{K_j}, $i < j$. Exponentiating this contribution, we then find the following relation between the partition function for branes on L_K and the partition function of single branes on $L_{\tilde{K}}$, treated as a general link:

$$\Psi_K^{[n]}(x_1, \ldots, x_n) = \Psi_{\tilde{K}}(x_1, \ldots, x_n)\Delta(x_1, \ldots, x_n),$$

where

$$\Delta(x_1, \ldots, x_n) = \prod_{i<j}(e^{x_i} - e^{x_j}).$$

Note that $\Psi_K^{[n]}$ encodes the HOMFLY of the knot K colored by representations with n rows:

$$\Psi_K(x_1, \ldots, x_n) = \left\langle \det\left(1 \otimes 1 - V^{-1} \otimes U(K)\right)^{(-1)} \right\rangle \Delta(x),$$

using the expansion

$$\det\left(1 \otimes 1 - V^{-1} \otimes U\right)^{-1} = \sum_R \mathrm{Tr}_R\, U \ \ \mathrm{Tr}_R\, V^{-1},$$

where the sum ranges over all $SU(N)$ representations R with at most n rows.

2.3 Large N duality

It was conjectured in [14] that $SU(N)$ Chern-Simons theory on S^3, or equivalently the topological A-model string on $X = T^*S^3$ with N D-branes on S^3, has a dual description in terms of the topological A-model on the resolved conifold Y which is the total space of the bundle

$$\mathcal{O}(-1) \oplus \mathcal{O}(-1) \to \mathbb{C}P^1.$$

Note that X and Y both approach the conical symplectic manifold $\mathbb{R} \times U^*S^3$, which is topologically $\mathbb{R} \times S^2 \times S^3$, at infinity. At the apex of the cone sits an S^3 in X, while in Y, there sits a $\mathbb{C}P^1 \approx S^2$. The duality arises from the geometric transition from X to Y that shrinks S^3 and replaces it by $\mathbb{C}P^1$ without altering the geometry at infinity. Furthermore, the N branes on S^3 in X disappear under the transition, but their number is related to the size (symplectic area) t of $\mathbb{C}P^1$ in Y as follows: $t = Ng_s$, where g_s is the string coupling constant. We write

$$Q = \exp(-t) = q^N. \tag{2.2}$$

The partition function of the closed topological string on Y counts holomorphic maps into Y. All such maps arise from perturbation of branched covers of the central $\mathbb{C}P^1 \subset Y$, and the variable Q keeps track of their degree. In [14], the partition function of the closed topological A-model string on Y was shown to agree with the partition function of $SU(N)$ Chern-Simons theory on S^3, and consequently with the partition function of A-model topological string on X in the background of N D-branes on S^3, provided that the Chern-Simons parameters, k and N, and the string coupling constant g_s (for topological strings in both X and Y) are related as $g_s = \frac{2\pi i}{k+N}$, and that (2.2) holds. In other words, for parameters related as described:

$$Z_{CS}(S^3; k, N) = Z_{GW}(X; g_s, N) = Z_{GW}(Y; g_s, Q).$$

We next describe how to include knots and links in this picture, see [13]. Let $K = K_1 \cup \cdots \cup K_n \subset S^3$ be a link. As described in Section 2.1, adding branes along the Lagrangian conormal L_K, we relate the Chern-Simons path integral with insertions corresponding to the link with open topological string on X with boundaries on either the N branes on S^3 or on the branes on L_K.

Recall that $L_K \subset X$ can be pushed off of the zero section S^3, corresponding to turning on $x_j \neq 0$ along L_{K_j}. We thus assume that $L_K \subset X$ is disjoint from $S^3 \subset X$ and consider the effect of the geometric transition from X to Y. Since the transition affects only a small neighborhood of the tip of the cone, corresponding to small neighborhoods of $S^3 \subset X$ and $\mathbb{C}P^1 \subset Y$, the Lagrangian L_K is canonically pushed through the transition as a Lagrangian in Y. Furthermore, in analogy with the closed string case discussed above, boundary conditions corresponding to worldsheet boundaries on the N branes on S^3 close up and disappear, while worldsheet boundaries on the branes on L_K remain unchanged. Thus, the partition function of branes on the components of L_K in Y

$$\Psi_{K_1,\dots,K_n}(x_1,\dots,x_n) = Z_{GW}(Y, L_{K_1},\dots,L_{K_n})/Z_{GW}(Y)$$

also gives the partition function of branes in X, wrapping the Lagrangian in X corresponding to $L_K \subset Y$ under transition, and N branes on $S^3 \subset X$, provided (2.2) holds.

2.4 SYZ mirror symmetry for knots

Consider the A-model topological string on the resolved conifold Y with a D-brane wrapping the Lagrangian conormal $L_K \subset Y$ of a knot $K \subset S^3$, see Section 2.3. There is then a contribution from short strings beginning and ending on L_K. Noting that a small neighborhood of L_K is symplectomorphic to a neighborhood of the zero section in T^*L_K and applying the construction of [12], we find that this contribution is given by the partition function of $U(1)$ Chern-Simons theory on $L_K \approx S^1 \times \mathbb{R}^2$. At infinity, L_K looks like $\mathbb{R} \times \Lambda_K$, with ideal boundary $\Lambda_K \approx T^2$, which means we should study $U(1)$ (or more precisely, $GL(1)$) Chern-Simons theory on the manifold $S^1 \times D^2$ with T^2 boundary. Let λ be the longitudinal cycle of the T^2, which determines the parallel of the knot K that links it trivially, and let μ be the meridional cycle. We denote by x the holonomy along λ, and by p the holonomy along μ:

$$\oint_\lambda A = x, \qquad \oint_\mu A = p,$$

where A is the connection 1-form. In Chern-Simons theory on L_K, the holonomies of λ and μ are canonically conjugate:

$$[p, x] = g_s,$$

where g_s is the string coupling constant. This means in particular that for a D-brane on L_K in Y, p and x are not independent; rather, if $\Psi_K(x)$ denotes the partition function of L_K, we have

$$p\,\Psi_K(x) = g_s \frac{\partial}{\partial x} \Psi_K(x).$$

In particular, we should view $\Psi_K(x)$ as a wave function with asymptotics

$$\Psi_K(x) \sim \exp\left(\frac{1}{g_s} \int p\,dx + \dots \right).$$

From the perspective of topological strings, in the classical limit $g_s \to 0$, maps of the disks dominate the perturbation expansion, with maps of more complicated Riemann surfaces giving contributions of higher order. This means that

$$\Psi_K(x) \sim \exp\left(\frac{1}{g_s} W_K(x, Q) + \dots \right),$$

where

$$W_K(x, Q) = \sum_{n,k \geq 0} a_{n,k}\, Q^k\, e^{-nx}$$

is the generating function of Gromov-Witten invariants $a_{n,k}$ corresponding to counting holomorphic disks in Y with boundary on L_K. Thus, at the level of the disk,

$$p = p(x) = \frac{\partial W_K}{\partial x}(x, Q). \tag{2.3}$$

We point out that (2.3) is consistent with the fact that, classically in L_K, the cycle μ bounds. Here μ bounding in L_K means that the holonomy around μ vanishes if we disregard the disk instanton corrections: $\frac{\partial W_K}{\partial x}$ vanishes up to contributions of instantons. This leads to the interpretation of the moduli space of the brane as a Lagrangian curve $V_K \subset T^*T^2$, with coordinates x, p and symplectic form $dx \wedge dp$, given by the equation

$$A_K(e^x, e^p, Q) = 0 \tag{2.4}$$

for a polynomial $A_K(e^x, e^p, Q)$ (where we view Q as a parameter). In general, for large x, (2.4) may have more than one solution $p = p(x)$ for p in terms of x. This corresponds to the fact that the theory may have more than critical point in this phase; to get the solution corresponding to L_K, we need to pick the one where $p \sim 0 + \mathcal{O}(e^{-x})$.

