Title
Developments in the mathematics of the A-model: constructing Calabi-Yau structures and stability conditions on target categories

Author
Takeda, Alex Atsushi

Publication Date
2019

Peer reviewed|Thesis/dissertation

Developments in the mathematics of the A-model: constructing Calabi-Yau structures and stability conditions on target categories

by

Alex Atsushi Takeda

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Vivek Shende, Co-chair
Professor Mina Aganagic, Co-chair
Professor Ori Ganor
Professor David Nadler

Summer 2019

Developments in the mathematics of the A-model: constructing Calabi-Yau structures and stability conditions on target categories

Abstract

Developments in the mathematics of the A-model: constructing Calabi-Yau structures and stability conditions on target categories

by

Alex Atsushi Takeda

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Vivek Shende, Co-chair

Professor Mina Aganagic, Co-chair

This dissertation is an exposition of the work conducted by the author in the later years of graduate school, when two main projects were completed. Both projects concern the application of sheaf-theoretic techniques to construct geometric structures on categories appearing in the mathematical description of the A-model, which are of interest to symplectic geometers and mathematicians working in mirror symmetry. This dissertation starts with an introduction to the aspects of the physics of mirror symmetry that will be needed for the exposition of the techniques and results of these two projects. The first project concerns the construction of Calabi-Yau structures on topological Fukaya categories, using the microlocal model of Nadler and others for these categories. The second project introduces and studies a similar local-to-global technique, this time used to construct Bridgeland stability conditions on Fukaya categories of marked surfaces, extending some results of Haiden, Katzarkov and Kontsevich on the relation between stability of Fukaya categories and geometry of holomorphic differentials.

To all the friends, family members and lovers that supported me along this journey.

Contents

List of Figures

Acknowledgments

I would like to especially thank my advisor Vivek Shende for being so willing in sharing his time, guidance and original ideas. I would also like to thank other members of UC Berkeley faculty such as my co-advisor Mina Aganagic, David Nadler, Ori Ganor and others for being so formative in my education here. I would like to thank as well Tom Bridgeland, Fabian Haiden, Chris Brav and Tobias Dyckerhoff for being so patient and clear in explaining to me some of their work, and Kyōji Saitō for very enjoyable lessons. Last but not least, I would like to thank my friends Dori Bejleri, Ilyas Bayramov, Benjamin Gammage, Tatsuki Kuwagaki, Michael Levy, Ryan Thorngren and Gjergji Zaimi, without whom mathematics would be no fun.

Chapter 1

Overview of mirror symmetry

1.1 Introduction

For the past 30 years, the study of mirror symmetry has been one a very active and fruitful front of the interaction between physics and mathematics. One can trace the lineage of mirror symmetry to the very fruitful study of QFT- and string-theoretic dualities of the 80s and 90s; the physical argument for mirror symmetry can be phrased as an application of T-duality in string theory to type II string theories compactified on Calabi-Yau threefolds.

Mirror symmetry provides a striking example of the power of physical arguments in mathematics, since it suggests a deep relation between different types of geometry that would not apparent for mathematicians working on either one of these fields. Perhaps the most classical and prominent example of this phenomenon is the pioneering work of Candelas, de la Ossa, Green and Parkes [26], where the authors conjecture a precise formula for the number of rational curves of a fixed degree in the quintic threefold $X \subset \mathbb{P}^4$, given in terms of periods of a different Calabi-Yau manifold (the mirror quintic). On the mathematical side, this formula was proven by the work of Lian, Liu and Yau [76].

For mathematicians, one of the important developments in the study of mirror symmetry is Kontsevich's homological mirror symmetry conjecture [71]. This conjecture relates the fields of symplectic geometry and complex geometry; roughly it can be interpreted to say that the symplectic geometry of a Calabi-Yau manifold is very closely related to the complex geometry of its mirror manifold.

One can assign certain categorical invariants to these manifolds, which should be non-trivially related for mirror pairs of manifolds. On the complex side, the category to be considered is the derived category of coherent sheaves, which is a very familiar object in algebraic geometry. On the symplectic side, the category to be considered is the Fukaya category, which first appeared in the work of Fukaya on Morse theory for Lagrangians. To a given symplectic manifold X, this construction assigns a category whose objects are given by Lagrangians in X, and whose morphisms come from intersection points; for the structure of this category to be made precise, one needs to qualify these statements properly.

This dissertation will focus on some recent mathematical developments on geometric structures on certain Fukaya categories. The results presented here are from my own work [100, 107] completed during my graduate studies in collaboration with my academic advisor. These papers address the construction of two kinds of geometric structures on Fukaya categories, namely *Calabi-Yau structures* and *Bridgeland stability conditions*. In this chapter we will review some of the background material on these topics, with an emphasis on the physics behind them; most of the material in here is based on the now-classical references [57, 8]

There are many different versions of Fukaya categories developed over the past decades, and to provide a comprehensive review of the field would be beyond the scope of this thesis; the interested reader can consult some more comprehensive monographs and reviews such as [46, 10]. Here we will limit ourselves to a schematic overview of some specific topics that will be important for the exposition of original research.

1.2 Mirror symmetry for Calabi-Yau manifolds

Generalities about SCFTs

Here we will follow the exposition in Kapustin and Orlov's 2003 lectures on mirror symmetry [63] and their work on this subject [65].

Mirror symmetry can be seen algebraically as a symmetry of the $\mathcal{N} = (2,2)$ superconformal algebra, without necessarily a relation to geometry. The full set of generators and relations for this superconformal algebra is quite complicated, and can be found eg. in [65] (where it is referred to as the $\mathcal{N} = 2$ super-Virasoro algebra). This algebra contains two copies $L_m, \bar{L}_m$ of the Virasoro generators, together with another set of even generators $J_n, \bar{J}_n$, and as odd generators, two copies of the set $\mathcal{N} = 2$ supersymmetry generators $Q_n^\pm, \bar{Q}_n^\pm$.

An $\mathcal{N} = (2,2)$ superconformal field theory (SCFT) is a 2d CFT endowed with an action of this algebra; such a theory assigns to the circle a space of states $\mathcal{H}$ which will have the structure of a module over the $\mathcal{N} = (2,2)$ superconformal algebra. One can understand this action by the fact that this algebra is a supersymmetric extension of the Virasoro algebra, which is itself a central extension of the algebra of vector fields on the circle. This algebra has an outer automorphism $\mathcal{M}$ which we will call the *mirror automorphism*

$$L_n \mapsto L_n, \quad J_n \mapsto -J_n, \quad Q_n^\pm \to Q_n^\mp$$

We will then say that two SCFTs $\mathcal{H}_1, \mathcal{H}_2$ are mirror if there is an isomorphism $\mathcal{H}_1 \leftrightarrow \mathcal{H}_2$ such that the mirror map intertwines the action of the superconformal algebra on the two sides.

The sigma model

The relation between this story to the phenomenon in Calabi-Yau geometry that interests mathematicians is given by the construction of the $\mathcal{N} = (2,2)$ sigma model. The data required to define such a model is a Calabi-Yau manifold equipped with a B-field, ie. a triple (X, g, B) where X is a compact complex manifold with trivial canonical class, g a Kähler metric on X and $B \in H^2(X,\mathbb{R})/H^2(X,\mathbb{Z})$ is the B-field.

This model [57, Ch. 13] can be defined classically by a Lagrangian that is a supersymmetric extension of a 2d sigma model; it is a theory of a map $\phi : \Sigma \to X$ of a surface Σ into the target space X, with action given by

$$S = \int_\Sigma \mathcal{L} = \int_\Sigma \left(-g_{ij}\partial^\mu \phi^i \partial_\mu \phi^j + \phi^* B + \text{ fermionic terms}\right)$$

ie. we use the map ϕ to pull back to Σ both the Kähler metric and the B-field.

It is believed that one can quantize this model preserving the $\mathcal{N} = (2, 2)$ symmetry to get an SCFT; this is known at least near the large volume limit. This gives a space of states $\mathcal{H}$ with an action of the superconformal algebra; this space also receives gradings from the actions of the even generators of the supeconformal algebra.

One can then ask what is the relation between this action and the geometry of the original Calabi-Yau manifold. The commutation relations above imply, for example, that the BRST operator $D = Q_0 + \bar{Q}_0$ squares to zero. This is interpreted as a differential on the space $\mathcal{H}$, and one can argue that the cohomology of D recovers the Hodge theory of the original manifold X; this cohomology gets two gradings from the generators $J_0, \bar{J}_0$ and therefore at degrees (p, q) one gets a cohomology group, and its dimension is given by the Hodge number $h^{p,q}(X)$.

If two different Calabi-Yau manifolds X, X' give mirror SCFTs in the sense we mentioned above, one can follow the gradings to deduce that we must have the following relation between the Hodge numbers

$$h^{p,q}(X) = h^{n-p,q}(X')$$

which, when plotted in the Hodge diamond, gives some justification for the name 'mirror symmetry'.

Topological twisting

The picture of mirror symmetry in terms of $\mathcal{N} = (2, 2)$ SCFTs is physically compelling, but mathematically quite hard to make sense of, since these theories are not yet fully axiomatized into mathematics. There is, however, another picture of mirror symmetry that captures its interesting phenomena, and whose two sides can be defined mathematically. This is obtained by a procedure called topological twisting, first proposed by Witten in [113]. Here we present a concise description of the two twists of the $\mathcal{N} = (2, 2)$ sigma model, based on the exposition in [57, Ch. 16].

Topological twisting can be seen as a way of modifying the theory so that one can compute its partition function on a curved worldsheet Σ, while preserving some amount of supersymmetry. This changes the value of the partition function of the theory on curved surfaces, but preserves its value on flat surfaces; in particular the Hilbert space assigned to a circle (which is the module over the superconformal algebra we discussed above) stays the same, since it can be calculated by a flat cylinder.

Twisting these theories requires first relating the Euclidean symmetry of the theory with its R-symmetry. For the $\mathcal{N} = (2, 2)$ sigma model, let us denote by $U(1)_E$ the rotation group inside the Euclidean symmetry group of the worldsheet. The R-symmetry group of the sigma model depends on the geometry of the target space X one is considering; for the case of immediate interest, where X is Calabi-Yau, the theory has two different unbroken $U(1)$ R-symmetries, the vector and the axial symmetries. These symmetries rotate the supercharges of the theory; the vector symmetry $U(1)_V$ acts by

$$Q^{\pm} \mapsto e^{-i\alpha} Q^{\pm}, \quad \bar{Q}^{\pm} \mapsto e^{i\alpha} \bar{Q}^{\pm}$$

and the axial symmetry $U(1)_A$ acts by

$$Q^{\pm} \mapsto e^{\mp i\alpha} Q^{\pm}, \quad \bar{Q}^{\pm} \mapsto e^{\pm i\alpha} \bar{Q}^{\pm}$$

Topologically twisting the sigma model involves choosing either one of these two $U(1)$ symmetries to use; we'll denote by $U(1)_R$ this choice. In the Calabi-Yau case, since we have the two choices above we will call the *A-twist* the case where we choose $U(1)_R = U(1)_V$ and the *B-twist* the case $U(1)_R = U(1)_A$.

The twisting procedure can be seen as a two-step process. We first change what we consider to be the 'Euclidean' symmetries of the theory; this means that instead of using the spin connection to gauge the symmetry $U(1)_E$, we use it to gauge the diagonal subgroup $U(1)_{E'}$ inside $U(1)_E \times U(1)_R$. Mathematically, this can be seen as a reinterpretation of the symbols $\phi, \psi_{\pm}, \bar{\psi}_{\pm}$ that denote the fields on Σ, ie. after the twist they will denote sections of different bundles, with different charges under the various $U(1)$ symmetries; the exact description of these changes is given eg. in [57, Sec.16.2.2].

The second step involves choosing a supercharge Q that transforms as a scalar after twisting; for the A-twist this charge is $Q_A = \bar{Q}_+ + Q_-$, and for the B-twist it is $Q_B = \bar{Q}_+ + \bar{Q}_-$ Then one declares that the only 'physical' operators of the theory are operators commuting with Q, and that the 'physical' Hilbert space is given by states annihilated by Q. We will denote by $\mathcal{O}$ a Q-closed operator, ie. an operator such that $[Q, \mathcal{O}] = 0$; the correlation function of a collection of operators $\mathcal{O}_i$ is then given by a path integral

$$\langle \mathcal{O}_1 \dots \mathcal{O}_n \rangle = \int \mathcal{D}\phi \mathcal{D}\psi^{\pm} \mathcal{D}\bar{\psi}^{\pm} e^{-S} \mathcal{O}_1 \dots \mathcal{O}_n$$

which can be computed by Q-localization.

Moreover, using the nilpotency of the supercharge Q, one can show that Q-exact operators, ie. operators of the form $\mathcal{O}_{\text{exact}} = \{Q, \mathcal{O}'\}$ are zero in this theory, in the sense that any correlation function involving it vanishes by integration by parts

$$\langle \mathcal{O}_{\text{exact}} \mathcal{O}_1 \dots \mathcal{O}_n \rangle = \int \mathcal{D}\phi \mathcal{D}\psi^{\pm} \mathcal{D}\bar{\psi}^{\pm} e^{-S} \{Q, \mathcal{O}'\} \mathcal{O}_1 \dots \mathcal{O}_n = 0$$

so that this represents the zero element in the algebra of operators.

An important feature of the twisted sigma model is that the stress energy tensor $T^{\mu\nu}$ is Q-exact; this implies that by the reason above that the theory is invariant under variations of the worldsheet metric. This is the reason for calling it a *topological* twist; historically these were some of the first topological theories that were discovered. The field of topological quantum field theory has since become very vast, and in more current literature, theories such as these are known as 'Witten-type TQFTs' or 'cohomological field theories', as opposed to other types of TQFTs such as Chern-Simons or Dijkgraaf-Witten theories, where the partition function is literally independent of variations of the metric, as opposed to independent up to cohomologically exact terms.

By now, many examples of twisted theories have appeared in the literature about supersymmetric field theories. All of these examples rely on the existence of exceptional homomorphisms from the Euclidean symmetry group to the R-symmetry group. Some of these are not fully topological in the sense mentioned above, but rather have a mix of topological and holomorphic dependence on different coordinates, a feature that appeared early in Kapustin's study of S-duality in gauge theories [62].

The A- and B-models with CY target

Let X denote a Calabi-Yau manifold of complex dimension n. As described above, there are two distinct topological twists of the $\mathcal{N} = (2, 2)$ sigma model with target X; let us discuss the main features of each one of these twists.