The curve (2.4) sums up disk instanton corrections to the moduli space of a D-brane on L_K probing Y. It was argued in [1] that this gives rise to a mirror Calabi-Yau manifold, X_K, given by the hypersurface

$$A_K(e^x, e^p, Q) = uv. \tag{2.5}$$

Mirror symmetry exchanges A-branes, i.e. Lagrangian submanifolds, of Y with B-branes, i.e. holomorphic submanifolds, of the mirror X_K. Moreover, quantum mirror symmetry implies that for every (special) Lagrangian brane in Y there is a B-brane in X_K such that the quantum corrected moduli space of L_K in Y is the same as the classical moduli space of the mirror B-brane in X_K. Here we see that the quantum dual B-branes are given by the line $\{u = 0, v \text{ arbitrary}\}$ over a point $q \in V_K$, i.e. if $q = (e^{x_0}, e^{p_0})$ then (2.4) holds for $(e^x, e^p) = (e^{x_0}, e^{p_0})$.

In the special case when K is the unknot, we obtain the same mirror of Y as in [15], but for more general knots, we get new mirrors. In fact, for every knot K in S^3 and its associated Lagrangian L_K in Y, we get a canonical mirror geometry (2.5), where the mirror of L_K is determined by a point on the mirror Riemann surface V_K. Of course, we expect each of the mirror Calabi-Yau manifolds (2.5) to contain not just the mirror of the Lagrangian L_K corresponding to the knot K, but also, by mirror symmetry, mirrors of all other Lagrangian branes $L_{K'}$ for knots $K' \neq K$. In general, however, these mirrors have more complicated descriptions.

3 Chern-Simons theory and Lagrangian fillings

Via large N duality, the Chern-Simons path integral on S^3 contains non-perturbative information about the A-model topological string in the resolved conifold Y. Before the transition, in $X = T^*S^3$, the data that go into specifying the theory are just the knot K, the number of branes on it, and the holonomy at infinity. We will set the number of branes on K to be 1 in this section, and consider the higher rank generalization in the next section. The data are imprinted at infinity of the ambient Calabi-Yau manifold, and are thus visible both before and after the transition. After the transition, the data at infinity may be compatible with more than one filling in the interior. In Section 2.4, we focused on the filling that gives the Lagrangian conormal L_K, but, as we shall see, Chern-Simons theory encodes information about other Lagrangian fillings as well; these correspond to different classical solutions of the topological string.

At infinity, the Calabi-Yau looks like $\mathbb{R} \times U^*S^3 \approx \mathbb{R} \times S^3 \times S^2$ and the Lagrangian L_K approaches $\mathbb{R} \times \Lambda_K$, where Λ_K, the conormal of the knot, is naturally identified with the boundary of a tubular neighborhood of K in S^3 and is thus topologically a torus [2]. There are two natural basic 1-cycles in Λ_K: the longitude λ, which is a parallel copy of the knot, and the meridian μ, which is the boundary of a fiber disk in the tubular neighborhood.

A vacuum of the topological string with these boundary conditions is determined once we specify the holonomy of the $GL(1)$ connection around a 1-cycle of the torus Λ_K at infinity and find a Lagrangian brane that fills in $\mathbb{R} \times \Lambda_K$ in Y with these moduli. Note that the holonomy of the $GL(1)$ connection encodes both the position of the brane and the $U(1)$-holonomy on it, see Section 2.4.[3] In the discussion so far we used the filling L_K and fixed the holonomy around the longitude cycle x. With a fixed connection at infinity, there may be more than one corresponding filling. Moreover, having found a vacuum with one choice of the filling, there is an $SL(2, \mathbb{Z})$ family of choices of flat connections, related by canonical transformations, that gives rise to the same filling. This corresponds to a choice of the framing of the Lagrangian, which we will discuss in more detail in Section 3.2.

---

[2] As will be further discussed in Section 5, $U^*S^3 \approx S^2 \times S^3$ is naturally a contact manifold, and Λ_K a Legendrian submanifold. These observations are the starting point for the knot contact homology approach to topological string in this background.

By taking the holonomy x around the longitude to be very large, we always get a choice of filling corresponding to the conormal Lagrangian L_K. The mirror Riemann surface

$$0 = A_K(e^x, e^p, Q), \tag{3.1}$$

which encodes the geometry of the mirror, also has information about fillings of $\mathbb{R} \times \Lambda_K$. As we explained, the Lagrangian L_K corresponds to a branch of the Riemann surface where

$$p = \frac{\partial W_K}{\partial x}(x) \sim 0.$$

Since p is the holonomy around the meridian of the knot, the fact that it vanishes classically means that the meridian cycle gets filled in in L_K, as is indeed the case topologically. The Lagrangian L_K has a one-dimensional moduli space, and the branch $p \sim 0$ is the branch where $x \to \infty$ and the worldsheet instanton corrections are maximally suppressed. As we vary x, we can probe all of the Riemann surface and go to a region where this is no longer the case.

---

[3] We cannot expect to specify holonomies around both cycles of Λ_K simultaneously, as the theory on the Lagrangian A-brane in a Calabi-Yau three-fold is Chern-Simons theory, and in Chern-Simons theory on a manifold with a T^2 boundary holonomies of the 1-cycles generating $H_1(T^2)$ are canonically conjugate.

There can be other ways of filling in. The region of p and x where the mirror geometry is well approximated by

$$mp + nx \sim \text{const} \tag{3.2}$$

should correspond to filling in which the $m\mu + n\lambda$ cycle in $\Lambda_K \approx T^2$ bounds. In all cases, "$\sim$" denotes "up to instanton corrections": if the approximation in (3.2) can be made arbitrarily good in the mirror, on the original Calabi-Yau Y, then the instanton correction can be made arbitrarily small, and the corresponding classical geometry exists in Y. The relation between the different phases should be akin to flop transitions in closed Gromov-Witten theory, in the sense that phase transitions change the topology of the cycles in the manifold: in this case, of the Lagrangian.

For example, filling in the longitude cycle λ gives

$$x \sim 0.$$

The Lagrangian M_K one obtains in this way is related to the knot complement $S^3 - K$. In the special case when the knot K is fibered, the complement can actually be constructed as a Lagrangian submanifold of Y asymptotic to $\mathbb{R} \times \Lambda_K$, see Section 6.4. In the general case we will give a conjectural construction of a Lagrangian filling with the same classical asymptotics, see Section 6.6. We will discuss this with further details in later sections. Here we simply let M_K denote a Lagrangian filling of Λ_K in Y that has the classical asymptotics $x \sim 0$. In this setting, the holonomy p around the meridian cycle is the natural parameter. In fact, the topological string partition function of Y with a single brane on M_K turns out to be

$$\tilde{\Psi}_K(p) = H_{p/g_s}(K), \tag{3.3}$$

computed by the HOMFLY polynomial $H_m(K)$ of the knot colored by the totally symmetric representation with m boxes. Here $m = p/g_s$ where p is fixed and g_s is small. This follows from existence of a geometric transition that relates M_K and L_K. Despite the fact that the transition changes the topology of the cycle that gets filled in, it is quantum mechanically completely smooth due to instanton corrections, in the case of knots. In later sections, we will study the case of links, where again there are different choices of filling in the cycle at infinity. However, there the phase structure of the theory becomes far more intricate.

3.1 Geometric transition between L_K and M_K

Let us clarify the relation between L_K and M_K and the nature of the transition between them. We start with N D-branes on the S^3 and a single brane on the conormal bundle L_K, intersecting along the knot K

$$L_K \cap S^3 = K.$$

The fact that L_K and S^3 intersect makes the configuration slightly singular, but one can remedy this by using the fact that L_K moves in a one-parameter family, parametrized by x. For any $\mathrm{Re}(x) \neq 0$, we can then move L_K off of S^3 so that they no longer intersect.

However, there is another way to smooth out the singularity by smoothing out the intersection between the S^3 and L_K. This requires breaking the gauge symmetry on S^3 by picking one of the N D-branes. Cut out a neighborhood of the knot K from both the D-brane on S^3 and on L_K. This gives $S^3 - K$ and $L_K - K$ both with a torus boundary. The two manifolds can then be glued together along their boundaries. Topologically, this results in $S^3 - K$, since gluing in $T^2 \times \mathbb{R}$ does not change the topology. (This way of obtaining $S^3 - K$ discussed, in a related context, in [16].) As mentioned above, it is not always possible to move the Lagrangian version of $S^3 - K$ off of the zero section, and we have a somewhat more involved construction with similar features to deal with this case. As in the previous section, we will use M_K

to denote a Lagrangian that classically gives $x \sim 0$.

Before the transition, the Lagrangian $S^3 - K$ and L_K project with different degrees to S^3 and no transition between them is possible. Also before the transition, N is finite, and hence $Q = e^{Ng_s} = 1$ classically. In this limit, the curve factorizes and contains

$$0 = A_K(e^x, e^p, Q = 1) = (1 - e^p) A_{K\ \text{class}}(e^x, e^p) \times \ldots \tag{3.4}$$

where $A_{K\ \text{class}}$ is the classical A-polynomial of the knot, describing the moduli space of flat connections on the knot complement. The $e^p - 1 = 0$ branch, corresponding to L_K, disconnects from the curve describing $S^3 - K$.

After the transition, the S^3 has shrunk, Q gets an expectation value $Q \neq 1$, the curve generally becomes irreducible

$$0 = A_K(e^x, e^p, Q), \tag{3.5}$$

and the distinction between L_K and M_K disappears. Since both $p \sim 0$ and $x \sim 0$ branches lie on the same Riemann surface, M_K and L_K are smoothly connected once we include disk instanton corrections, with no phase transition between them.

To make this point even clearer, it is useful to consider one very simple example of this when K is the unknot. We will study this both before and after the transition. In the mirror geometry L_K is given by choosing a point λ in the curve

$$Q - \lambda - \mu + \lambda\mu = 0 \tag{3.6}$$

where $\lambda = e^x$ and $\mu = e^p$. When $|\lambda| \gg 1$ this gives the Lagrangian mirror of L_K for the unknot. Note that this is consistent with $\mu = 1$, i.e., $p = 0$. The Lagrangian mirror to M_K is given by $|\mu| \gg 1$, which requires $\lambda = 1$, i.e., $x = 0$ as expected. Note that going from L_K to M_K can be done via a smooth path in the mirror curve, even though in the original geometry they are classically distinct.[4]

The transition between L_K and M_K described above has a simple interpretation in the topological string. At the intersection of the D-branes wrapping L_K and S^3 lives a pair of complex scalars A, B transforming in the bifundamental representation $(N,-1)$, $(\bar{N},1)$, corresponding to strings with one boundary on the zero section and the other on L_K. As long as L_K and S^3 intersect, the bifundamental is massless (the real mass vanishes). In one phase, we move the Lagrangian L_K off of the zero section. In this phase, the bifundamental hypermultiplet gets a mass $r = \mathrm{Re}(x)$ corresponding to the modulus of moving off. From this perspective, the partition function $\Psi_K(x)$ arises as follows. Integrating out the massive bifundamental of mass x generates the determinant [13] $\det(1 \otimes 1 - e^{-x} \otimes U)^{-1}$, where, as usual, $U = Pe^{i\oint_K A}$ is the holonomy along the knot and x is identified with holonomy on L_K. In Chern-Simons theory, we integrate over A and thereby compute the expectation value of the determinant $\Psi_K(x) = \langle \det(1 \otimes 1 - e^{-x} \otimes U)^{-1} \rangle_{S^3}$, which is in turn computed by the colored HOMFLY polynomial

$$\Psi_K(x) = \sum_{n=0}^{\infty} H_n(K)\, e^{-nx},$$

as explained earlier. This is in fact the derivation of the partition function given in [13].

The theory has a second phase, where the bifundamental hypermultiplet remains massless and gets an expectation value p. This requires the gauge group on the S^3 to be broken from $U(N)$ to $U(N-1) \times U(1)$, and then the $U(1)$ factors on the S^3 and on L_K to be identified. Giving the expectation value to the bifundamental gives rise to the smoothing of the intersection between L_K and S^3, which we described above. In the next subsection we will show that the partition function in this phase is given by (3.3).

---

[4] Note also that before the transition, to go from L_K to M_K, one must add one copy of S^3, as is clear from the path in the mirror geometry going from one asymptotic region to the other, taking into account that the periods get modified by one unit around the cycle mirror to the blown-up $\mathbb{C}P^1$.

3.2 Framing of the Lagrangians

To write down the partition function of the theory, in addition to choosing the vacuum by picking a point on the Riemann surface, we have to choose a framing of the Lagrangian: a flat connection at infinity. Since the holonomies around different cycles of the $\Lambda_K \approx T^2$ do not commute, different choices are related by wave-function transforms. In L_K, the natural variable is the holonomy along the longitude of the knot x, and the wave function is given by

$$\Psi_K(x) = \sum_{S_n} H_n(K)\; e^{-nx} = \exp\left(\frac{1}{g_s} W_K(x) + \dots\right).$$

In M_K, the natural variable is p, the holonomy around the meridian. Since x and p are canonically conjugate in the theory on the brane,

$$[p, x] = g_s,$$

and moreover since there is no distinction between M_K and the L_K phase, we can identify the Fourier transform of $\Psi_K(x)$ with the partition function of M_K:

$$\tilde{\Psi}_K(p) = \int_{-\infty}^{\infty} e^{\frac{1}{g_s}px}\Psi_K(x).$$

We can simply deform the contour, at least in perturbation theory, to find (3.3):

$$\tilde{\Psi}_K(p) = H_{p/g_s}(K) = \exp\left(\frac{1}{g_s} U_K(p) + \dots\right).$$

In particular, the Gromov-Witten “potentials” $W_K(x)$ and $U_K(p)$ that we previously associated to L_K and M_K are simply dual to each other,

$$U_K(p) \;=_{\text{crit}}\; px + W_K(x),$$

related by Legendre transformation. In Section 6.8, we will provide a different path to this result by directly comparing contributions of disk instantons to Gromov-Witten potentials of M_K and L_K.

4 Large N duality, mirror symmetry, and HOMFLY invariant for links

Consider Chern-Simons theory on S^3 with an n-component link K,

$$K = K_1 \cup \ldots \cup K_n,$$

where K_j are the knot components of K. The conormal L_K of K in $X = T^* S^3$ is the union of the n conormals of the knot components,

$$L_K = L_{K_1} \cup \cdots \cup L_{K_n},$$

and $L_K \cap S^3 = K$. Under large N transition as in Section 2.3, X is replaced by the total space Y of the bundle $\mathcal{O}(-1) \oplus \mathcal{O}(-1) \to \mathbb{C}P^1$. As we explained in Section 3.2, it is natural to view the boundary data as the only fixed data, in which case we should consider the geometry of L_K far from the apex of the cone. Here the Lagrangian L_K approaches $\mathbb{R} \times \Lambda_K$, where Λ_K is the union of the n conormal tori Λ_{K_j} of the components of K that describe the imprint of the corresponding branes at the $S^2 \times S^3$ at infinity of the Calabi-Yau. As in the knot case, there are different ways to fill in Λ_K in the interior of Y. As we shall see, the link case is more involved than the case of a single knot: there is a larger number of ways to fill in $\mathbb{R} \times \Lambda_K$, connecting different numbers of components of Λ_K in the interior.