The A-model

Point operators in the A-model are assembled out of the fields that have become scalars after the topological twist, ie. the fields ψ_- and $\bar{\psi}_+$, and the $2n$ (real and imaginary) components of the field $\phi : \Sigma \to X$. That is, a general point operator is of the form

$$\mathcal{O} = f(\phi)_{i_1 \dots i_p} \psi_-^{i_1} \dots \psi_-^{i_p} \bar{\psi}_+^{j_1} \dots \bar{\psi}_+^{j_q}$$

where the f is some (not necessarily holomorphic) function on the target space. One can argue that these components transform under change of coordinates in a way that these operators should be identified with a (p, q)-form on X; explicitly in a holomorphic chart in X we can identify

$$\psi_-^i \mapsto dz^i, \quad \bar{\psi}_+^i \mapsto d\bar{z}^i$$

Under this identification, the supersymmetry generators $Q_-, \bar{Q}_+$ act as $\partial, \bar{\partial}$, respectively, so that the chosen supersymmetry Q_A acts as the de Rham differential $d = \partial + \bar{\partial}$. Therefore the physical operators get identified with the group of de Rham cohomology classes $H^*_{dR}(X, \mathbb{C})$ of the target space.

The sum of all the cohomology groups of X have a natural ring structure given by the ordinary cup product; here they also acquire another ring structure coming from the chiral ring structure of the underlying SCFT; ie. the product of cohomology classes given by operators $\mathcal{O}_1, \ldots, \mathcal{O}_n$ is their correlation function. This, in turn, is given by a path integral over all maps $\phi : \Sigma \to X$. We can split this into a sum over the integral homology class $\beta \in H_2(X, \mathbb{Z})$ of the image of Σ.

$$\langle \mathcal{O}_1 \ldots \mathcal{O}_n \rangle = \sum_{\beta \in H_2(X,\mathbb{Z})} \langle \mathcal{O}_1 \ldots \mathcal{O}_n \rangle_\beta$$

The important feature of these topological models is that supersymmetric localization implies that this integral only received contributions from a very restricted subset of maps $\Sigma \to X$, which makes possible a rigorous mathematical definition of this path integral. In this case, Q_A localization means that only holomorphic maps $\Sigma \to M$ contribute; therefore this correlation function can be expressed as an integral over a well-defined moduli space of maps $\mathcal{M}_\Sigma(X, \beta)$.

$$\langle \mathcal{O}_1 \ldots \mathcal{O}_n \rangle_\beta = e^{-(\omega - iB)\beta} \int_{\mathcal{M}_\Sigma(X,\beta)} \mathrm{ev}_1^* \, \omega_1 \ldots \mathrm{ev}_n^* \, \omega^n$$

with an integrand in terms of tautological classes on the moduli space.

When properly defined, these integrals give the correlation function in terms of the *Gromov-Witten invariants* of X; and endow the (de Rham) cohomology $H^*(X, \mathbb{C})$ of X with a new product, which can be shown to be a deformation of the ordinary cup product. The resulting ring is referred to as instanton-corrected cohomology or, more commonly in the mathematical literature, *quantum cohomology* of X.

It is important to remark that the A-model can be defined on any Kähler manifold X, and not necessarily on a Calabi-Yau; it will then have a mirror that is not necessarily a geometric sigma model, instead being given by a Landau-Ginzburg model. Also, up to a correct notion of equivalence, the A-model depends only on the symplectic structure of X; different choices of complex structures required to define the moduli space of holomorphic maps give equivalent models. So the A-model can be said to capture the 'symplectic side' of the geometry of X.

The B-model

For the B-twist of the sigma model, let us take again X to be a Calabi-Yau, since one requires the axial symmetry to be part of the R-symmetry group of the SCFT.

Similarly to what we described above for the A-model, after performing the B-twist, one can identify what (combinations) of the fields give scalars, and then from them assemble the point operators of the theory. It turns out that certain combinations of components of the fields ϕ and $\bar{\psi}$ transform as holomorphic vector fields $\partial/\partial z^i$ and others transform as anti-holomorphic one-forms $d\bar{z}^i$; therefore a general point operator is given by an expression

$$\mathcal{O} = \omega^{j_1,\ldots,j_q}_{i_1,\ldots,i_p} dz^{i_1} \ldots dz^{i_p} \left(\frac{\partial}{\partial z^{j_1}}\right) \ldots \left(\frac{\partial}{\partial z^{j_q}}\right)$$

ie. an form in $\Omega^{0,p}(X, \wedge^q T)$ where T is the holomorphic tangent bundle. Under this identification the supersymmetric charge Q_B acts as the Dolbeault differential $\bar{\partial}$ so the space of physical states is given by the total sum of Dolbeault cohomology groups $H^{0,*}(X, \wedge^* T)$.

As for the A-model, the path integral giving the n-point correlator of the B-model can be computed by supersymmetric localization. However, the action of the supercharge Q_B on the coordinates ϕ_i of the map $\phi : \Sigma \to X$ is $\delta\phi^i = 0$, which in turn means that the path integral localizes to constant maps $\partial_\mu \phi^i = 0$. So the correlation is given by an integral over the target space X itself:

$$\langle \mathcal{O}_1 \dots \mathcal{O}_n \rangle = \int_X \langle \omega_1 \wedge \dots \wedge \omega_n, \Omega \rangle \wedge \Omega$$

where Ω is the given holomorphic n-form.

The statement of mirror symmetry for Calabi-Yau manifolds is that given a Calabi-Yau manifold X of dimension d, there is another Calabi-Yau X' of the same dimension, such that the A-model on X is equivalent to the B-model on X' and vice-versa; this implies, in particular, a deep relation between the Gromov-Witten invariants of X, which are notoriously hard to calculate, and periods of certain differentials on X', which were a subject of classical interest in complex geometry; this is the starting point of the calculations of Candelas et al. [26].

1.3 The homological mirror symmetry conjecture: D-branes, Fukaya categories and coherent sheaves

Kontsevich's 1994 Homological Mirror Symmetry conjecture [71] suggested that the phenomenon of 'closed string' mirror symmetry as we sketched above could be explained as a shadow of a categorical equivalence, relating the symplectic and complex geometries of mirror spaces. The enumerative statements of mirror symmetry could then be derived from this categorical equivalence by taking certain invariants on both sides.

More specifically, this conjecture states that given a mirror pair of Calabi-Yau manifolds X, X', there is an equivalence of categories between the derived Fukaya category of X and the derived category of coherent sheaves of X'. At the time of this conjecture, while the derived category of coherent sheaves was a familiar object to algebraic geometers, the Fukaya category was a more obscure construction, coming from Lagrangian Floer theory. This object was initially defined by Fukaya as an A_∞-category, motivated by the existence of natural product structures in Lagrangian Floer homology. Very roughly, the objects of the Fukaya category are Lagrangians and the morphisms are generated by intersection points between these.

Decades later, there is still quite a bit of discussion on what are the most natural 'correct' structures on Fukaya categories for general manifolds (eg. whether one should include immersed objects or define it as a curved A_∞-category etc.). This is not so much of a problem to the HMS conjecture since most of these differences seem to disappear once one passes to the derived category. In this dissertation we will later deal with some classes of examples where particular combinatorial models for the Fukaya category exist.

Some years after the proposal of the HMS conjecture, with the development of the theory of D-branes in string theory, it became clear that this correspondence should be understood in terms of an equivalence of the categories of boundary conditions of the A- and B-models on mirror spaces. More specifically, from the point of view of the open string topological theories, the category of boundary conditions for the A-model (category of A-branes) should be given by some version of the Fukaya category, and on the B-side the category of B-branes should be given by the derived category of coherent sheaves.

There is some difficulty in reconciling these two points of view; for example it takes a non-trivial amount of effort to argue that on a general coherent sheaf should provide a consistent boundary condition for the B-model. There are also problems on the other side; for example, Kapustin and Orlov [64] argued that the A-model can also admit other non-Lagrangian boundary conditions, the so-called coisotropic branes. These have appeared prominently in applications in high-energy physics [87], but on the mathematical side their existence and relation to the statement of homological mirror symmetry is still quite mysterious; though there is evidence that their understanding is fundamental in making precise some constructions that are of interest in mathematics, especially in the area of quantization [55].

In this section we will give a cursory introduction to these objects, with out applications of later chapters in mind. For a more comprehensive exposition of the topic of D-branes in mirror symmetry, the reader can consult the references [57, Ch.37-39] and [8, Ch.3]; here we will mostly follow these references.

Branes in the supersymmetric sigma model

We will analyze the possible boundary conditions for the sigma-model geometrically, roughly following the first reference mentioned above. Let us first fix X a CY n-fold, and now let us consider a boundary condition given by a submanifold $C \subset X$; this means that we will be quantizing a theory of open strings that are constrained to end on C. It turns out to be more natural to generalize this point of view and also allow a $U(1)$ gauge field A on C.

The action for an open string traveling along a worldsheet $\Sigma \subset X$ with boundary $\partial\Sigma \subset C$ is then given by

$$S = \int_\Sigma \mathcal{L} = \int_\Sigma \left(-g_{ij}\partial^\mu \phi^i \partial_\mu \phi^j + \phi^* B\right) + \int_{\partial\Sigma} \phi^* A + \text{ fermionic terms}$$

The equations of motion for the action above can be written succinctly in terms of holomorphic/antiholomorphic coordinates $z, \bar{z}$ on the worldsheet σ by

$$\frac{\partial \phi^i}{\partial z} = R^i_j \frac{\partial \phi^j}{\partial \bar{z}}$$

when we set the fermion and B-fields to zero. The matrix R_{ij} is defined as an orthogonal matrix-valued function on C, defined by

$$R = (g|_{TC} - F)^{-1}(g|_{TC} + F)$$

where F is the curvature two-form of the gauge field A. This equation of motion gets deformed in the presence of fermions, but one can also use this matrix R for the equations of motion of the fermionic fields $\psi_\pm$; these relate the left and right moving fermions as a reflection relation

$$\psi_+^i = R_j^i \psi_-^j$$

along C. A full treatment of these equations of motion, including special cases where more boundary conditions are present due to 'generalized geometry', can be found in [75].

This sketch of the equations of motion above is enough to illustrate that it is not possible to preserve all the of the $\mathcal{N} = (2, 2)$ supersymmetry with such a boundary condition, since that algebra rotates between the left and right-moving fermions and the reflection relation above must break part of that supersymmetry. It turns out that for each of the topological twists of this sigma model it is possible to preserve the necessary amount supersymmetry, as long as the submanifold C and gauge field A satisfy appropriate equations.

A-branes

We use the complex structure J on the Calabi-Yau target X and pick local holomorphic and anti-holomorphic coordinated indexed by $i = 1, \ldots, m$ and $\bar{i} = \bar{1}, \ldots, \bar{m}$. To preserve the supersymmetry under the generator Q_A, the diagonal blocks R_j^i and $R_{\bar{j}}^{\bar{i}}$ have to be zero, and the matrix-valued function R needs to be block off-diagonal.

Writing $R = (g|_C - F)^{-1}(g|_C + F)$, one can interpret this condition in terms of the geometry of C and the curvature of the $U(1)$ bundle on it; one possible solution is given by

$$J^2|_{TC} = -1, \quad \omega|_C = 0, \quad F = 0$$

where J exchanges the tangent and normal directions to C, so J^2 is well-defined as an endomorphism of TC. Therefore C is a middle-dimensional Lagrangian submanifold with a flat $U(1)$ connection (in the presence of background field B this gets modified to $F = B|_C$). This describes a rank 1 Lagrangian brane; one can consider also higher-rank branes carrying $U(N)$-vector bundles for $N > 1$; as usual in D-brane constructions, this arises from stacks of rank 1 branes, where the $U(1)^N$ gauge group gets enhanced to $U(N)$ by the presence of massless states.

In general one needs to worry about quantum effects that might make such a boundary condition anomalous; the most important example of such an effect for us is the axial anomaly. From Subsection 1.2, remember that when the target X is Calabi-Yau there are two $U(1)$ R-symmetries of the sigma model; and upon performing the A-twist we used up the vector symmetry $U(1)_V$; the resulting axial symmetry $U(1)_A$ is still unbroken in the closed string A-model. In order for the A-brane supported on the Lagrangian C to have vanishing anomaly, we have the following condition. Consider the holomorphic CY m-form Ω; at any point $p \in C$ we define $\xi(p) \in S^1$ to be the phase such that

$$\Im(e^{-i\pi\xi(p)}\Omega|_C) = 0$$

this defines a map $\xi : C \to S^1$ which represents a class in $H^1(C, \mathbb{Z})$, called the Maslov class of the Lagrangian cycle C [1]. The axial symmetry is anomalous if this class is non-vanishing. So we will restrict attention to Lagrangian branes of vanishing Maslov class.

This motivates the identification between A-branes and the Fukaya category of Lagrangian submanifolds; the vanishing Maslov class condition also plays a role in lifting the grading in Lagrangian Floer homology to a $\mathbb{Z}$-grading [10]. There are still many subtle facts about the A-model that have not been formalized completely into the mathematical language. Most notably, in the solution above we could have tried to solve the supersymmetry equations with $F \neq 0$, ie. with a non-flat $U(N)$ vector bundle. This leads to the discovery of coisotropic branes [64], supported on submanifolds of dimension greater than the middle dimension n. The resulting equations for ϕ and F often do not have non-trivial solutions; notable exceptions are the cases of hyperkähler targets analyzed in works such as [55]. The exact relation of these other boundary conditions with the mathematical setting of the HMS conjectures is still quite mysterious; very recently [69] gave some constructions of large classes of examples of coisotropic branes, conjectured to be related to a number of conjectures in algebraic geometry. For the rest of this dissertation we will restrict attention to Lagrangian A-branes.

Strings between A-branes

Here we sketch the general arguments on why the spaces of strings in the A-model should be related to the morphism spaces of the Fukaya category. Consider two rank one A-branes supported on Lagrangians L_0, L_0 intersecting transversely at a finite set of points $p_1, \dots, p_N$.

According to the principle of cohomological field theories, in the A-model we should interpret the supercharge Q_A as a differential acting on the space of string states stretching between the branes, and the corresponding physical Hilbert space will be given by the cohomology with respect to Q_A. Let us sketch how this works in the case at hand.