As in Section 2.4, we are led to consider holonomies of a $GL(1)$ gauge field on the brane around the 1-cycles of Λ_K. This gives a phase space of the system, which is the cotangent bundle of a $2n$-dimensional torus:

$$\mathcal{M}_n = (\mathbb{C}^*)^{2n} = T^*(T^{2n}).$$

The torus is the torus of holonomies of the $U(1)$ gauge field on the brane around the $2n$ 1-cycles of Λ_K. The cotangent direction arises from the moduli of the Lagrangians. Equivalently, the phase space $\mathcal{M}_n$ arises from holonomies of the $GL(1)$ gauge field on $\mathbb{R} \times \Lambda_K$ around the cycles of n copies of T^2 that comprise the infinity Λ_K. Thus on $\mathcal{M}_n$ we have $(\mathbb{C}^*)^{2n}$-coordinates

$$e^{x_i}, \quad e^{p_j}, \qquad i, j = 1, \dots, n$$

associated to the longitudes and the meridians of knots K_i, and a symplectic form

$$\omega = \sum_i dx_i \wedge dp_i.$$

Note that $\mathcal{M}_n$ is in fact hyper-Kähler, a fact which we will make use of later in this paper.

Inside $\mathcal{M}_n$ there is a Lagrangian subvariety V_K associated to the link (i.e. V_K is half-dimensional and $\omega|_{V_K} = 0$). As we shall see, V_K is in general reducible:

$$V_K = \bigcup_P V_K(P),$$

where the subvarieties correspond to different fillings of Λ_K labeled by partitions P of $\{1, \dots, n\}$. We want to identify the fillings that can be obtained by smoothly varying the holonomies, since these give rise to the same Lagrangian submanifold $V_K(P)$. We will initiate the study of the varieties in this section, using Chern-Simons theory and large N duality. In Sections 5 and 6 we will use another approach based on knot contact homology, where V_K will be identified with the augmentation variety of the knot. In Section 4.4, we will study quantum mirror symmetry, where we explain how to quantize this variety.

4.1 Conormal bundle L_K

In the simplest case, $\mathbb{R} \times \Lambda_K$ is filled by the conormal bundle of the link

$$L_K = L_{K_1} \cup \ldots \cup L_{K_n}.$$

This is simply a union of n disconnected Lagrangians L_{K_j}, which we move off of S^3 independently and then transition to Y. This was studied previously in [17].

On Y, in the presence of L_K, there are no holomorphic disks with boundary on more than one component of L_K, as the conormals L_{K_j} are disjoint from each other. So the disk amplitude is simply a sum of k contributions, coming from disks with boundaries on one of the L_{K_j}:

$$W_{K,1^n}(x_1,\ldots,x_n) = W_{K_1}(x_1) + \cdots + W_{K_n}(x_n). \tag{4.1}$$

From the perspective of large N dualities, this arises as follows. We start with $X = T^*S^3$, with the n Lagrangians L_{K_j} intersecting the zero section along the link K. Pushing L_{K_j} off of S^3 by x_j for each j, we give a mass x_j to the bifundamental hypermultiplet at the intersection of K_j and S^3. Integrating these k massive hypermultiplets out, we end up computing the expectation value of

$$\begin{aligned}&\prod_{i=1}^n \det(1\otimes 1 - e^{-x_i}\otimes U(K_i))^{-1}\\ &= \sum_{m_1,\ldots,m_n} \mathrm{Tr}_{m_1}U(K_1)\ldots\mathrm{Tr}_{m_n}U(K_n)\; e^{-m_1x_1-\cdots-m_nx_n}\end{aligned}$$

where m_i denotes the m_i-th symmetric representation of the fundamental representation, in the topological string or Chern-Simons theory. This can be written as

$$\Psi_{K,1^n}(x_1,\ldots,x_n,g_s,N) = \sum_{m_1,\ldots,m_n} H_{S_{m_1},\ldots,S_{m_n}}(K)\; e^{-m_1x_1-\cdots-m_nx_n}, \tag{4.2}$$

where

$$H_{m_1,\ldots,m_n}(K) = \langle \mathrm{Tr}_{m_1}U(K_1)\ldots\mathrm{Tr}_{m_n}U(K_n)\rangle. \tag{4.3}$$

As was shown in [17], the classical limit of this is

$$\Psi_{K,1^n}(x_1,\ldots,x_n,g_s,Q) = \exp\left(\frac{1}{g_s}W_{K,1^n}(x_1,\ldots,x_n) + \ldots\right),$$

where $W_{K,1^n}$ is as in (4.1). In this case, the link indeed gives rise to a n-dimensional variety, but one which is simply a direct product of n 1-dimensional curves,

$$V_K(1^n): \qquad p_i = \frac{\partial W_{K_i}}{\partial x_i}(x_i).$$

We will denote this variety by $V_K(1^n)$, to indicate that it factors as a product of n pieces, each of dimension 1.

4.2 Link-complement-like fillings M_K

Just like in the knot case, there are other fillings of $\mathbb{R} \times \Lambda_K$ in Y. Here we focus on fillings that in the classical limit looks like the link complement. We denote a general such filling M_K, see Section 6.6 for geometric constructions of such fillings.

Physically, M_K is obtained by giving expectation values $p_i = m_i g_s$ to the n hypermultiplets at the intersections of L_{K_i} and S^3 in $X = T^*S^3$. This smooths the intersections of the D-branes on L_{K_i} with one of the D-branes on the S^3, to give a single Lagrangian M_K. This Lagrangian has first Betti number $b_1(M_K) = n$ so we still have n moduli that allow us to move M_K off of the zero section. These n moduli are the p_i's, which are also the holonomies around the meridians. After that, the geometric transition relates the Lagrangian M_K to a copy of itself on Y.

The disk amplitude depends on n meridian holonomies p_i, and is obtained from the colored HOMFLY polynomial of the link,

$$H_{m_1,\dots,m_n}(K) = \langle \mathrm{Tr}_{m_1} U(K_1) \dots \mathrm{Tr}_{m_n} U(K_n) \rangle, \tag{4.4}$$

by rewriting it in terms of $p_i = m_i g_s$:

$$\begin{aligned} \tilde{\Psi}_{K,n}(p_1, \dots, p_n, g_s, Q) &= H_{p_1/g_s,\dots,p_n/g_s}(K, g_s, Q) \\ &= \exp\left(\frac{1}{g_s} U_K(p_1, \dots p_n, Q) + \dots\right). \end{aligned}$$

Equivalently, we can write this as

$$U_K(p_1, \ldots p_n, Q) = \lim_{g_s \to 0} [g_s \log(H_{p_1/g_s, \ldots, p_n/g_s}(K, g_s, Q))].$$

The quantum corrected moduli space of the brane on M_K in Y is an n-dimensional variety $V_K(n)$ given by

$$V_K(n): \qquad x_i = \frac{\partial U_K}{\partial p_i}(p_1, \ldots, p_n, Q). \tag{4.5}$$

By the above construction, $V_K(n)$ can be viewed as a Lagrangian subspace in the $2n$-dimensional space of $(x_i, p_i) \in (\mathbb{C})^{2n}$ relative to the symplectic form $\sum_i dx_i \wedge dp_i$.