The Dirichlet condition along the branes L_0, L_1 means that to compute the Hilbert space $\mathcal{H}(L_0, L_1)$ we will quantize the theory of maps from a strip

$$\phi : [0, 1] \times \mathbb{R} \to X$$

Parametrizing this strip by coordinates (σ, τ), the boundary conditions read

$$\phi(0, \tau) \in L_1, \quad \partial_\sigma \phi(0, \tau) \in \phi^* N L_0$$

[1]The Maslov class (or index) of a Lagrangian can be defined for any symplectic manifold, without reference to complex or CY structures; one can prove that the definition given in the text is independent of those choices and agrees with the definition using the topology of the Lagrangian Grassmannian

and similarly for the other end at $\sigma = 1$. There are similar conditions for the fermions $\psi_\pm, \bar{\psi}_\pm$. Let us denote the space of such paths by $\Omega(L_1, L_2)$. In terms of variational derivatives, the supercharge Q_A acts on this space as

$$iQ_A = \delta + \delta h \wedge$$

where h is a function measuring 'how far' the map ϕ is from a constant fixed map ϕ_0. This function itself depends on these choices but its variation is given by

$$\delta h = \int_{[0,1]} \left(-ig_{ij}\delta\phi^i d\bar{\phi}^j + ig_{ij}d\phi^i \delta\bar{\phi}^j\right) d\sigma$$

The closed states under Q_A then correspond to critical points of the function h which by the formula above are constant maps localized at a given intersection point $p \in L_0 \cap L_1$. A similar analysis in the higher rank case shows that the states then should be spanned by morphisms between the stalks of the vector bundles at the intersection points, so our model for the space of states is the cohomology complex

$$\mathcal{H}(L_0, L_1) = \left(\bigoplus_{p \in L_0 \cap L_1} \mathrm{Hom}_{\mathbb{C}}((E_1)_p, (E_2)_p)\right) \Big/ \mathrm{Im}(Q_A)$$

One can then ask two things: whether the axial R-symmetry gives a consistent grading on the set of intersection points by ghost number, such that this is an honest $\mathbb{Z}$-graded complex of vector spaces, and whether this is truly a chain complex ie. $Q_A^2 = 0$. The answer to the first question is quite delicate in the general case due to the anomaly of the axial symmetry [2], but in the CY case we are considering this anomaly vanishes and one can choose a well-defined grading of the critical points such that Q_A has homogeneous degree 1. The second question, of whether the differential squares to zero, is connected to some delicate questions of bubbling of Lagrangians and appropriate perturbations; in particular it can be shown that these problems do not happen if L_0 and L_1 bound no holomorphic disks of positive area so this is the setting we will assume.

This differential can then be computed by supersymmetric localization, just as the closed string n-point functions were. Here localization implies that we can restrict attention to maps $\phi : \Sigma \to X$ that are holomorphic with respect to the complex structure of X. To calculate the differential, let us look at the moduli of holomorphic strips; index considerations imply that the only nonzero contributions to the path integral come from strips connecting intersection points of adjacent degree, ie. maps ϕ satisfying

$$\bar{\partial}\phi^i = 0, \quad \phi(0,\tau) \in L_0, \quad \phi(1,\tau) \in L_1, \quad \lim_{\tau \to \pm\infty} \phi(\sigma,\tau) = p_\pm$$

[2] In particular, for non CY targets the discussion in the text needs to be appropriately modified, considering instead covers of the path space $\Omega(L_0, L_1)$; for a discussion of these issues see [57, p. 39.4.2].

where $p_\pm$ are intersection points with degree satisfying $\mu(p_+) - \mu(p_-) = 1$. The differential Q acts then by mapping the classes supported at the intersection points as $[p_-] \to [p_+]$, up to a factor encoding the area of the disk [3].

Having in mind the proposed relation between the A-model and the Fukaya category, we will denote this complex of states by $\mathrm{Hom}^*(L_0, L_1)$, with grading by ghost number. One can interpret the map above, instead as a map from a strip, as a map from a disk with two marked points along its boundary. Then it becomes natural to consider the analogous question with $n \geq 2$ marked points along the boundary, constrained to map respectively to Lagrangian branes $L_0, \dots, L_{n-1}$; a similar localization argument as above indicates that the path integral should count holomorphic maps from this disk, and these counts should be assembled into "higher structure maps" giving relations between the morphism spaces $\mathrm{Hom}(L_0, L_1), \dots, \mathrm{Hom}(L_{n-2}, L_{n-1}), \mathrm{Hom}(L_{n-1}, L_0)$. Later, in Section 2.2 we will see how these higher structure maps are formalized in the Fukaya category setting as the notion of an A_∞-structure.

B-branes

Let us discuss the boundary conditions on the other side of mirror symmetry, for the B-twist of the sigma model with a Calabi-Yau target. The main parts of this thesis, in all subsequent chapters, will not really need to discuss B-branes, as they will all concern the Fukaya category, so we will limit ourselves to a cursory introduction to B-branes, mostly in the spirit of motivating the appearance of the derived category in mirror symmetry.

Consider again a brane supported on a submanifold C, together with a bundle $E \to C$, and let us look at the supersymmetry conditions. Again using local holomorphic and anti-holomorphic coordinates on the target space indexed by $i = 1, \dots, m$ and $\bar{i} = \bar{1}, \dots, \bar{m}$, the equations for the matrix R that must be satisfied to preserve the supersymmetry generated by Q_B imply that R is block-diagonal with respect to J, ie. $R^i_{\bar{j}} = R^{\bar{i}}_j = 0$. So J now preserves the tangent and normal directions to C, which implies that C is a holomorphic submanifold of X. The equations for the connection now imply that its curvature satisfies $F \in \Omega^{1,1}(C, \mathrm{End}(E))$, meaning that E is a holomorphic vector bundle [4].

Strings between B-branes

For simplicity let us discuss the case where the submanifold C is all of X, and two different B-branes will just be specified by two holomorphic vector bundles $E_0, E_1 \to X$. A general discussion of the space of states in the B-model would involve a discussion of the problem of 1d supersymmetric quantum mechanics such as in [57, Sec.10.4], but fortunately for our purposes when the target X is CY, one can show that the space of states reduces to the space of fermionic zero modes.

[3] depending on which ground ring one is using to define the Fukaya category; see the review [10]

[4] This is technically only strictly true when the B-field is zero.

This means that if we are parametrizing the worldsheet of this string as a strip $\phi : [0,1] \times \mathbb{R} \to X$ with coordinates σ, τ as before, and with fermionic fields $\psi_\pm, \bar{\psi}_\pm$, we will consider only states where $\phi, \psi_\pm, \bar{\psi}_\pm$ have no σ-dependence. The boundary condition forces

$$\psi_- - \psi_+ = \bar{\psi}_- - \bar{\psi}_+ = 0$$

and the algebra of the two remaining fermionic degrees of freedom generates the Hilbert space of the theory

$$\bigoplus_{p=0}^{n} \Omega^{0,p}(X, E_0^* \otimes E_1) = \bigoplus_{p=0}^{n} \Omega^{0,p}(X, \mathrm{End}(E_0, E_1))$$

as $\bar{\psi}_+^i + \bar{\psi}_-^i \mapsto d\bar{z}^i$ and $\psi_+^i + \psi_-^i \mapsto g_j^i \frac{\partial}{\partial \bar{z}^j}$.

Just as in the case of the closed string B-model, under this identification the supersymmetry generator Q_B acts on the vector bundle $\mathrm{End}(E_0, E_1)$ as the Dolbeault differential $\bar{\partial}$, so the physical Hilbert space of strings between these branes is given by the total Dolbeault cohomology

$$\mathcal{H}(E_0, E_1) = \bigoplus_{p=0}^{n} H^{0,p}(X, \mathrm{End}(E_0, E_1))$$

Also analogous to the closed string B-model, the correlators here can also be computed simply in terms of the wedge product of differential forms, integrated against the Calabi-Yau holomorphic n-form Ω. That is, given forms

$$a \in \Omega^{0,p}(X, \mathrm{Hom}(E_0, E_1)), \quad b \in \Omega^{0,q}(X, \mathrm{Hom}(E_1, E_2)), \quad c \in \Omega^{0,n-p-q}(X, \mathrm{Hom}(E_2, E_0))$$

with corresponding insertion point operators W_a, W_b, W_c, the correlation function up to all orders and genera is simply given by integrating the trace over the space of constant maps:

$$\langle W_a W_b W_c \rangle = \int_X \mathrm{Tr}(a \wedge b \wedge c) \wedge \Omega$$

One can argue also by localization under Q_B that there are no instanton corrections to these correlation functions; this is one of the aspects in which the B-model appears much simpler than its mirror.

The appearance of the bounded derived category of coherent sheaves

Even in the original formulation of the HMS conjecture by Kontsevich [71], the category appearing on the B-side was already the bounded derived category of coherent sheaves. One can motivate this by the arguments that 1) even for the most classical setting of mirror symmetry between Calabi-Yau manifolds, comparing the set of objects of the Fukaya category to the set of holomorphic vector bundles on the mirror, one notices that there are not enough vector bundles and 2) even a naive characterization of B-branes as above shows that one must also include objects encoding holomorphic vector bundles supported on lower-dimensional submanifolds.

Still, one must motivate the appearance of the features of the derived category from physical arguments; this is the main topic of investigations for a series of papers by Sen [97], Douglas [38], Sharpe [98], Aspinwall-Lawrence [7] and others around the year 2000.

Let us sketch some of the main ideas concerning the appearance of the grading and the importance of quasi-isomorphisms, following mostly [8, Sec.5.3]. From the perspective of the B-model as a topological open string theory, this extra grading comes from the $U(1)$ R-symmetry and takes integral values. One can combine boundary conditions of different degrees; if $E^n, n \in \mathbb{Z}$ are vector bundles on X, a boundary condition for the B-model with target X can be given by

$$E = \bigoplus_n E^n$$

It is important to note that from the perspective of the SCFT coming from the (untwisted) sigma model with CY target X, there is a slightly different description of this grading, which takes values in $U(1)$ or $\mathbb{R}$ instead of $\mathbb{Z}$. In the sigma model, one can compute the R-charge from the geometric data on X; for example, for some vector bundle $E \to X$, the corresponding charge is [7]:

$$Z(E) = \frac{1}{\pi} \Im \log \int_X e^{B+iJ} ch(E) \sqrt{td(T_X)} + \dots$$

which, though naively valued in $U(1)$, turns out to be more naturally valued in the universal cover $\mathbb{R}$ by a monodromy argument. The relation between this $\mathbb{R}$-valued charge in the SCFT and the $\mathbb{Z}$-valued charge in the B-model has been historically explained by the fact that not every superposition of bundles is a boundary condition preserving the necessary supersymmetry Q_B; for this to happen the charges defined above need to differ by integer amounts, so by an overall degree shift they can all be made $\mathbb{Z}$-valued. This is a partially satisfying explanation; a more complete explanation can be found in the discussion of the concept of Π-stability in 2d SCFTs [39, 6].

As for the appearance of chain complexes, this comes from the fact that chiral operators of R-charge one give deformations of boundary conditions. For example, for a B-type a single holomorphic bundle E and a class $\delta A^{(0,1)} \in \mathrm{Ext}(E, E)$, we can deform this boundary condition by changing the BRST operator acting on the states by

$$\bar{\partial} \mapsto \bar{\partial} + \delta A^{(0,1)}$$

A similar argument can be used to understand any R-charge one operator; consider now the superposition $E = \bigoplus_n E^n$ as above, and operators $d_n \in \mathrm{Hom}(E^n, E^{n+1}), d = \sum_n d_n$. One can write a boundary deformation of the action which corresponds to a deformation of the BRST charge by

$$Q \mapsto Q + d$$

and now asking that this operator squares to zero corresponds to $d_{n+1}d_n = 0$ which is the condition for a complex.

Moreover, a similar argument shows that given a homotopy equivalence between complexes E and F, that is a pair of maps composing to the identity up to a homotopy, the spaces of string states differ by a Q-exact quantity; this motivates the appearance of the homotopy category $K(\mathrm{Coh}(X))$. One needs to extend this argument a bit to get to the derived category. This has been explained early on the history of the subject as a consequence of the existence of open string tachyonic states associated with extension morphisms [97]; after condensation flowing under the RG flow should parallel the corresponding equivalence in the derived category.

There is another argument for the appearance of the derived category due to Aspinwall and Lawrence, involving the idea of 'physical equivalence'; this basically says that if two boundary conditions have equal correlators with every other boundary condition, then they should be considered the same. More precisely, given two objects E, E' of the homotopy category $K(\mathrm{Coh}(X))$, they are said to be physically equivalent if for any other object F we have $\mathrm{Ext}^*(E, F) \cong \mathrm{Ext}^*(E', F)$ and $\mathrm{Ext}^*(F, E) \cong \mathrm{Ext}^*(F, E')$ and moreover the products given by the 3-point functions agree; then one can prove [7] that $K(\mathrm{Coh}(X))$ modulo these relations is equivalent to the derived category.

Chapter 2

TCFTs, Calabi-Yau structures, microlocal models for Fukaya categories

In this chapter, we will introduce the general mathematical formalism of open-closed 2d TQFTs, and also discuss some extensions of that concept, designed to accommodate theories such as the A and B-models with CY targets that we discussed in Ch.1. For us, this model will come in the notion of a *topological conformal field theory*, or TCFT for short.

The purpose of providing exposition of this topic is to give a context for the presentation of one of this dissertation's main results, which concerns the construction of *Calabi-Yau structures* on certain categories. As we will see below, the boundary conditions of a TQFT form a category, and given some category $\mathcal{C}$, a compatible pair of a A_∞ structure and a Calabi-Yau structure determines a TCFT with $\mathcal{C}$ as its boundary conditions.

As part of the background for the results of Chapter 3, we will also introduce the generalizations of the Fukaya category for non-compact manifolds that will appear there, and discuss a particular model for this category, in terms of categories of microlocal sheaves, that has been a very active topic of research in recent years, and which we will use as the main application of the techniques of Chapter 3.

2.1 The TQFT/TCFT perspective

All the mathematical models for topological theories that we will be working with are in the spirit of the Atiyah-Segal axioms [9, 95], which formalize the algebraic structure of partition functions of a topological theory. The subject of TQFTs has been studied so much, and from so many perspectives, that by now one can find in the literature a great number of introductory texts; the reader can consult [45] for an earlier reference, and [108] for an introductory account of the recent developments on the mathematics of TQFTs.

We will not attempt to give a general account of TQFTs but instead pass directly into the context of open-closed 2d TQFTs, which is the setting that is most relevant to us.

2d open-closed TQFTs

Here we follow the exposition in [8, Sec.2.1]. The starting point of an open-closed theory is a category $\mathcal{B}$ of boundary conditions. Note that in the case of a theory obtained from a 2d sigma model of maps from the worldsheet Σ of a string, the objects $a, b, \dots$ of this category would label branes on which the string is allowed to end. For convenience we will often just refer to these objects in general as branes.

The open-closed theory assigns two types of spaces of states: spaces of states $\mathcal{O}_{ab}$ on open strings stretching between branes a and b, and spaces of states $\mathcal{C}$ on the closed string.