Note that if the link K is a completely split link (i.e., one can find disjoint solid balls $B_1, \ldots, B_n \subset S^3$ such that $K_i \subset B_i$ for each i), then in fact $V_K(n)$ and $V_K(1^n)$ coincide. This is because the HOMFLY polynomial of the link factor in this case,

$$H_{m_1,\ldots,m_n}(K) =_{K \text{ split}} H_{m_1}(K_1) \ldots H_{m_n}(K_n), \tag{4.6}$$

and consequently the brane partition function, is an exact product of $V_{K_i}(1)$. This suggests that in this case, one can go smoothly from the phase where we have a single Lagrangian M_K, to the conormal Lagrangian L_K consisting of n disconnected Lagrangians. Since $V_K(n)$ would have no new information in this case, we will simply identify it with $V_K(1^n)$ and reserve the notation of $V_K(n)$ for a non-split link K of n components: the fact that the link is non-split translates into the fact that $V_K(n)$ is not a direct product of some lower-dimensional varieties.

Just as in the case of the knots it is natural to expect $V_K(n)$ to correspond to moduli space of $SL(2, \mathbb{C})$ flat connections on $S^3 - K$. More precisely, we expect this to be a Q-deformed version of it, where $SL(2, \mathbb{C})$ is embedded in the canonical way in $SL(N, \mathbb{C})$. This is natural from the viewpoint of the Higgsing construction discussed before.

4.3 General fillings

We will now describe the set of physically distinct ways of filling in the torus Λ_K in the interior of Y. As we will explain, the physically distinct fillings are labeled by the ways to partition an n-component link K into sublinks. The different ways to do this are labeled by integer partitions P of $\{1, \ldots, n\}$, where we view the parts P_a of the partition P as giving corresponding sublinks K_a of K. An alternative way to label the fillings is by a link $K(P)$ which is a split union of sublinks K_a:

$$K(P) = \bigcup_{P_a \subset P} K_a.$$

(That is, move each of the sublinks away from the others, so that the result is split.) We conjecture that the general distinct fillings are labeled by the *primitive* partitions P, namely those for which each sublink K_a is non-split, or, equivalently, those that have no refinement P' such that $K(P) = K(P')$. The set of primitive partitions of an n-component link K is a subset of the set of all partitions of $\{1, \ldots, n\}$. For example, for a fully split link, there is a single primitive partition $P = \{\{1\}, \{2\}, \ldots, \{n\}\}$, and a single distinct partition of K into sublinks: $K = K(P)$.

Let us now explain why this is reasonable. The Lagrangian fillings of Λ_K, taken crudely, describe how components of Λ_K come together in the interior of Y. What one misses in this description is any information about which cycles get filled in; however, as we have seen, this fact is not invariant (one should recall the knot case, where phases with different cycles filled in get smoothly connected, due to instanton corrections). We want to identify those fillings which, while they may result in distinct classical geometries of the Lagrangian, are indistinguishable once disk instanton corrections are taken into account. Since one can interpolate smoothly between such Lagrangians, they lead to the same mirror variety. If this were the only consideration, we would simply label the fillings by all partitions P. However, we also have to take linking information into account. If some sublinks are split, since the amplitudes should not depend on where the corresponding Lagrangians are, we can separate them infinitely far from one another, which suggests there could not have existed a single Lagrangian with those asymptotics. Even if such Lagrangians existed, one could smoothly interpolate between the would-be connected ones and the disconnected phases of the Lagrangian by BRST trivial deformation as we separate them infinitely apart. Thus, partitions P that are not primitive contain only redundant information. Thus, for an n-component link K, the physically inequivalent fillings are labeled by primitive partitions P of $\{1, \ldots, n\}$, the ones that result in sublinks K_a that do not split further.

Given such a primitive partition P and the associated collection of sublinks $K(P)$, we will denote the corresponding Lagrangian filling by

$$F_{K(P)} \qquad \subset \qquad Y.$$

The filling $F_{K(P)}$ gives rise to an n-dimensional Lagrangian submanifold

$$V_K(P) = \prod_{P_a \subset P} V_{K_a} \subset \mathcal{M}_n,$$

which is a product of the varieties V_{K_a} associated to the sublinks K_a of $K(P)$. Each sublink K_a is by assumption non-split; suppose that K_a consists of n_a knot components, where $n = \sum_a n_a$. We now explain the meaning of $V_K(P)$.

Consider the disk amplitude

$$U_{K(P)}(p_1, \ldots, p_n)$$

obtained by studying the topological A-model string on Y, together with the Lagrangian brane $F_{K(P)}$. We define the “D-mirror” variety $V_K(P)$ as

$$V_K(P): \qquad x_i = \frac{\partial U_{K(P)}}{\partial p_i}(p_1, \ldots, p_k).$$

We have written $U_{K(P)}$ in terms of the meridian variables, which is a convenient choice when the meridian cycles get filled in. Since there cannot be any disk instantons ending on disconnected components of a Lagrangian filling, the disk potential is a sum

$$U_{K(P)} = \sum_a U_{K_a}.$$

Each disconnected component M_{K_a} of the filling is homologically like the complement of the link K_a in the classical limit, and U_{K_a} is the corresponding potential. This together with the fact that the potentials U_{K_a} give rise to varieties $V_{K_a}(n_a)$ gives the result we claimed. This is consistent with what we expect from large N duality: since the link $K(P)$ is a split link, with K_a its non-split sublinks, the corresponding potential should be captured by the classical limit of the HOMFLY polynomial of $K(P)$:

$$U_{K(P)}(p_1, \dots p_{n_a}) = \lim_{g_s \to 0} g_s \log(H_{p_1/g_s, \dots, p_{n_a}/g_s}(K(P))).$$

It is important to note that the HOMFLY polynomial $H_{K(P)}$ of the link $K(P)$ is *not* the full topological string amplitude for this filling; P and $K(P)$ only label the filling. The branes in the filling labeled by P do remember the full data of the link K they came from. The partitions P for which we get nontrivial fillings correspond to saddle (critical) points of the HOMFLY polynomial H_K of the link, and the full partition function of topological string with the Lagrangian branes on the filling $F_{K(P)}$ is encoded by the expansion of H_K, the HOMFLY polynomial of the original link K, around its saddle point P. We will discuss this in more detail below.

To summarize, for each different splitting of K into non-split sublinks K_a, labeled by the primitive partition P, we get a Lagrangian filling $F_{K(P)}$ of the Legendrian tori Λ_K at infinity. This filling is a union of link-complement-like Lagrangians M_{K_a}, one for each sublink of $K(P)$. The mirror variety corresponding to this filling is $V_K(P)$, which is a product of the corresponding varieties $V_{K_a}(n_a)$.

4.4 D-mirror variety

Given the link K, it is natural to define the mirror variety so that it depends on the data at infinity of the Calabi-Yau alone, and not the specific filling. With this in mind, we propose to identify

$$V_K = \bigcup_P V_K(P)$$

as the mirror variety to the link K. This is a natural proposal as it contains information both about knots that comprise the link, and the way they are linked together. Moreover, this is exactly the same data that goes into defining the dual Chern-Simons partition function on S^3 with the link K. In Section 8, we propose a way to quantize this variety.