The structure of a 2d open-closed TQFT (over $\mathbb{C}$) is given by:

1. $\mathbb{C}$-vector spaces $\mathcal{C}$, and $\mathcal{O}_{ab}$ for any pair of branes a, b
2. A commutative Frobenius algebra structure on $\mathcal{C}$ given by a multiplication map with identity $1_{\mathcal{C}}$ and a trace map $\theta_{\mathcal{C}} : \mathcal{C} \to \mathbb{C}$)
3. Associative bilinear products $\mathcal{O}_{ab} \otimes \mathcal{O}_{bc} \to \mathcal{O}_{ac}$
4. A nondegenerate trace map $\theta_a : \mathcal{O}_{aa} \to \mathbb{C}$ giving a non-necessarily commutative Frobenius algebra structure on $\mathcal{O}_{aa}$
5. Linear maps $\iota_a : \mathcal{C} \to \mathcal{O}_{aa}$ and $\iota^a : \mathcal{O}_{aa} \to \mathcal{C}$

This collection of maps satisfies a number of relations [8, Sec.2.1.2] that can be encoded in a finite number of surface diagrams; then the 'sewing theorem' implies that this data uniquely defines the TQFT.

One consequence of these relations is that for any pair a, b, the composition

$$\mathcal{O}_{ab} \otimes \mathcal{O}_{ba} \to \mathcal{O}_{aa} \xrightarrow{\theta_a} \mathbb{C}$$

is a nondegenerate pairing, giving a canonical duality between $\mathcal{O}_{ab}, \mathcal{O}_{ba}$. This defines the structure of a *Calabi-Yau* category. We will later discuss this concept in greater detail.

TCFTs

We would like to regard the A- and B-models in this functorial formalism. However there is a subtle distinction, that turns out to be essential. This formalism of 2d open-closed TQFTs is better suited for theories that are topologically invariant *on the nose*. On the other hand, the topological twists we care about are 'Witten-type' or cohomological field theories; ie. in each case there is a BRST operator Q such that a deformation of the worldsheet leaves the action invariant just up to Q-exact terms.

Therefore to get a better functorial picture of these models, one must work in a properly cohomological setting, where one remembers eg. the chains on the worldsheet instead of its mere homeomorphism type. The model that we will use is usually called in the literature by the name of *topological conformal field theories* (TCFTs). These were initially discussed by Segal and described by Getzler in [53]. Later, Costello [30, 31] developed this framework and applied it to context of higher-genus mirror symmetry. Here we will follow the notation and formalism of these papers.

Let us fix a set Λ of branes, ie. boundary conditions. A TCFT is defined as an appropriate functor from a geometrically-defined category $\mathcal{OC}^d_\Lambda$ depending on Λ and an integer d which should be thought of as the dimension of the target space. This functor is a cochain-level enhancement of the data of a 2d open-closed TQFT as we defined above.

The category $\mathcal{OC}^d_\Lambda$ has as objects collections of boundary circles, which can be closed strings or open strings. That is, its objects are quadruples (C, O, s, t), where $C \in \mathbb{Z}_{\geq 0}$ is the number of closed strings, $O \in \mathbb{Z}_{\geq 0}$ is the number of open strings, and $s, t : O \to \Lambda$ are the source and target maps for each open string.

For a pair of such objects $(C_\pm, O_\pm, s_\pm, t_\pm)$, consider the moduli space $\mathcal{M}_\Lambda(+,-)$ of Riemann surfaces with incoming closed/open boundaries C_+, O_+ and outgoing closed/open boundaries C_-, O_-. There is a 'determinant' local system det on this moduli space, whose fiber is defined from from the topology of the corresponding surface [30, Sec.3]. Let us then define the morphisms between these two objects to be the twisted chains on this moduli space, ie.

$$\mathrm{Hom}_{\mathcal{OC}^d_\lambda}(+,-) := C_*(\mathcal{M}_\Lambda(+,-), \det{}^d)$$

One can prove that this defines a category enriched over the category of chain complexes. This is only true after one takes care of some technicalities involved in this definition; see the reference above for a full treatment.

Definition 1. [30] An open-closed TCFT of dimension d on a set Λ of boundary conditions (for some field k) is a symmetric monoidal functor

$$\Phi : \mathcal{OC}^d_\Lambda \to \mathrm{Mod}_k$$

to the category of differential graded k-vector spaces.

One can define appropriate subcategories $\mathcal{C}^d_\Lambda, \mathcal{O}^d_\Lambda$ of closed and open string states and write analogous definitions of closed or open TCFT.

This definition above gives a chain-level enhancement of the notion of an 2d open-closed TQFT; for a more modern perspective on the subject, in the setting of the cobordism hypothesis of Baez-Dolan [11] later proved by Lurie, these TCFTs should be regarded as 'extended' 2d TQFTs [78, Sec.4.2].

2.2 A-infinity categories and Calabi-Yau structures

Comparing the definition of a TCFT with our previous definition of a TQFT, one can ask whether there is an intrinsic structure on the set Λ that uniquely characterizes a given TCFT, the same way that the structure of a Frobenius algebra/category characterizes a 2d TQFT. The answer turns out to be yes; but one needs first to define an appropriate homotopical enhancements of the associativity condition and the trace function. These enhancements are captured respectively by the notions of an A_∞-category and a compatible Calabi-Yau structure.

A_∞-categories

A A_∞-category is a generalization of a category to a structure where composition of morphisms is only associative up to higher coherence maps. The degree conventions for defining an A_∞ structure vary a bit throughout the literature; for consistency with later parts of this dissertation let us use a different convention from [30]; in our formulas the complexes are graded cohomologically, and not homologically.

Let Λ be a set of objects as before.

Definition 2. A A_∞-category with Λ as a set of objects is the data of $\mathbb{Z}$-graded k-vector spaces $\mathrm{Hom}(X, Y)$ for any pair $X, Y \in \Lambda$ and for every integer $n \geq 1$ and tuple X_i a map

$$\mu^n : \mathrm{Hom}(X_0, X_1) \otimes \mathrm{Hom}(X_1, X_2) \otimes \cdots \otimes Hom(X_{n-1}, X_n) \to \mathrm{Hom}(X_0, X_n)$$

of degree $2 - n$, satisfying the relations

$$\sum_{i+j+k=n} (-1)^{|a_1|+\cdots+|a_k|-k} \mu^{i+1+k}(a_1, \ldots, \mu^{1+i+k}(a_k, \ldots, a_{n-i}), \ldots, a_n) = 0$$

The A_∞-categories we will use are all (strongly) unital, ie. there is a distinguished identity morphism $1_X \in \mathrm{Hom}(X, X)$ of degree zero satisfying appropriate identity relations.

Calabi-Yau structures on A_∞-categories

A Calabi-Yau structure on a A_∞-category is a generalization of a Frobenius structure on an associative algebra. Calabi-Yau structures will play a very prominent role in Chapter 3 of this dissertations; there we will present a more modern and precise definition, in the context of dg categories.

For the current application to A_∞-categories and TCFTs, we will use the following simpler definition. In the terminology we will introduce later, this would correspond to a A_∞ version of a *right* or *proper* Calabi-Yau structure.

Definition 3. A Calabi-Yau structure on a A_∞-category $\mathcal{D}$ is a collection of nondegenerate pairings

$$\mathrm{Hom}(X, Y) \otimes \mathrm{Hom}(Y, X) \xrightarrow{\langle,\rangle} k[-d]$$

which are symmetric (under switching the factors) and compatible with the A_∞ structure maps, satisfying

$$\langle \mu^n(a_1, \ldots, a_{n-1}), a_n \rangle = (-1)^{n+|a_0|\sum |a_i|} \langle \mu^n(a_1, \ldots, a_n), a_0 \rangle$$

for all possible choices of n and morphisms.

In the context of the TCFT, these pairings will be given by what the theory assigns to a cap that closes an incoming state of two strings. The relevance of the definition of Calabi-Yau structure to the study of TCFTs is evident from the main theorem of Costello's first paper on the subject:

Theorem 1. *([30, Thm.A]) The category of open TCFTs of dimension d with set of branes Λ is equivalent to the category of Calabi-Yau A_∞-categories with set of objects Λ.*

Moreover there is a universal way of associating to such a theory an open-closed TCFT and a closed TCFT; this construction uses Hochschild homology of categories, which we will introduce later in Chapter 3.

This theorem motivates the main discussion in the next chapter; constructing Calabi-Yau structures on categories of geometric interest is useful since it allows one to use the formalism of TCFTs to define invariants, in the same way that one can use it to (in principle) calculate higher Gromov-Witten invariants. We will be concerned in the following chapter with constructing Calabi-Yau structures on 'topological Fukaya categories', which are a model for Fukaya categories of non-compact symplectic manifolds. To the writer's knowledge, a further study of the use of TCFT methods in this context has not yet been conducted.

2.3 Liouville manifolds and the wrapped Fukaya category

In this section we will describe some generalizations of the Fukaya category for dealing with non-compact spaces. These categories appear since it has been long understood that the most natural setting of mirror symmetry is not necessarily restricted to sigma models with CY targets; instead one should consider the more general setting of Landau-Ginzburg models which appear as mirrors of non-CY manifolds. We will skip a discussion of Landau-Ginzburg models in this dissertation, since they will not directly appear in later sections, but the interested reader can consult [57, Ch.13,39] for a thorough introduction to their appearance in physics and their mathematical structures.

However, one consequence of this generalization of mirror symmetry to Fano and general type varieties is that, when one tries to apply the Fukaya/coherent sheaf picture of mirror symmetry to cases where the mirror is not proper or smooth, it becomes clear that the ordinary Fukaya category (with compact Lagrangians as objects) cannot be the right A-side category. One way of seeing this is that the derived category of a singular variety has objects with infinite-dimensional endomorphism algebra, which cannot correspond to the endomorphisms of any compact Lagrangian.

A solution to this problem was proposed by Abouzaid and Seidel in [4], in the form of the *wrapped Fukaya category.* This is a category defined for a class of non-compact symplectic manifolds called Liouville manifolds; in particular this includes Weinstein manifolds as an important special case. In this section we will review the general ideas behind the construction of this category.

Wrapped Fukaya categories

Liouville manifolds and their skeleta

The following discussion is taken mostly from [2]; we will use the notations and conventions of that reference. A Liouville manifold is an open symplectic manifold $(M, \partial M, \omega)$ with a one-form λ (the *Liouville form*) such that $d\lambda = \omega$ and the vector field $Z = \omega^{-1}\lambda$ is outward-pointing on the boundary ∂M. We also require that away from a compact set M^{in}, the manifold M can be expressed as

$$M = M^{in} \cup_{\partial M} (\partial M \times [1, \infty))$$

and if we denote by $r \in [1, \infty)$ the coordinate along the cylindrical boundary, we require the proportionality $\lambda = r\lambda|_{\partial M}$ with $\lambda|_{\partial M}$ a contact structure on ∂M. We will call this neighborhood of the boundary the cylindrical end or collar of M.

Given a Liouville manifold (M, λ, Z), consider the subset

$$\mathrm{Sk}(M) = \bigcap_{t<0} \phi^t_Z(X)$$

ie. the attracting set of the negative flow of the vector field Z; we will call this the *skeleton* of the Liouville manifold M.

Example. As a simple example of a Liouville manifold, consider $M = T^*X$ for some compact n-dimensional manifold X, with Liouville form locally given by $\lambda = \sum_i p\ dq_i$ (in Darboux coordinates) and vector field Z given by scaling of the fiber; the skeleton of M is this case is just the zero section X.

In general, the topology of Liouville manifolds can be difficult to understand because the Liouville vector field may have complicated dynamics; one can avoid that by using a Lyapunov function. A *Weinstein manifold* is a Liouville manifold (M, λ, Z) together with a Morse function $h : M \to \mathbb{R}$ such that Z is gradient-like with respect to h. For M a Weinstein manifold, the skeleton of M is composed of the union of the stable manifolds of each zero of the Liouville vector field; and compatibility with the Liouville form implies that each such stable manifold is isotropic.

For a generic Weinstein structure on a manifold, the top-dimensional cells of the skeleton $L = \mathrm{Sk}(M)$ will be middle-dimensional; therefore the skeleton is a (possibly singular)

Lagrangian. We will be working in this case, so we will just refer to L as the *Lagrangian skeleton*. For a more comprehensive exposition of these and other topics in open symplectic geometry, one can consult eg. the reviews [28, 42].

Example. [82] Following the previous example of the cotangent bundle $M = T^*X$, pick a Morse function $f : X \to \mathbb{R}$ and a Riemannian metric g on the base. Consider the function

$$F_X(x, \xi) = \xi(\nabla_g f|_x)$$

One can check that for small enough $\epsilon + 0$, the one-form $\lambda = \lambda_0 + dF_X$, where λ_0 is the standard one-form on T^*X, is a Liouville form compatible with the Morse function (on T^*X)

$$h(x, \xi) = g(\xi, \xi) + f(x)$$

. The skeleton of this is still the zero-section X, but now stratified by the Morse cells coming from the function f (which are isotropic submanifolds of T^*X).

Wrapped Floer cohomology and wrapped Fukaya category

In order to discuss Floer theory with non-compact Lagrangians, one must make some choices about the behavior at infinity. We will fix a Hamiltonian function H on M that is quadratic in the cylindrical end; that is we have $H(y, r) \propto r^2$ in $\partial M \times [1, \infty)$. We will also need to consider almost-complex structures J that are compatible with ω and are of 'contact type' in the collar. For technical reasons we have to make appropriately generic choices of Liouville form, Hamiltonian function and almost-complex structure; it can be shown that the category we will construct does not depend on these choices, as long as they are made appropriately generically.

We will only consider Lagrangians $L \subset M$ that are exact and cylindrical near infinity, ie. in the cylindrical end they are given by

$$\partial L \times [1, \infty) \subset \partial M \times [1, \infty)$$

where ∂L is a Legendrian in the contact manifold ∂M. Moreover, just like in the case of the Fukaya category of compact manifolds, in order for our categories to be $\mathbb{Z}$-graded we need to put further restrictions on the Lagrangian; namely we require that L be spin and that (twice) its relative Chern class vanish: $2c_1(L, M) = 0$.

Given two Lagrangians L_1, L_2 satisfying the conditions above, we consider time 1 chords of the Hamiltonian flow ϕ_H between L_1 and L_2; under the genericity assumptions one can show that these chords are all non-degenerate and the set of such chords is discrete; let us denote this set by $X(L_1, L_2)$.

Under the vanishing conditions on the Lagrangians, the Maslov index defines an integral grading on the space $X(L_1, L_2)$, which we denote by $\deg(x)$. For our purposes, it will suffice to use the simpler description of the set as the transverse intersections between L_1 and L_2 *wrapped* by the time 1 flow under the Hamiltonian vector field, ie. $X(L_1, L_2) = L_1 \cap \phi_H(L_2)$.

For any coefficient field k, we define the *wrapped Floer complex* $CW^*(L_1, L_2)$ to be the k-vector space spanned by the elements of $X(L_1, L_2)$, and graded by the Maslov degree. Analogously to the holomorphic strips defining the differential in the ordinary Lagrangian Floer homology, there is then a notion of a moduli space $R(x_0, x_1)$ of between two elements x_0, x_1, which basically counts pseudo-holomorphic strips. The actual definition of this moduli space is subtle since to ensure transversality one needs to consider families of compatible almost-complex structures instead of fixing one; for the full definition see [2, Sec.2.1].

There is a translation $\mathbb{R}$-action on the space of flows; we will quotient by this action to define $R(x_0, x_1)$. Then for generic choices, $R(x_0, x_1)$ is a compact manifold of dimension $\deg(x_0) - \deg(x_1) - 1$, so when this dimension is zero, counting points defines a differential d of degree $+1$, which can be proven to square to zero, defining the wrapped Floer cohomology $HW^*(L_0, L_1)$.

It can then be shown [2] that there are composition and higher structure maps between wrapped Floer complexes; this defines a A_∞-category $\mathcal{W}(M)$, called the wrapped Fukaya category of M, whose objects are exact Lagrangians, conical at infinity as above, and morphisms are given by the wrapped Floer complex construction.

Example. Consider the cylinder $M = T^*S^1$, with standard exact symplectic form and Hamiltonian given by $H = p^2$, where p is the coordinate along the fiber. Consider the exact Lagrangian L given by the fiber over a point $x \in S^1$.

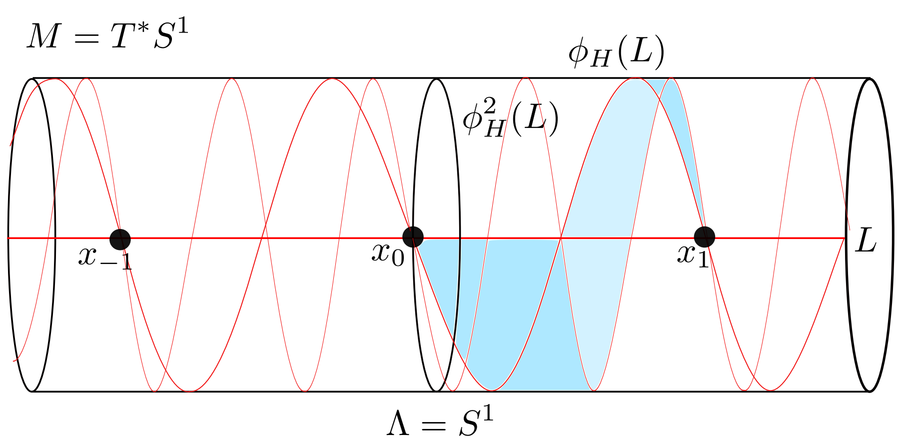

Figure 2.1: Wrapped Fukaya category of the cylinder T^*S^1. The red lines are a Lagrangian L given by the fiber together with its time 1 and 2 flows. The shaded triangle gives a nontrivial product $x_0 \cdot x_1 = x_1$.

The intersections between L and the time 1 flow $\phi_H(L)$ can be indexed by $\mathbb{Z}$, with zero corresponding to the point x on the base. Let us denote then these points by $x_i, i \in \mathbb{Z}$. There are no holomorphic disks (in this case bigons on the surface M) so as a graded vector space the wrapped Floer cohomology of L is given by $CW^*(L, L) = span(\{x_i\})$.

To understand the A_∞ structure, one must also calculate products; the shaded triangle in Figure 2.1 gives a non-trivial product $x_0 \cdot x_1 = x_1$ (see [10, Sec 4.2] for a more detailed explanation), and likewise we get products

$$x_i \cdot x_j = x_{i+j}$$

for any i, j. One can also check that all the higher morphisms μ^n vanish, so as a A_∞-algebra, the wrapped Floer cohomology is given by

$$CW^*(L, L) = k[x, x^{-1}], \quad \deg(x) = 0$$

Since the fiber is the only exact Lagrangian up to Hamiltonian isotopy, it follows that it generates the category $\mathcal{W}(M)$, which is in this case equivalent to the category of A_∞-modules over the A_∞-algebra $CW^*(L, L)$.

In fact, something similar to the example above happens for the cotangent bundle $M = T^*X$ of any closed manifold X. If L is a cotangent fiber, then we have the following result of Abouzaid.

Theorem 2. *[3] There is an isomorphism of A_∞-algebras*

$$CW^*(L, L) \cong C_{-*}(\Omega_* X)$$

between the wrapped Floer complex and chains on the based loop space of the base X.

It has also been proven by Abouzaid [1] that the cotangent fiber generates $\mathcal{W}(M)$, which implies that this category embeds fully faithfully into the category of A_∞-modules over the A_∞-algebra $C_{-*}(\Omega_* X)$.

2.4 Topological Fukaya categories

In 2009, Kontsevich [72] suggested that the wrapped Fukaya category of certain types of Liouville domains could be computed by a sheaf-theoretic calculation, localized to a certain singular Lagrangian. Since this proposal, there has been quite a lot of work [50, 49, 81, 85] (among others) making this proposal precise and proposing explicit models for these categories. In this section we will introduce the part of these definitions and results that will be relevant for us.

These (co)sheaf-theoretic models for Fukaya categories, to which we will refer by the name of *topological Fukaya categories*, actually also provide models for another type of Fukaya-type category one can associate to an open symplectic manifold, the so-called infinitesimal Fukaya category. In this section we will provide a brief introduction to these models; the literature on topological Fukaya categories is by now extensive. We will just mention the main elements that will be necessary to discuss the original results of the work presented in the next chapter, such as infinitesimal vs. wrapped Fukaya categories, microlocal sheaves and arboreal singularities.

Infinitesimal Fukaya categories

In order to introduce the microlocal sheaf models for Fukaya categories, let us first discuss another type of Fukaya category one can associate to some open symplectic manifolds, called the *infinitesimal Fukaya category.* This category appears in the work of Seidel [96] and Nadler-Zaslow [86], among others.

This category is a variant of the category we describe in the previous section; one still considers some kinds of exact Lagrangians as objects, but the difference is roughly that instead of considering a large wrapping, given by a time one flow under a Hamiltonian, one instead wraps *infinitesimally.* An advantage of this category is that because of this infinitesimal deformation, together with compactness restrictions on the wrapping Hamiltonian, there are finitely many intersections between Lagrangians, making their Hom spaces finite-dimensional (which is not the case for the wrapped Fukaya category).

It was eventually understood that these categories can be seen as a special case of a construction of Seidel of a Fukaya-type category for a Lefschetz fibration; we will not need to discuss this so-called Fukaya-Seidel category, but the details about this specific kind of Fukaya-Seidel category (for a Stein manifold) can be found in [96, Sec.III.19].

The cotangent bundle

While the existence of a relation between Fukaya categories and microlocal/constructible sheaves had been suggested by earlier works, the work of Nadler and Zaslow [86] provided the first explicit description of this correspondence, for the case of the Fukaya category of a cotangent bundle. We will summarize the definition given in [86, Sec.5], with some simplifications.

Given a cotangent bundle $M = T^*X$ of dimension $2n$, one considers Hamiltonian functions $H : T^*X \to \mathbb{R}$ that are 'controlled' in the sense that outside of some compact set $K \subset T^*X$, the Hamiltonian is given by $H(x,\xi) = |\xi|$. We also fix some topological data on T^*X, such as a trivialization of the bicanonical class and a pin class.

One can define the infinitesimal Fukaya category $\mathcal{F}(T^*X)$ as a union of smaller categories. Let us fix a conical Lagrangian $\Lambda \subset T^*X$, and define a category $\mathcal{F}(T^*X)_\Lambda$. An object in this category is given by an exact Lagrangian $L \subset T^*X$, appropriately conical at infinity and asymptotic to Λ, equipped with a vector bundle $E \to L$ with flat connection. [1]

Just as before, one can use the Maslov index to define a degree of each intersection point between such Lagrangians. To calculate the morphism space between two objects (L_1, E_1) and (L_2, E_2), one uses the Hamiltonian function to flow the second Lagrangian a 'small amount of time' δ and consider the intersection $L_1 \cap \phi_H^\delta(L_2)$. The morphism space is then generated by those points

$$\mathrm{Hom}_{\mathcal{F}(T^*X)_\Lambda}(L_1, L_2) = \bigoplus_{p \in L_1 \cap \phi_H^\delta(L_2)} \mathrm{Hom}(E_1|_p, E_2|_p)[-\deg(p)]$$

[1] One is also required to specify a 'brane structure' and 'perturbation datum', but this technical discussion won't matter for our exposition.

One can then define composition and higher A_∞ structure maps; for this purpose it is better to allow for more general perturbations than above, see [86, Sec.5.4] for details.

The category $\mathcal{F}(T^*X)$ can be defined as an appropriate union of all these categories; more specifically, for any finite collection of objects, there is a conical Lagrangian such that $\mathcal{F}(T^*X)_\Lambda$ contains them. For convenience let us assume that our conical Lagrangian Λ always contains the zero section X (which may turn it into a singular Lagrangian). Therefore in this convention the smallest choice of conical Lagrangian will be $\Lambda = X$.

Example. Let us return to the setting of Example 2.3 where $X = S^1$. We must fix a Hamiltonian that agrees with $H = |p|$ outside of some compact set; let us then just pick a smooth function $H(p)$ that interpolates between $-p, p < K$ and $p, p > K$ for some positive constant K.

Consider the trivial choice of conical Lagrangian given by the base S^1. Since this is compact, objects of $\mathcal{F} * (T^*X)_{S^1}$ are given by compact Lagrangians equipped with flat vector bundles; since these are all Hamiltonian isotopic to the zero section, each object is equivalent to the circle S^1 with a finite-dimensional local system; so we have an equivalence of categories

$$\mathcal{F}(T^*S^1)_{S^1} \cong \mathrm{Loc}(S^1)$$

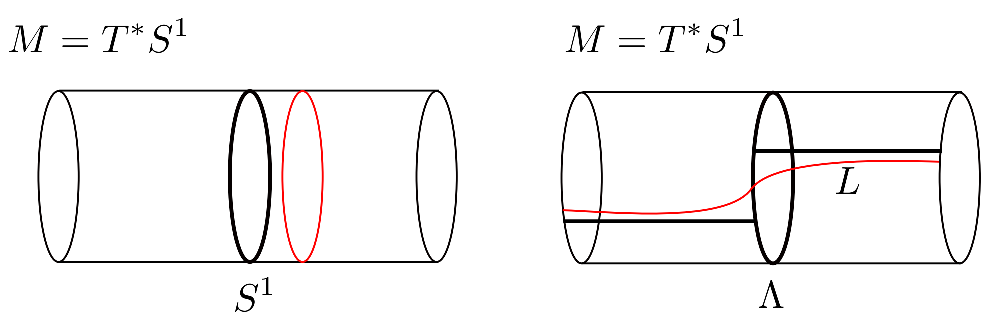

Figure 2.2: Left: the cylinder with Lagrangian skeleton S^1. Every Lagrangian in $\mathcal{F}(T^*S^1)_{S^1}$ is Hamiltonian isotopic to the zero section. Right: the cylinder with Lagrangian skeleton Λ. There are more objects, eg. the Lagrangian L which asymptotes to Λ at infinity

For comparison, consider the case where Λ has two extra half-lines, ie. is the union of the zero section with the cone of over two points, one in each component of $\partial(T^*S^1)$ as in the right side of Figure 2.2. Here there are extra Lagrangians, which wrap around the base S^1 some number of times and asymptote to the two points at infinity. For concreteness, consider one such Lagrangian L that does not wrap around the base. Perturbing it to calculate the infinitesimal Floer complex, we see that there is a single intersection between L and $\phi_H^\delta(L)$. Therefore we have

$$\mathrm{Hom}_{\mathcal{F}(T^*S^1)}(L, \phi_H^\delta(L)) \cong k$$

where we equipped L with a trivial rank one local system.

One can show [85] that there is an equivalence of categories

$$\mathcal{F}(T^*S^1)_\Lambda \cong \mathrm{Mod}(K)$$

where K is the Kronecker quiver $\bullet \rightrightarrows \bullet$. Under this equivalence, the object L above gets sent to the object $k \rightrightarrows 0$, and there is an equivalence between the subcategory $\mathcal{F}(T^*X)_{S^1}$ (given by local systems on the base) and representations of K where both maps are isomorphisms. In particular the trivial rank one object of $\mathcal{F}(T^*X)_{S^1}$ corresponds to $k \underset{1}{\overset{1}{\rightrightarrows}} k$.

Microlocal sheaves

The motivation behind the idea of studying these Fukaya categories by sheaf-theoretic methods lies in the fact that negative time flow of the the Liouville vector field Z contracts the manifold to a lower-dimensional space, namely its skeleton $\mathrm{Sk}(M)$. The fact that the negative flow preserves the important features of Fukaya categories can be seen as the idea behind Kontsevich's [72] initial suggestion to model the (wrapped) Fukaya category in terms of an object supported only on the skeleton L, namely a cosheaf of categories.

The work of Nadler [81, 85, 86] and others [50, 51, 102] has identified categories of *microlocal sheaves* as candidates for these (co)sheaves of categories. Microlocal sheaves are objects of microlocal geometry defined and studied by Kashiwara and Schapira [66]; this formalism packages facts about the study of the local geometry of cotangent bundles into a sheaf-theoretic framework.

We will not attempt to give an exposition of microlocal sheaves in this chapter, since we will not need them for the results of the next chapter; the interested reader can consult the original reference [66] for more details. We will instead just introduce some of the formal properties of the categories of microlocal sheaves, and present the results in the literature that relate them to Fukaya categories. For the more careful reader, the technical terminology relating to dg categories and sheaves of categories can be found in Section 3.2 in the next chapter, where we will introduce the concepts in greater detail.

Here we follow some of the exposition in [85, Sec.3]. Consider the manifold $M = T^*X$, with some closed (possibly singular) Lagrangian subvariety Λ. To any conic (ie. invariant under scaling) open subspace $U \subset T^*X$, there is dg category $\mu\,\mathrm{Sh}_\Lambda(U)$ of microlocal sheaves on U with singular support along Λ; its objects are locally given by constructible sheaves along the base $X \cap U$ with singular support described by the Lagrangian $\Lambda|_U$.