Information about all the fillings P is contained in the HOMFLY polynomial of the link. In particular, quantum mechanically, $V_K(P)$ are the Lagrangian submanifolds of $\mathcal{M}_n$ that are associated to different classical saddle points of the wave function[5]

$$\Psi_K(x_1, \ldots, x_n) = \sum_{m_1,\ldots,m_n} H_{m_1,\ldots,m_n}(K) e^{-m_1 x_1 - \ldots - m_n x_n}.$$

In general, different saddle points contribute to Ψ_K:

$$\Psi_K(x_1, \ldots, x_n) = \sum_P c_P \Psi_K^P(x_1, \ldots, x_n),$$

where Ψ_K^P is a wave function canonically associated to the corresponding saddle point P, and the coefficients c_P are integers. Which saddle point dominates the Chern-Simons path integral depends on the values of the parameters x_i, g_s, and N, and the classical action at the saddle point

$$\Psi_K^P(x_1, \ldots, x_n) = \exp\left(\frac{1}{g_s} W_K(P)(x_1, \ldots, x_n) + \ldots\right),$$

---

[5] Non-perturbative aspects of topological strings have been studied in [18, 19, 20].

where $W_K(P)$ is the potential defining the variety $V_K(P)$ via

$$V_K(P): \qquad p_i = \frac{W_K(P)}{\partial x_i}.$$

In the regime where one of the saddle points dominates the others, are exponentially suppressed.

The above way of writing Ψ_K is a little bit crude because it neglects the fact that in general, a single P may give rise to more than one wave function, as $V_K(P)$ may give rise to several vacua. We will disregard this fact for two reasons: notational simplicity, and the fact that it is $V_K(P)$ that plays the crucial role in quantizing the theory. (While the subtlety affects the possible $W_K(P)$, it does not affect $V_K(P)$.)

The existence of a single function $\Psi_K(x_1, \ldots, x_n)$ that has all these different classical limits imposes constraints on the structure of V_K, viewed as a reducible variety. As we will discuss in Section 8, quantization of the variety V_K gives rise to a D-module on $\mathcal{M}_n$. The wave function $\Psi_K(x_1, \ldots, x_n)$ is a section of this D-module. The fact that the D-module that arises in this way is irreducible (it corresponds to a single irreducible quantum system, rather than being a direct sum of super-selection sectors) implies that different components $V_K(P)$ of V_K must fit together in a specific way, encoded by a graph that we now describe.

To every link K we can associate the graph Γ_K as follows. The vertices of the graph correspond to the primitive partitions P of $\{1, \ldots, n\}$, corresponding to the ways of partitioning K into non-split sublinks. The set of all such partitions is a subset of partitions of the set $\{1, \ldots, n\}$. Two vertices P and P' of the graph are connected by an edge if P' is a refinement of P and there is no other graph vertex P'' between the two (where P'' is a refinement of P and P' is a refinement of P'').

We claim that the graph Γ_K obtained in this way captures the geometry of V_K as follows: for any pair of vertices of the graph that are connected by an edge, we conjecture that

$$\operatorname{codim}(V_K(P) \cap V_K(P')) = 1,$$

where codimension is counted inside either of the n-dimensional varieties $V_K(P), V_K(P')$. Note that, based on counting dimensions, a generic intersection would be over points (codimension n). The presence of an edge means that the intersection of $V_K(P)$ and $V_K(P')$ is as non-generic as possible without them coinciding. In a generic situation thus the graph would consist of vertices alone. The existence of a single function Ψ_K implies that this is never the case for multi-component links. Furthermore, it is natural to generalize our conjecture as follows: for any two primitive partitions P and P', the codimension of the intersection of $V_K(P)$ and $V_K(P')$ is at most the distance between the partitions, defined as the minimum number of edges in Γ_K needed to pass from P to P':

$$\operatorname{codim}(V_P \cap V_{P'}) \leq d(P, P').$$

If all partitions are primitive and thus give vertices of the graph, then we can connect the coarsest partition[6] $P_c = (12 \cdots n)$ with the finest partition $P_f = (1) \cdots (n)$ by a sequence of partitions $P_c = P_1, P_2, \ldots, P_n = P_f$, where $P_i = (1) \cdots (i-1)(i \cdots n)$. Passing successively from P_1 to P_n involves removing the knot components of K one by one, starting with the first component, until we have removed all of them. We expect that for each i, $V_K(P_i)$ and $V_K(P_{i+1})$ intersect in codimension 1, and this allows us to move between components of V_K in such a way that each move is between components whose intersection has codimension 1. See Section 5.4 for further discussion for general links and a reformulation of this discussion in terms of the augmentation variety.

---

[6] For convenience, we will sometimes abbreviate $P_c = (12 \cdots n)$ by n and $P_f = (1) \cdots (n)$ by 1^n when writing $V_K(P_c)$ or $V_K(P_f)$.

Let us now explain why the irreducibility of the D-module implies the codimension condition. A way to characterize the fact that Ψ_K corresponds to an irreducible D-module is that, first of all, it satisfies a set of differential (or, more precisely, difference) equations generated by a finite set of operators

$$\mathcal{A}_\alpha \Psi_K = 0, \tag{4.7}$$

where each of $\mathcal{A}_\alpha$ is of the form

$$\mathcal{A}_\alpha = \sum_{k_1,\ldots,k_n} a_{k_1,\ldots,k_n}(e^{x_1},\ldots,e^{x_n}) e^{g_s \sum_i k_i \partial_{x_i}}.$$

where $a_{k_1,\ldots,k_n}$ are polynomials in e^{x_i}. The fact that the equation (4.7) is linear implies that not only Ψ_K but also each of Ψ_K^P satisfy the equation. It was proven in [21] that the colored HOMFLY polynomial of any link satisfies such a finite set of difference equations, and that the set is holonomic (or more precisely, q-holonomic). What this means is that, in the classical limit $g_s \to 0$,

$$\mathcal{A}_\alpha(e^{x_i}, e^{g_s \partial_j}) \qquad \longrightarrow \qquad A_\alpha(e^{x_i}, e^{p_j}),$$

when x_i, $p_j = g_s \partial_{x_j}$ become numbers, the set of equations (4.7) parametrize a Lagrangian submanifold of $\mathcal{M}_n$. This Lagrangian submanifold is just what we called V_K,

$$V_K : \qquad\qquad A_\alpha(e^{x_i}, e^{p_j}) = 0.$$

The fact that irreducible components of $V_K = \bigcup_P V_K(P)$ intersect over the varieties of codimension 1 is a consequence of a general theory of systems of differential equations of this type [22, 23, 24].[7]. We will discuss this point further in Section 8, where we discuss quantization of the system from the physical perspective.

---

[7] The system at hand is q-holonomic, rather than holonomic. The theorems we need are best developed in the ordinary holonomic case. However, we can view the q-holonomic system as holonomic, by working locally, so presumably the distinction is immaterial.