For any inclusion $U' \in U$ of conical open sets, there is a restriction functor $\mu\,\mathrm{Sh}_\Lambda(U) \to \mu\,\mathrm{Sh}_\Lambda(U')$; moreover all these restriction functors assemble into a sheaf of categories. We have the following theorem of Nadler.

Theorem 3. *[83] There is a quasi-equivalence $\mu\,\mathrm{Sh}_\Lambda(X) \xrightarrow{\sim} \mathcal{F}(T^*X)_\Lambda$. Moreover if one considers all possible choices of Λ, there is a quasi-equivalence*

$$\mathrm{Sh}_c(X) \xrightarrow{\sim} \mathcal{F}(T^*X)$$

between constructible sheaves on the base and the infinitesimal Fukaya category.

There is a similar microlocal model for the wrapped Fukaya category, proposed in [85]. There is a *cosheaf* of dg categories $\mu\,\mathrm{Sh}^w_\Lambda$, which Nadler calls *wrapped microlocal sheaves.* Note that due to this being a cosheaf, there are corestriction maps $\mathcal{W}(U) \to \mathcal{W}(V)$ for inclusions $U \hookrightarrow V$ of conical open sets.

Example. Consider again the example $M = T^*S^1$, and let us pick $\Lambda = S^1$ to be the zero-section. Locally, on a contractible open set U, the cosections $\mu\,\mathrm{Sh}^w_\Lambda(U)$ are given by the category Mod_k of dg vector spaces (just like the sections of the sheaf $\mu\,\mathrm{Sh}_\Lambda$). However, over the whole circle, the cosections are given by

$$\mu\,\mathrm{Sh}^w_\Lambda(S^1) \cong \mathrm{Perf}(k[x, x^{-1}])$$

ie. the category of dg $k[x, x^{-1}]$-modules that are perfect over the algebra $k[x, x^{-1}]$ itself.

This category is bigger than the category $\mathrm{Mod}^{(fd)}(k[x, x^{-1}]) \cong \mathrm{Loc}(S^1)$ of finite-dimensional (over k) modules. For example, the module given by $k[x, x^{-1}]$ over itself is not finite-dimensional over k, but it is a perfect object over $k[x, x^{-1}]$.

Note that in the example above the global cosections agrees with our calculation of the wrapped Fukaya category. The correspondence between wrapped microlocal sheaves and the wrapped Fukaya category conjectured by Nadler has recently been proven in the work of Ganatra, Pardon and Shende in the form of the following result, which is stated for more general choices of Λ, which give a *partially wrapped* version of the Fukaya category. [2]

Theorem 4. *([50, Thm.1.1]) Let Λ be the Lagrangian skeleton of a Liouville domain M. The triangulated completion of the (partially) wrapped Fukaya category $\mathcal{W}(M)$ is equivalent to the global cosections $\mu\,\mathrm{Sh}^w_\Lambda(M)$.*

2.5 Arboreal singularities

Arboreal singularities were introduced by Nadler [81, 84], with the goal of providing combinatorial models for the microlocal sheaf categories discussed above. Arboreal singularities make up a particularly simple class of Legendrian (or Lagrangian, depending on the chosen embedding) singularities.

We will below give a combinatorial description of arboreal singularities, following Nadler's definitions but also establishing some calculations that were presented in [100]. First let us recall the following result by Nadler [84], which establishes the algorithm for expanding general singularities into arboreal singularities. This result is presented in terms of Legendrian singularities. Let S^*X denote the spherical projectivization of the tangent bundle T^*X; this is a contact manifold, and consider a singular Legendrian $\Lambda \subset S^*X$. We are interested

[2]Note that this result applies more generally than to cotangent bundles, which is the context we introduced microlocal sheaves earlier

in 'non-characteristic deformations' of Λ, ie. deformations that induce equivalences of the categories of microlocal sheaves. Under mild assumptions on Λ, Nadler proves that:

Theorem 5. *[84, Thm.1.1] There is an algorithm that produces a non-characteristic deformation that takes Λ to a singular Legendrian Λ_{arb} with only arboreal singularities.*

Combined with results cited previously, this suggests that one can always reduce the calculation of the wrapped Fukaya category to microlocal sheaves on locally arboreal spaces; recent results on the relation between such non-characteristic deformations and Lagrangian/Legendrian geometry can be found in the work of Starkson [105].

Each arboreal singularity corresponds to a tree (nonempty, finite, connected and acyclic graph) and one can construct an explicit local model for the singularity, realizing it as a singular Legendrian. In this section we recall definitions and results of [81].

Gluing construction of arboreal singularities

Let T be a tree, and $V(T), E(T)$ its sets of vertices and edges, with $n = |V(T)|$. For each vertex $\alpha \in V(T)$, let us denote by

$$\mathbb{T}(\alpha) = \mathbb{R}^{V(T)\setminus\{\alpha\}}$$

the Euclidean space of dimension $n-1$ with coordinates $x_\gamma(\alpha), \gamma \neq \alpha$.

For each edge $\alpha - \beta$, we define an equivalence relation on $\mathbb{T}(\alpha) \sqcup \mathbb{T}(\beta)$ given by identifying the points with coordinates $\{x_\gamma(\alpha)\}$ and $\{x_\gamma(\beta)\}$ when

$$x_\beta(\alpha) = x_\alpha(\beta) \geq 0 \qquad \text{and} \qquad x_\gamma(\alpha) = x_\gamma(\beta), \quad \gamma \neq \alpha, \beta$$

In words, this equivalence relation glues the Euclidean spaces $\mathbb{T}(\alpha)$ and $\mathbb{T}(\beta)$ along the respective half-spaces $\{x_\beta(\alpha) \geq 0\}$ and $\{x_\alpha(\beta) \geq 0\}$.

The arboreal singularity $\mathbb{T}$ associated to the tree T is the quotient space

$$\mathbb{T} = \left. \bigsqcup_{\alpha \in V(T)} \mathbb{T}(\alpha) \middle/ \sim \right.$$

where $\sim$ denotes the equivalence relation generated by all the edges in $E(T)$.

As defined above, the arboreal singularity $\mathbb{T}$ is a singular space of top dimension $n-1$. It is easy to see that this space admits some natural stratifications, coming from this gluing construction. In [81], Nadler gives a way of associating each stratum of this space with a *correspondence of trees.*

Trees and correspondences

Here we will also use correspondences of trees to define a stratification on the arboreal singularity $\mathbb{T}$. However we will proceed in a different way from the original construction of Nadler, giving this stratification in more functorial terms.

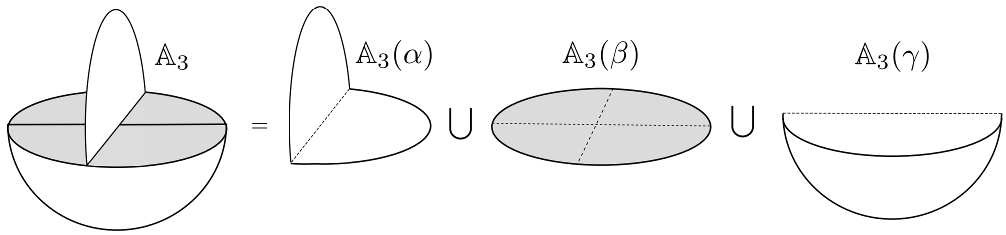

Figure 2.3: Gluing of $\mathbb{A}_3$ from the discs (or equivalently, Euclidean planes) $\mathbb{A}_3(\bullet)$.

A correspondence of trees $\mathfrak{p}$ is a diagram

$$R \overset{q}{\twoheadleftarrow} P \overset{i}{\hookrightarrow} T$$

where R, P, T are trees, and the maps q, i are respectively surjective and injective maps of graphs. An isomorphism of correspondences $R \overset{q}{\twoheadleftarrow} P \overset{i}{\hookrightarrow} T$ and $R' \overset{q'}{\twoheadleftarrow} P' \overset{i'}{\hookrightarrow} T'$ is the data of isomorphisms $R \cong R', P \cong P', T \cong T'$ which intertwine the maps.

Given trees and maps $Q \hookrightarrow R \twoheadleftarrow S$, the fiber product graph $Q \times_R S$ is the subtree of S whose vertices map to the image of Q. Thus correspondences can be composed:

$$(P \twoheadleftarrow Q \hookrightarrow R) \circ (R \twoheadleftarrow S \hookrightarrow T) = (P \twoheadleftarrow Q \times_R S \hookrightarrow T)$$

Definition 4. The category *Arb* has correspondences of trees as its objects. The Hom sets are

$$\mathrm{Hom}(\mathfrak{p}', \mathfrak{p}) := \{\mathfrak{q} \,|\, \mathfrak{p} = \mathfrak{q} \circ \mathfrak{p}'\}/\text{isomorphism}$$

Composition is given by composition of correspondences, which makes sense since given $\mathfrak{q} \in \mathrm{Hom}(\mathfrak{p}', \mathfrak{p})$ and $\mathfrak{q}' \in \mathrm{Hom}(\mathfrak{p}'', \mathfrak{p}')$, we have $\mathfrak{p} = \mathfrak{q} \circ \mathfrak{p}' = \mathfrak{q} \circ \mathfrak{q}' \circ \mathfrak{p}''$.

Lemma 6. *Arb is a poset – i.e., for any correspondences $\mathfrak{p}', \mathfrak{p}$ there is, up to isomorphism, at most one $\mathfrak{q}$ such that $\mathfrak{p} = \mathfrak{q} \circ \mathfrak{p}'$.*

Proof. Suppose $(P \twoheadleftarrow Q \hookrightarrow R) \circ (R \twoheadleftarrow S \hookrightarrow T) = (P \twoheadleftarrow N \hookrightarrow T)$. We want to reconstruct $(P \twoheadleftarrow Q \hookrightarrow R)$ from just $(R \twoheadleftarrow S \hookrightarrow T)$ and $(P \twoheadleftarrow N \hookrightarrow T)$.

The map $N = Q \times_R S \to S$ determined by taking the (necessarily unique) factorization of $N \hookrightarrow T$ as $N \hookrightarrow S \hookrightarrow T$. From this we can characterize Q as the image of N under the map $S \twoheadrightarrow R$. The map $Q \twoheadrightarrow P$ is determined by the (necessarily unique) factorization of $N \twoheadrightarrow P$ into $N \twoheadrightarrow Q \twoheadrightarrow P$. □

Definition 5. Let T be a tree. We write Arb_T for the subcategory of Arb of correspondences $R \twoheadleftarrow S \hookrightarrow T$, or equivalently, for the subcategory of objects admitting a map from $\mathfrak{p}_T = (T \twoheadleftarrow T \hookrightarrow T)$.

Functorial construction of the arboreal singularity

Recall that the nerve of a category C is the simplicial complex whose vertices are the objects of C, morphisms are the edges, triangles are commuting triangles giving compositions, etc. When C is a poset, this is also called the order complex.

Definition 6. If T is a tree, we write $\mathbb{T}$ for the nerve of the category Arb_T. This is a stratified space, the *arboreal singularity* of T that will serve as our local model for locally arboreal spaces. We write $\mathbb{T}^{int}$ for the union of simplices containing $\mathfrak{p}_T$, and $\mathbb{T}^{\mathrm{link}}$ for the complement of this union.

The space $\mathbb{T}$ is conical; the initial object $\mathfrak{p}_T \in \mathrm{Arb}_T$ gives the cone point and $\mathbb{T}^{\mathrm{link}}$ gives the link. It follows from the results of [81] that the definition of $\mathbb{T}$ above agrees with Nadler's gluing construction we presented above, at the level of homeomorphism type. The advantage of presenting the arboreal singularity as the nerve of a category is that it will allow us to easily construct sheaves on the arboreal singularity by giving functors out of the category Arb, as we do below in Section 3.4.

Example. We write A_2 for the tree $\bullet - \bullet$. We label the vertices α and β. There are four correspondences: the trivial correspondence $\mathfrak{p}_0 = (\alpha - \beta) \twoheadleftarrow (\alpha - \beta) \hookrightarrow (\alpha - \beta)$, the correspondence $\bullet \twoheadleftarrow (\alpha - \beta) \hookrightarrow (\alpha - \beta)$, and two correspondences $\bullet \twoheadleftarrow \bullet \hookrightarrow (\alpha - \beta)$ for inclusions of α or β.

We abbreviate these by enclosing in parenthesis those vertices of A_2 which get identified in the quotient $R \twoheadleftarrow S$. So, for example, we will denote the three nontrivial correspondences by $\alpha, \beta, (\alpha\beta)$ and the trivial correspondence simply by $\alpha\beta$.

In the poset structure, the three nontrivial correspondences are incomparable, and the correspondence $\mathfrak{p}_0$ is smaller than all of them. Thus there are 7 strata in the order complex $\mathbb{A}_2$: the four 0-simplices $[\alpha\beta], [\alpha], [\beta], [(\alpha\beta)]$, and three 1-simplices $[\alpha\beta \to \alpha], [\alpha\beta \to \beta], [\alpha\beta \to (\alpha\beta)]$. This can be realized as the following stratified space of dimension one (Figure 2.4). Note that $\mathbb{A}_2$ is the disjoint union of 3 points, each labeled by a 0-simplex $[\mathfrak{q}] \neq [\mathfrak{p}_0]$

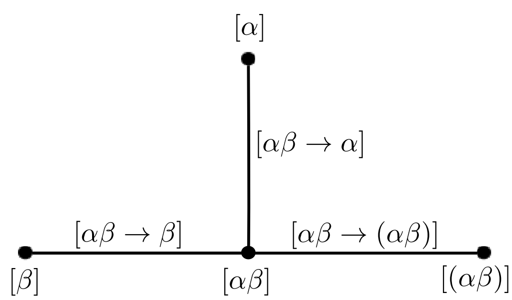

Figure 2.4: Arboreal singularity $\mathbb{A}_2$. For simplicity we use the notation described above for each correspondence

Example. We write A_3 for the tree $\bullet - \bullet - \bullet$, whose vertices we label $\alpha - \beta - \gamma$. There are eleven correspondences: the trivial one, four correspondences of the form $[\bullet \to \bullet] \twoheadleftarrow \dots$ and six of the form $[\bullet] \twoheadleftarrow \dots$. There are 45 strata in the order complex $\mathbb{A}_3$: 11 zero-dimensional, 22 one-dimensional and 12 two-dimensional strata, assembled as in Figure 2.5.

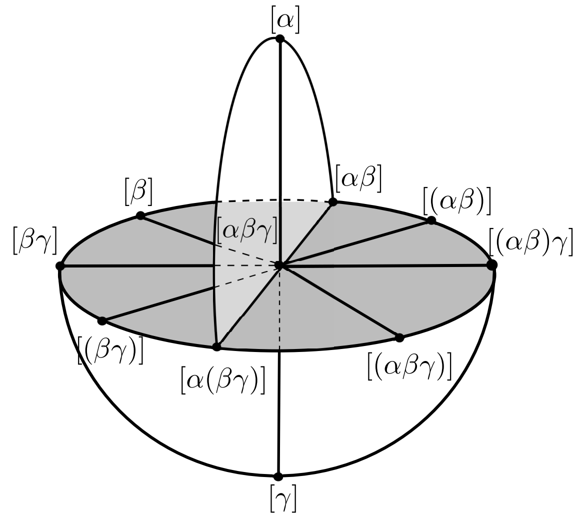

Figure 2.5: Arboreal singularity $\mathbb{A}_3$. For simplicity we label only the 0-simplices; the labels on all other simplices can be deduced from their vertices

The link $\mathbb{A}_3^{link}$ can be glued out of copies of $\mathbb{A}_2$. Note also that the realization of this space admits a coarser stratification, with 5 zero-dimensional, 10 one-dimensional and 6 two-dimensional strata.

Chapter 3

Calabi-Yau structures on topological Fukaya categories

In this chapter, we will present the results and proofs of [100], which were results obtained by the author of this dissertation jointly with his academic advisor. That paper develops a technique for constructing Calabi-Yau structures on categories that admit a description in terms of global sections of certain constructible (co)sheaves of categories. Thus we will first develop some local-to-global methods for constructing Calabi-Yau structures, and then apply them in the context of interest.

As presented in the previous chapter, Calabi-Yau categories are generalizations of Frobenius algebras suitable to capture the algebraic structure of enriched 2d topological field theories. Among the examples of Calabi-Yau categories one finds: categories of sheaves on Calabi-Yau algebraic varieties (the topological B-model), categories of Lagrangians in symplectic manifolds (the topological A-model) [88], and cluster-tilted categories in algebra [70].

Such categories give rise to a host of numerical invariants: the partition function, or Gromov-Witten type invariants [31], and (for CY3 categories) the BPS, or the Donaldson-Thomas type invariants [73]. Moreover, their moduli spaces of objects carry shifted symplectic structures and quantizations thereof [21, 91, 25].

Our goal here is to further expand the inventory of Calabi-Yau categories to include the topological Fukaya categories as defined in the previous chapter, using the microlocal model for these categories. The motivation for taking these combinatorial models for the Fukaya category is twofold. One is that it is easy to make computations and prove structural results. For example, the global cosections of a constructible cosheaf on a finite type space are computed by a homotopy finite colimit; hence for such a cosheaf whose stalk categories are finite, the global cosections are finite as well.

Secondly, the combinatorial description, or rather its relation to categories of sheaves with prescribed microsupport, often allows the identification of (topological) Fukaya categories with categories of interest elsewhere in mathematics. In particular, many authors have formed moduli spaces from variously decorated surfaces, and constructed symplectic or

Poisson structures on these spaces. These include wild character varieties [16, 17, 18], multiplicative Nakajima varieties [114, 15], the cluster varieties "from surfaces" [52] and their higher rank generalizations [92, 43]. It is explained in [103, 102] that all of these constructions can be realized naturally as moduli spaces of sheaves with prescribed microsupport. Another sort of example comes from the augmentation variety of knot contact homology [41]; the relation to sheaves with prescribed microsupport is explained in [40].

3.1 Summary of results

As mentioned in Section 2.2, the data of a Calabi-Yau structure on a category $\mathcal{F}$ is a collection of pairings; these can be seen as an isomorphism of the diagonal bimodule with a shift of its dual. In fact there are two such notions, corresponding to two ways of dualizing the diagonal. The first, called a proper or right Calabi-Yau structure, corresponds to isomorphism of the diagonal bimodule with its shifted linear dual. This is often discussed in terms of the corresponding system of isomorphisms $\mathrm{Hom}(x,y) \cong \mathrm{Hom}(y,x)^*[-d]$. The other variant, called a smooth or left Calabi-Yau structure, corresponds to an isomorphism between the diagonal bimodule and a shift of its bimodule dual. Both these structures can be specified in terms of a trace on or element of the Hochschild complex $\mathrm{HH}(\mathcal{F})$, this being the (derived) tensor product of the diagonal bimodule of $\mathcal{F}$ with itself, and the state space associated to the circle by the topological field theory. In fact one must therefore require that everything respect the natural circle action on the Hochschild complex; we will denote by $\mathrm{HH}(\mathcal{F})^{S_1}$ the homotopy S^1-invariants (fixed points) and by $\mathrm{HH}(\mathcal{F})_{S^1}$ the homotopy S^1-coinvariants (orbits) with respect to this action.

From a constructible sheaf $\mathcal{F}$ of categories on $\mathbb{X}$, we wish to produce from local data a map $\mathrm{HH}(\mathcal{F}(\mathbb{X})) \to k[-d]$; from a constructible cosheaf of categories $\mathcal{F}$ on $\mathbb{X}$, we wish to produce from local data a map $k[d] \to \mathrm{HH}(\mathcal{F}(\mathbb{X}))$.

Hochschild homology is functorial, so $U \mapsto \mathrm{HH}(\mathcal{F}(U))$ determines a presheaf. In general it is not a sheaf: when $\mathcal{C}$ is the constant sheaf of categories, $\mathrm{HH}(\mathcal{F}(U))$ can be identified with cochains on the free loop space of U. Nevertheless we may form the sheafification $\mathcal{HH}(\mathcal{F})$ of this presheaf; the resulting sheaf carries a suitable S^1 action inherited from the local S^1 actions. We then seek a local construction of a morphism $\mathcal{HH}(\mathcal{F})(\mathbb{X}) \to k[-d]$. This sheaf also carries a suitable S^1 action inherited from the local S^1 actions. By composition with the natural morphism $\mathrm{HH}(\mathcal{F}(\mathbb{X})) \to \mathcal{HH}(\mathcal{F})(\mathbb{X})$, such a morphism would determine an orientation. The local data which integrates to such a map is precisely a morphism to the Verdier dualizing sheaf. Moreover, it is possible to formulate a local version of the condition that the resulting trace induces a perfect pairing.

Definition 7. (Local orientations) Let $\mathbb{X}$ be a stratified topological space of dimension d; we write $\omega_{\mathbb{X}}$ for its Verdier dualizing sheaf. Let $\mathcal{F}$ be a sheaf of categories on $\mathbb{X}$. Then an orientation of $(\mathbb{X}, \mathcal{F})$ is a morphism

$$\mathcal{HH}(\mathcal{F})_{S^1} \to \omega_{\mathbb{X}}[-d]$$

For an open subset $U \subset \mathbb{X}$, and objects $x, y \in \mathcal{F}(U)$, we write $\mathcal{F}_\Delta(x,y)$ for the constructible sheaf of morphisms on U. An orientation induces a morphism

$$\mathcal{F}_\Delta(x,y) \to \mathscr{H}om_U(\mathcal{F}_\Delta(y,x), \omega_U)[-d]$$

We say the orientation is non-degenerate if this morphism is an isomorphism; this is only possible if the stalks of $\mathcal{F}$ are proper categories, in the sense that its morphism spaces are perfect complexes of k-modules.

Example. Let M be a manifold and k a field. Let $\mathcal{L}oc$ be the sheaf of categories of local systems of k-vector spaces on M. Then $\mathcal{HH}(\mathcal{L}oc) = k_M$. An orientation of $(M, \mathcal{L}oc)$ is a choice of isomorphism $k_M \cong \omega_M[-\dim M]$, i.e., an orientation of M in the sense of topology.

We show (Prop. 19) that if $\mathbb{X}$ is compact, a non-degenerate local orientation indeed induces a proper Calabi-Yau structure on $\mathcal{F}(\mathbb{X})$, and more generally (Prop. 20) that if the noncompactness of $\mathbb{X}$ is exhausted by $\partial\mathbb{X}$, then $\mathcal{F}(\mathbb{X}) \to \mathcal{F}(\partial\mathbb{X})$ carries a relative right Calabi-Yau structure.

Often, proper categories arise as the finite dimensional module categories over some category with infinite dimensional Hom spaces – e.g., finite rank $\mathbb{C}$-local systems on M as representations of the algebra of chains $C_*(\Omega M)$, generally infinite-dimensional over $\mathbb{C}$. In the cases of interest, these infinite dimensional preduals have a different sort of finiteness: they are homologically smooth, in the sense that the diagonal bimodule is perfect as a bimodule. In this case one can ask for a cotrace inducing an isomorphism of the diagonal bimodule with its bimodule dual. Such is called a *smooth* or *left* Calabi-Yau structure. It is more fundamental, in particular inducing a right Calabi-Yau structure on the finite-dimensional-module category. There is also a relative version of this notion [20]. If the sheaf $\mathcal{F}$ is the pseudo-perfect modules over a cosheaf of smooth categories $\mathcal{W}$, we show that an orientation of $\mathcal{F}$ also induces smooth Calabi-Yau structures on sections of $\mathcal{W}$:

Theorem 7. *Let $(\mathbb{X}, \partial\mathbb{X})$ be a stratified space of dimension d with compact boundary $\partial\mathbb{X}$, with a locally saturated constructible cosheaf of smooth categories $\mathcal{W}$ and its sheaf of pseudo-perfect modules $\mathcal{W}^{pp}$. Then a non-degenerate local orientation on $\mathcal{W}^{pp}$ induces a relative proper (or right) Calabi-Yau structure of dimension d on the restriction*

$$\mathcal{W}^{pp}(\mathbb{X}) \to \mathcal{W}^{pp}(\partial\mathbb{X})$$

together with a compatible (absolute) proper Calabi-Yau structure of dimension $(d-1)$ on $\mathcal{W}^{pp}(\partial\mathbb{X})$, and a relative smooth (or left) Calabi-Yau structure of dimension d on the corestriction

$$\mathcal{W}(\partial\mathbb{X}) \to \mathcal{W}(\mathbb{X})$$

together with a compatible (absolute) smooth Calabi-Yau structure of dimension $(d-1)$ on $\mathcal{W}(\partial\mathbb{X})$.

In [112], the authors construct moduli spaces of objects in dg categories with suitable finiteness conditions; which are satisfied for a homotopy-finite colimit of smooth proper categories, hence for the global sections of our cosheaves. While the moduli construction takes as input a category like $\mathcal{W}(\mathbb{X})$, the resulting space $\mathcal{M}(\mathcal{W}(\mathbb{X}))$ parameterizes objects in $\mathcal{W}^{pp}(\mathbb{X})$. In particular, the inclusion of the boundary $\partial\mathbb{X} \to \mathbb{X}$ gives a map $\mathcal{W}(\partial\mathbb{X}) \to \mathcal{W}(\mathbb{X})$ hence $\mathcal{W}(\mathbb{X})^{pp} \to \mathcal{W}(\partial\mathbb{X})^{pp}$ and correspondingly on moduli $\mathcal{M}(\mathcal{W}(\mathbb{X})) \to \mathcal{M}(\mathcal{W}(\partial\mathbb{X}))$. If $\mathcal{W}$ carries a relative orientation, then, as we have shown, the map $\mathcal{W}(\partial\mathbb{X}) \to \mathcal{W}(\mathbb{X})$ is (left) Calabi-Yau.

By the main theorem of [21], this implies that the corresponding map on moduli $\mathcal{M}(\mathcal{W}(\mathbb{X})) \to \mathcal{M}(\mathcal{W}(\partial\mathbb{X}))$ is a derived Lagrangian morphism in the sense of [91, 25]. From such a result we can deduce:

Corollary 8. *Let $(\mathbb{X}, \partial\mathbb{X})$ be a stratified space of dimension d with compact boundary $\partial\mathbb{X}$, with a locally saturated constructible cosheaf of smooth categories $\mathcal{W}$ and its sheaf of pseudo-perfect modules $\mathcal{W}^{pp}$. Then a non-degenerate local orientation on the sheaf of proper categories $\mathcal{W}^{pp}$ gives a $(3-n)$-shifted symplectic structure on the moduli space $\mathcal{M}_{\mathcal{W}(\partial\mathbb{X})}$ parametrizing objects in $\mathcal{W}^{pp}(\partial\mathbb{X})$ and a Lagrangian structure on the morphism*

$$\mathcal{M}_{\mathcal{W}(\mathbb{X})} \to \mathcal{M}_{\mathcal{W}(\partial\mathbb{X})}$$

corresponding to restriction of objects in $\mathcal{W}^{pp}(\mathbb{X})$ to the boundary. In the case where $\partial\mathbb{X} = \emptyset$ this gives a $(2-d)$-shifted symplectic structure on the moduli of global objects $\mathcal{M}_{\mathcal{W}(\mathbb{X})}$.

The main motivation and example will be the co/sheaves which microlocal sheaf theory associates to a (singular) Legendrian. We restrict attention to the 'arboreal' setting of [81]. As we survey in Section 3.6, many categories of interest arise in this setting. In fact, by the results of [84], results in the arboreal setting suffice to treat the general case; we will not however explain the reduction in detail.

The arboreal singularities of [81] are a certain class of local models of Legendrian singularities. The microlocal sheaf theory equips the underlying topological spaces with certain constructible co/sheaves of categories. The resulting co/sheaves can be built directly from the representation theory of tree quivers.

We adopt here the latter combinatorial point of view. In brief, to each tree T is attached a certain stratified topological space $\mathbb{T}$ such that each stratum is labelled by a tree and each attaching map is labelled by a correspondence of trees. There is a constructible cosheaf of smooth categories $\mathcal{W}_{\mathbb{T}}$ on $\mathbb{T}$, whose stalks are quiver representation categories, and whose cogenerization maps are given in terms of correspondences of quivers. By construction, the cosheaf $\mathcal{W}_{\mathbb{T}}$ is locally saturated and exact. We will also be interested in the sheaf of categories formed by its pseudo-perfect modules $\mathcal{W}_{\mathbb{T}}^{pp}$; this is a sheaf of proper categories. We will show:

Theorem 9. *Let $\vec{T}$ be a rooted tree, and $\mathbb{T}$ the corresponding arboreal singularity. Then $\mathcal{HH}(\mathcal{W}^{pp})$ and the dualizing complex $\omega_{\mathbb{T}}[-\dim \mathbb{T}]$ are isomorphic. The isomorphism is unique up to a scalar, and induces a non-degenerate orientation on $\mathcal{W}^{pp}$.*

Let us now indicate how these notions globalize.

Definition 8. A locally arboreal space is a stratified space $\mathbb{X}$ equipped with a locally saturated and exact cosheaf of dg categories $\mathcal{W}$ such that $(\mathbb{X}, \mathcal{W})$ is locally modelled on $(\mathbb{T} \times \mathbb{R}^n, \mathcal{W}_{\mathbb{T}})$.