4.5 An example: Hopf link

The simplest nontrivial 2-component link is the Hopf link $K = \text{Hopf} = K_1 \cup K_2$, with $K_1, K_2 = \bigcirc$. In this case, we expect two distinct fillings, corresponding to partitions $P = (12)$ (which we will abbreviate as 2) and $P = (1)(2)$ (which we will abbreviate as 1^2). In the first case $P = (12)$, Lagrangian filling of Λ_{Hopf} is the complement of the entire Hopf link in S^3,

$$M_{\text{Hopf}} = S^3 - \text{Hopf},$$

and this gives rise to $V_{\text{Hopf}}(2)$. The filling is easy to see from the toric diagram, before and after the transition. The Hopf link is realized by taking two Lagrangian branes on opposite toric legs before the transition. They can each glue to one copy of the brane on S^3 to move off the S^3. The topology of the resulting Lagrangian is $T^2 \times \mathbb{R}$. This has $b_1 = 2$, so the moduli space is 2-dimensional. Moreover, there are no disk instanton corrections to it[8], so the potential is determined classically. It is easy to see that

$$U_{\text{Hopf}}(p_1, p_2) = p_1 p_2, \tag{4.8}$$

since the critical points associated to this simply state that the meridian of one knot in the Hopf link gets identified with the longitude of the other, once we glue M_K from pieces, as we described. Namely, in terms of

$$e^{p_i} = \mu_i, \qquad e^{x_i} = e^{\partial_{p_i} U_{\text{Hopf}}} = \lambda_i,$$

we have

$$V_{\text{Hopf}}(2): \qquad \lambda_1 = \mu_2, \qquad \lambda_2 = \mu_1.$$

The other filling corresponding to $P = (1)(2)$ has two branes L_{K_i}, each simply a conormal to the unknot, probing the conifold geometry after the transition. At the level of the disk potential, they cannot talk to each other, so

---

[8] This is the case because any holomorphic map with an S^1 boundary on the T^2 comes in an S^1 family. The fact that the Euler characteristic of the S^1 vanishes leads to vanishing of the corrections.

$$U_{\text{Hopf}}(1^2)(p_1, p_2) = U_{\bigcirc}(p_1) + U_{\bigcirc}(p_2),$$

leading to a direct product of two copies of Riemann surfaces mirror to the unknot:

$$V_{\text{Hopf}}(1^2): \qquad Q - \lambda_1 - \mu_1 + \lambda_1\mu_1 = 0, \qquad Q - \lambda_2 - \mu_2 + \lambda_1\mu_1 = 0.$$

The variety

$$V_{\text{Hopf}} = V_{\text{Hopf}}(1^2) \cup V_{\text{Hopf}}(2)$$

coincides with the augmentation variety of the Hopf link, as we will see in the next section.

The same result can be obtained in a number of different ways, two of which we mention here. First, the result can be obtained by calculating the HOMFLY polynomials of the Hopf link, colored by totally symmetric representations, and taking the classical limit. We will show this in Appendix A in detail. Second, as we will brush on in Section 8, one can obtain the same result by considering a pair of branes on the Riemann surface mirror to the unknot.

These two approaches allow one not only to recover the classical variety V_{Hopf}, but also give a prediction for its quantization. Namely,

$$\Psi_{\text{Hopf}}(x_1, x_2) = \sum_{m_1, m_2} H_{m_1, m_2} e^{m_1 x_1 + m_2 x_2}$$

is given in terms of the $SU(N)_n$ WZW S matrix

$$H_{m_1, m_2} = S_{m_1, m_2}/S_{00}$$

evaluated at $q = e^{2\pi i/(n+N)}$and $Q = q^N$. Here $S_{m_1 m_2}$ is the matrix element corresponding to the totally symmetric representations with m_1, m_2 boxes.

Using the well-known explicit expressions (see Appendix A), one can show that Ψ_{Hopf} satisfies a set of difference equations

$$\mathcal{A}_\alpha \Psi_{\text{Hopf}} = 0, \qquad \alpha = 1, 2, 3,$$

where

$$\begin{aligned}
\mathcal{A}_1 &= -e^{x_1} + e^{x_2} - (1 - Qe^{x_1})e^{g_s \partial_1} + (1 - Qe^{x_2})e^{g_s \partial_2}, \\
\mathcal{A}_2 &= (1 - q^{-1}e^{x_2} - (1 - Qe^{x_2})e^{g_s \partial_2})(e^{g_s \partial_1} - e^{x_2}), \\
\mathcal{A}_3 &= (1 - q^{-1}e^{x_1} - (1 - Qe^{x_1})e^{g_s \partial_1})(e^{g_s \partial_2} - e^{x_1}).
\end{aligned}$$

In the classical limit, these reduce to $A_\alpha = 0$, where

$$\begin{aligned}
A_1 &= -\lambda_1 + \lambda_2 - (1 - Q\lambda_1)\mu_1 + (1 - Q\lambda_2)\mu_2, \\
A_2 &= (1 - \lambda_2 - \mu_2 + Q\lambda_2\,\mu_2)(\mu_1 - \lambda_2), \\
A_3 &= (\mu_2 - \lambda_1)(1 - \lambda_1 - \mu_1 + Q\lambda_1\,\mu_1)(\mu_2 - \lambda_1).
\end{aligned}$$

The Lagrangian solutions to this are precisely the variety $V_{\text{Hopf}} = V_{\text{Hopf}}(1^2) \cup V_{\text{Hopf}}(2)$. (There are other solutions that do not lead to Lagrangians. These are not of interest, as they do not describe saddle points.) Finally, note that the intersection $V_{\text{Hopf}}(1^2) \cap V_{\text{Hopf}}(2)$ is indeed codimension 1: it is a curve, given by

$$\mu_2 = \lambda_1, \qquad \lambda_2 = \mu_1,$$

where μ_1, λ_1 lie on the unknot curve

$$1 - \lambda_1 - \mu_1 + Q\lambda_1\mu_1 = 0.$$

The equations are simple enough that we can solve them exactly. Two linearly independent solutions, with different asymptotics, can be obtained by considering two different contours of integration $\mathcal{C}_{(2)}$, $\mathcal{C}_{(1^2)}$ in

$$\Psi^P_{\text{Hopf}}(x_1, x_2) = \int_{\mathcal{C}_P} dp_2\, \frac{1}{1 - e^{-p_2}e^{x_1}} \Psi_{\bigcirc}(x_1)\Psi^{-1}_{\bigcirc}(p_2)\; e^{p_2 x_2 / g_s}$$

where $\Psi_{\bigcirc}(x)$ is the partition function of the unknot.[9] The wave function corresponding to $V_{\mathrm{Hopf}}(2)$ is obtained by taking a contour $\mathcal{C}_{(2)}$ which is a small circle $p_2 = x_1$; this gives

$$\Psi^{(2)}_{\mathrm{Hopf}}(x_1, x_2) = e^{x_1 x_2/g_s}.$$

The wave function corresponding to $V_{\mathrm{Hopf}}(1^2)$ can be obtained by taking the contour $\mathcal{C}_{(1^2)}$ to run along the real axis. Using the fact that Fourier transform $\Psi^{-1}_{\bigcirc}(p)$ is $\Psi_{\bigcirc}(x)$, this gives

$$\Psi^{(1^2)}_{\mathrm{Hopf}}(x_1, x_2) = \exp(W_{\bigcirc}(x_1)/g_s + W_{\bigcirc}(x_2)/g_s + \ldots).$$

4.6 Knot parallels and higher rank representations

Consider a link $\tilde{K}$ obtained by taking n parallels of a single knot K. This is closely related to studying n branes on a single Lagrangian associated to the knot K. As explained in Section 2.2, the partition functions are related by

$$Z_{\tilde{K}}(x_1, \ldots, x_n) = Z^{(n)}_K(x_1, \ldots, x_n)/\Delta(x),$$

where

$$\Delta(x) = \prod_{1 \le i < j \le n} (e^{(x_i - x_j)/2} - e^{(x_j - x_i)/2})$$

arises from integrating out short strings, with boundary on L_K alone and with no boundaries on the S^3. This is a sum over annuli and hence has no g_s dependence.

This relation between $Z_{\tilde{K}}$ and $Z^{(n)}_K$ implies that for every saddle point of $Z_{\tilde{K}}$ we get a saddle point of $Z^{(n)}_K$ as well, at least as long $\Delta(x)$ has no zeroes there. Thus we expect the saddle points of Z_K to be a subset of saddle points of $Z_{\tilde{K}}$. The ones that may be missing are those where short bifundamental strings get expectation values. These are missing in $Z^{(n)}_K(x_1, \ldots, x_n)$, where by assumption these strings are massive and we have integrated them out to get $\Delta(x)$. We will give a rigorous mathematical proof of this statement in Section 5, from the perspective of knot contact homology. It is satisfying that the same picture emerges from Chern-Simons theory and large N duality.

---

[9] We have simplified things slightly by shifts of variables. For details see Appendix A.

5 Knot contact homology

In this section we discuss the connection between knot contact homology and topological strings in the context of knot and link invariants. We first give a brief discussion of this connection in the next subsection, before giving a more detailed discussion in the following subsections and in Section 6.

5.1 Brief summary of relation between topological strings and knot contact homology

Knot contact homology uses the Lagrangian brane L_K in $X = T^*S^3$, but with the zero section $S^3 \subset T^*S^3$ deleted. In other words, one considers the geometry far from the apex of the cone where the geometry is $\mathbb{R} \times U^*S^3 \approx \mathbb{R} \times S^3 \times S^2$; this shares the basic feature of the geometry of Y after the transition, where S^2 does not bound.

For simplicity of notation, we temporarily assume that K is a single-component knot. Then L_K is $\mathbb{R} \times \Lambda_K \approx \mathbb{R} \times T^2$ and it is natural to view the $\mathbb{R}$-factor as the "time" direction. Here the torus Λ_K can be viewed as the geometry of the Lagrangian brane at infinity. The knot leaves its imprint on how the Legendrian submanifold Λ_K is embedded in the contact manifold $U^*S^3 \approx S^3 \times S^2$.

Consider the physical open string states ending on Λ_K. Among these curves with endpoints on Λ_K, the paths that are stationary for the action $\int p\,dq$ are "Reeb chords" a_i, which are string trajectories that are flow segments of the Reeb vector field. (Reeb chords also have the property that holomorphic maps can end on them.) In the physical setup we would say that the a_i are classically annihilated by the BRST symmetry Q_B:

$$Q_B \cdot a_i = 0.$$

However, disk instantons modify the operation of Q_B in a way similar to Witten's formulation of Morse theory [25], where the critical points of the Morse function correspond to vacua, but instanton corrections, which correspond to gradient flows (rigid up to translation), modify the supersymmetry algebra. In the case at hand the role of the gradient flows are played by disk instantons, which are disks D (rigid up to $\mathbb{R}$-translation) that at $t \to +\infty$ start with the Reeb chord a_i and at $t \to -\infty$ approach the Reeb chords $a_{j_1}, \ldots, a_{j_{k_D}}$ as ordered multi-pronged strips with $k_D + 1$ punctures. The boundary of the disk maps to $\mathbb{R} \times \Lambda_K$ and hence (upon choosing some fixed paths capping the ends of the Reeb chords) represents an element in $H_1(\Lambda_K)$, and thus picks up the holonomy factor $\exp(l_D x + m_D p)$ from the Wilson lines on the Lagrangian brane. Let n_D denote the intersection of this disk (again suitably capped off) with the 4-cycle dual to S^2, i.e. $\mathbb{R} \times S^3$. Then there is a deformed Q_B operator (called ∂ below), which we interpret as the quantum corrected Q_B, given by

$$Q_B \cdot a_i = \sum_{\text{rigid disks } D} Q^{n_D} e^{l_D x + m_D p} a_{j_1} \cdots a_{j_{k_D}},$$

where x, p correspond to the holonomies of the probe brane around the longitudinal/meridional directions of the knot. Thus the open string states correspond to nontrivial elements of the deformed Q_B cohomology. The existence of a 1-dimensional representation of this cohomology leads to an algebraic constraint $A(e^x, e^p, Q) = 0$. Such 1-dimensional representations can be interpreted as the disk corrected moduli space of one such Lagrangian brane after the large N transition; see Section 6.2, where it is explained how a Lagrangian filling with associated moduli space of disks gives an augmentation. Using the large N duality, this should be the same as the Q-deformed A-polynomial we defined earlier, which is the quantum corrected moduli space of the single brane L_K. The fact that we find agreement between these two setups can be interpreted as further evidence for large N duality.

The same setup applies to the case of links. The main novelty there is that we have n Legendrian tori and thus we have n^2 sectors of open strings associated to the n^2 choices for the endpoint components of the Reeb chords. As in the knot case, we can use disk instantons to find the Q_B corrected cohomology. The representation theory of the resulting algebra, where there are nontrivial open string states between all pairs, leads to an n-dimensional variety $V(n)$, which we identify with the corresponding moduli space $V_K(n)$ discussed earlier; see Section 6.6 for possible geometric constructions of corresponding Lagrangian fillings. Moreover, we can consider other representations where part of the open strings between some pairs of Legendrian tori are mapped to zero. This gives other varieties $V_K(P)$ as discussed in Section 4.

We now turn our attention to a more mathematical description of knot contact homology. We begin with a brief general introduction to contact homology, an object associated to contact manifolds and to their Legendrian submanifolds. Next we describe knot contact homology, which is the contact homology associated to the Legendrian torus $\Lambda_K \subset U^*\mathbb{R}^3$ (which can be viewed as a patch of $U^*S^3 = S^2 \times S^3$), first geometrically and then algebraically. We then present augmentations from an algebraic perspective and use these to define a polynomial knot invariant, the augmentation polynomial, or a variety in the case of a multi-component link. The augmentation variety is conjectured to agree with the mirror-symmetry variety V_K from previous sections. Besides agreeing with V_K in a number of examples, we will see that the augmentation variety shares many properties of V_K. This will be discussed further in Section 6.

5.2 Mathematical overview of knot contact homology

Here we provide a summary of the mathematics behind knot contact homology and augmentations. This goes into more detail than the previous subsection and also places the construction in the context of contact and symplectic geometry. The full story is rather long, and we will omit many technical details in the interest of readability. As a result, some of the discussion below is rather imprecise, but we will provide references to more detailed treatments in the literature for the interested reader. A slightly less brief overview can be found in [26].

Knot contact homology is a special case of *Legendrian contact homology*, which is itself a small part of the more elaborate Symplectic Field Theory package in symplectic topology introduced in [27]. The general setup begins with a contact manifold, which for our purposes is a $(2m-1)$-dimensional manifold V equipped with a 1-form α such that $\alpha \wedge (d\alpha)^{m-1}$ is a volume form on V. The 1-form α determines the Reeb vector field R_α on V given by $d\alpha(R_\alpha, \cdot) = 0$ and $\alpha(R_\alpha) = 1$. Associated to the contact manifold V is its symplectization, the $2m$-dimensional manifold $\mathbb{R} \times V$ equipped with the symplectic form $\omega = d(e^t \alpha)$, where t is the coordinate in the additional $\mathbb{R}$-factor; on the symplectization we may choose an $\mathbb{R}$-invariant almost complex structure J pairing $\partial/\partial t$ and the Reeb vector field R_α.

Legendrian contact homology is associated to the contact manifold V along with a *Legendrian submanifold* $\Lambda \subset V$, which is a submanifold along which α is identically 0, of maximal dimension; the contact condition on α forces this maximal dimension to be $m-1$. In this case, $\mathbb{R}\times\Lambda$ is a Lagrangian submanifold of $\mathbb{R} \times V$.

Define a *Reeb chord* of Λ to be a flowline of R_α that begins and ends on Λ. We define the algebra $\mathcal{A} = \mathcal{A}(\Lambda)$ to be the free (tensor) algebra over the ring $\mathbb{Z}[H_2(V, \Lambda)]$ generated by Reeb chords of Λ. That is, if $a_1, \ldots, a_r$ are the