Thus if $(\mathbb{X}, \mathcal{W})$ is a locally arboreal space, the obstruction to its global orientability is the nontriviality of the rank one local system $\mathcal{H}om(\mathcal{HH}(\mathcal{W}^{pp}), \omega_{\mathbb{X}}[-\dim \mathbb{X}])$. This is classified by the corresponding element of $H^1(\mathbb{X}, k^*)$. We will show that this is in fact an element of $H^1(\mathbb{X}, \pm 1)$, which we term the first Stiefel-Whitney class of the locally arboreal space $(\mathbb{X}, \mathcal{W})$. When this obstruction vanishes, a choice of isomorphism $\mathcal{HH}(\mathcal{W}^{pp}) \cong \omega_{\mathbb{X}}[-\dim \mathbb{X}]$ determines a relative left Calabi-Yau structure on $\mathcal{W}(\partial \mathbb{X}) \to \mathcal{W}(\mathbb{X})$, and a right Calabi-Yau structure on $\mathcal{W}^{pp}(\mathbb{X}) \to \mathcal{W}^{pp}(\partial \mathbb{X})$.

3.2 DG categories

The categories of microlocal sheaves appearing in this model are *differential graded categories*; here we will first specify the scope of our definitions, relying on notations and definitions of [48], and will later review some facts relevant to us, following in some parts [20, 21].

We will use the $(\infty, 1)$-categories of dg categories introduced in [48, Sec. 10.3]. Following that reference, we denote by $\mathrm{DGCat}_{\mathrm{cont}}$ the $(\infty, 1)$-category whose objects are presentable complete and cocomplete (ie. containing all limits and colimits) dg categories over k and whose morphisms are (exact) continuous (ie. colimit-preserving) functors. This category is endowed with a symmetric monoidal structure given by a tensor product $\otimes$, and an internal Hom functor right-adjoint $\otimes$ which we will denote by $\underline{\mathrm{Hom}}$. We have another category DGCat of presentable dg categories with the same objects, but whose morphisms include non-continuous (exact) functors.

We also denote by dgcat the category of small, idempotent-complete dg categories over k, with morphisms given by exact functors. Taking compact objects gives a functor $(-)^c$: $\mathrm{DGCat}_{\mathrm{cont}} \to \mathrm{dgcat}$, and taking ind-completion gives an adjoint functor $\mathrm{Ind} : \mathrm{dgcat} \to \mathrm{DGCat}_{\mathrm{cont}}$. If $\mathcal{C}$ is a compactly generated category [48, Def. 7.1.3], then there is a canonical equivalence $\mathrm{Ind}(\mathcal{C}^c) \to \mathcal{C}$.

For ease of notation we will call by 'dg category' an object of $\mathrm{DGCat}_{\mathrm{cont}}$ and by 'small dg category' and object of dgcat.

Limits and colimits of dg categories

Throughout this chapter and the next, we will need to take many limits and colimits of dg categories; fortunately the following construction (rephrased from [48, Sec. 7.2.6] in terms of dg categories) allows us to freely exchange colimits for limits and vice-versa, provided we take them in the appropriate category.

Let I be an index category and $\mathcal{C}_I : I \to \mathrm{DGCat}_{\mathrm{cont}}$ be a diagram of dg categories. Suppose that each one of the functors $\mathcal{C}_i \to \mathcal{C}_j$ also preserves compact objects; this gives rise to a diagram

$$\mathcal{C}_I^c : I \to \mathrm{dgcat}, \quad i \to \mathcal{C}_i^c$$

Taking right adjoints to the morphisms $\mathcal{C}_i \to \mathcal{C}_j$ gives us a diagram

$$\mathcal{C}_I^R : I^{op} \to \mathrm{DGCat}$$

where the morphisms may not be continuous. The following result means that the colimit of small dg categories can be computed in terms of a *limit* of dg categories.

Lemma 10. *[48, Cor. 7.2.7] There is a canonical equivalence* $\mathrm{Ind}\,(\mathrm{colim}_I\, \mathcal{C}_I^c) \to \lim_{I^{op}} \mathcal{C}_I^R$

Constructible sheaves and cosheaves of categories

A $\mathcal{C}$-valued presheaf on a stratified space $\mathbb{X}$ is a functor $\mathcal{F} : \mathrm{Opens}(\mathbb{X})^{op} \to \mathcal{C}$; it is a sheaf if carries covers to limits; ie. for any open set U and cover $U = \bigcup_i U_i$ the natural map $\mathcal{F}(U) \to \lim_i \mathcal{F}(U_i)$ is an equivalence. Likewise, a $\mathcal{C}$-valued cosheaf is a functor $\mathcal{F} : \mathrm{Opens}(\mathbb{X}) \to \mathcal{C}$; it is a cosheaf if it carries covers to colimits, ie. the natural map $\mathrm{colim}_i\, \mathcal{F}(U_i) \to \mathcal{F}(U)$ is an equivalence for any cover. When $\mathcal{C}$ is an $(\infty, 1)$-category, the functor, limits and colimits should all be understood in the $(\infty, 1)$-categorical sense.

In the rest of this dissertation, by *(co)sheaf of small categories* we mean a (co)sheaf valued in the $(\infty, 1)$-category dgcat of small idempotent-complete dg categories with exact functors. As a consequence of Lemma 10, if we consider instead the category DGCat of large dg categories we can freely exchange between limits and colimits. Inspired by this we make the following definition.

Definition 9. (Co/sheaves of categories) A co/sheaf of categories $\mathcal{F}$ is a sheaf valued in the $(\infty, 1)$-category DGCat of dg categories, whose restriction maps $\mathcal{F}(U) \to \mathcal{F}(V)$ for any inclusion of opens $U \subseteq V$ preserve limits and colimits.

It is a general fact that if a functor ρ preserves limits and colimits, it has left adjoint ρ^ℓ which preserves compact objects. Using this fact together with Lemma 10 we have the following equivalent definition, which motivates the name "co/sheaf".

Lemma 11. *The data of a co/sheaf of categories is the same as the data of a cosheaf valued in* $\mathrm{DGCat}_{\mathrm{cont}}$ *of dg categories with continuous functors, whose corestriction functors moreover preserve compact objects.*

Given any co/sheaf of categories $\mathcal{F}$ we can restrict to compact objects to get a cosheaf of small categories $\mathcal{F}^c$. Conversely, given a cosheaf of small categories $\mathcal{W}$ we can ind-complete to get a co/sheaf $\mathrm{Ind}\,\mathcal{W}$. There is a canonical equivalence of cosheaves of small categories $(\mathrm{Ind}\,\mathcal{W})^c \to \mathcal{W}$.

The sheaf of pseudo-perfect modules over a cosheaf of small categories

Let Perf denote the full subcategory of Vect whose objects are the perfect complexes of vector spaces. This is small category, ie. an object of the $(\infty, 1)$-category dgcat.

Definition 10. For any small dg category $\mathcal{C}$, its category of pseudo-perfect modules $\mathcal{C}^{pp}$ is the object of dgcat defined by

$$\mathcal{C}^{pp} := \underline{\mathrm{Hom}}(\mathcal{C}^{op}, \mathrm{Perf})$$

By [48, Lem. 10.5.6], the ind-completion of DG categories can be described by modules: there is a canonical equivalence $\mathrm{Ind}\,\mathcal{C} \xrightarrow{\sim} \mathrm{Mod}\text{-}\mathcal{C} = \underline{\mathrm{Hom}}(\mathcal{C}, \mathrm{Vect})$. Therefore $\mathcal{C}^{pp}$ can be identified with a full subcategory of $\mathrm{Ind}\,\mathcal{C}$.

Let us now apply this definition to sheaves and cosheaves. We assume the regularity condition that $\mathbb{X}$ can be stratified by finitely many strata, and that this stratification can be refined to a finite simplicial complex. We fix this stratification from now on, and restrict attention to sheaves and cosheaves constructible with respect to it.

Definition 11. Let $\mathbb{X}$ be a stratified space and $\mathcal{W}$ a cosheaf of small categories on $\mathbb{X}$, with corestriction maps $\iota_U^V : \mathcal{W}(U) \to \mathcal{W}(V)$ for $U \subseteq V$. Then the sheaf of pseudo-perfect modules $\mathcal{W}^{pp}$ is a sheaf of categories assigning

$$U \mapsto \mathcal{W}(U)^{pp}$$

with restriction maps given by the pullback $(\iota_U^V)^*$.

$\mathcal{W}^{pp}$ is easily seen to be a sheaf, since for a cover $U = \cup_i U_i$ we have

$$\mathcal{W}^{pp}(U) = \underline{\mathrm{Hom}}(\mathcal{W}(U), \mathrm{Perf}_k) \cong \underline{\mathrm{Hom}}(\operatorname*{colim}_i \mathcal{W}(U_i), \mathrm{Perf}_k) \cong \lim_i \underline{\mathrm{Hom}}(\mathcal{W}(U_i), \mathrm{Perf}_k) = \lim_i \mathcal{W}^{pp}(U_i)$$

One can also identify the sheaf of pseudo-perfect modules $\mathcal{W}^{pp}$ with a subsheaf of co/sheaf $\mathrm{Ind}\,\mathcal{W}$. Let us now assume that every category of sections $\mathcal{W}(U)$ is a smooth category.

Lemma 12. *[112] If $\mathcal{A}$ is a smooth dg category, then every pseudo-perfect module over $\mathcal{A}$ is also perfect (i.e. compact in* $\mathrm{Mod}\text{-}\mathcal{A}$*). Conversely, if $\mathcal{A}$ is proper, then every perfect module over $\mathcal{A}$ is pseudo-perfect.*

The Yoneda embedding gives a canonical fully faithful functor $\mathcal{C} \to \mathrm{Ind}\,\mathcal{C}$, whose essential image lies in $(\mathrm{Ind}\,\mathcal{C})^c$ when $\mathcal{C}$ is any stable category. By general facts, the Yoneda map is an equivalence to its image when $\mathcal{C}$ is idempotent-complete; therefore for any $\mathcal{C} \in$ dgcat we have a canonical equivalence $\mathcal{C} \xrightarrow{\sim} (\mathrm{Ind}\,\mathcal{C})^c = \mathrm{Perf}\text{-}\mathcal{C}$. Thus when all the sections $\mathcal{W}(U)$ are smooth categories, for any open set U we have a fully faithful embedding $\mathcal{W}^{pp}(U) \to \mathcal{W}(U)$. Note that this is *not* a map of either sheaves or cosheaves, since it maps sections of a sheaf to cosections of a cosheaf. By the lemma above, this is an equivalence if and only if $\mathcal{W}(U)$ is smooth.

Definition 12. Let $\mathcal{W}$ be a cosheaf of small categories on $\mathbb{X}$. We say that $\mathcal{W}$ is locally saturated if all stalks are smooth and proper.

Lemma 13. *Let $\mathcal{W}$ be a locally saturated cosheaf of small categories. Then the maps $\mathcal{W}^{pp}(U_\epsilon) \to \mathcal{W}(U_\epsilon)$ are quasi-isomorphisms for all small enough open sets U_ϵ. Also, for* any *open set U, $\mathcal{W}(U)$ is smooth and finite type and $\mathcal{W}^{pp}(U)$ is proper.*

Remark. Note that by the regularity condition that our stratification can be refined to a finite simplicial complex, each open simplex in that complex is contractible and therefore "small enough" in the sense that all the stalks in it are isomorphic to the sections over that open simplex.

Proof. By constructibility around any point there is a small open U_ϵ such that the stalk is isomorphic to $\mathcal{W}(U_\epsilon)$; together with Lemma 12 this proves the first assertion. Moreover, any open set U can be covered by finitely many such small open sets U_ϵ (eg. neighborhoods of the open simplices of the finite simplicial refinement); the lemma follows from the fact that finite colimits of smooth categories are smooth and finite limits of proper categories are proper. □

3.3 Calabi-Yau structures and local orientations

Calabi-Yau structures on dg categories

Let us denote by Vect the category of dg vector spaces over k; this is the unit object in $\mathrm{DGCat}_{\mathrm{cont}}$ for the tensor product $\otimes$.

Modules and bimodules

We define the (underived) categories of right $\mathcal{A}$-modules, left $\mathcal{A}$-modules and $\mathcal{A},\mathcal{B}$-bimodules by using the internal Hom functor in $\mathrm{DGCat}_{\mathrm{cont}}$:

$$\mathrm{mod}\text{-}\mathcal{A} = \underline{\mathrm{Hom}}(\mathcal{A}^{op}, \mathrm{Vect}), \qquad \mathcal{A}\text{-}\mathrm{mod} = \underline{\mathrm{Hom}}(\mathcal{A}, \mathrm{Vect}), \qquad \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} = \underline{\mathrm{Hom}}(\mathcal{A}\otimes\mathcal{B}^{op}, \mathrm{Vect})$$

We will denote by $\mathrm{Mod}\text{-}\mathcal{A} := \mathcal{D}(\mathrm{mod}\text{-}\mathcal{A})$ etc. the corresponding derived categories of modules. These categories are 'categories of presheaves' valued in Vect, so by a general fact of category theory are both complete and cocomplete. Moreover all limits and colimits can be computed objectwise, ie. taking the limit or colimit commutes with evaluation of the (bi)module at a fixed object of $\mathcal{A}$.

We will say that an object P of $\mathrm{Mod}\text{-}\mathcal{A}$ is perfect or compact if the functor $\mathrm{Mod}\text{-}\mathcal{A} \to \mathrm{Vect}$ given by $M \mapsto \mathrm{Hom}_{\mathrm{Mod}\text{-}\mathcal{A}}(P, M)$ commutes with filtered colimits (equivalently, direct sums thus and all colimits). We will denote as $\mathrm{Perf}\text{-}\mathcal{A}$ the full dg subcategory of $\mathrm{Mod}\text{-}\mathcal{A}$ spanned by such objects, and analogously we can define $\mathcal{A}\text{-}\mathrm{Perf}$ as a full dg subcategory of $\mathcal{A}\text{-}\mathrm{Mod}$.

There is a tensor product $\otimes_{\mathcal{B}} : \mathcal{A}\text{-}\operatorname{mod}\text{-}\mathcal{B} \otimes \mathcal{B}\text{-}\operatorname{mod}\text{-}\mathcal{C} \to \mathrm{Vect}$ pairing right and left modules; this gives rise to a left derived tensor product $\otimes^L_{\mathcal{B}} : \mathcal{A}\text{-}\operatorname{Mod}\text{-}\mathcal{B} \otimes \mathcal{B}\text{-}\operatorname{Mod}\text{-}\mathcal{C} \to \mathrm{Vect}$ pairing derived categories.

We will need to compute such derived functors; for that we can use any appropriate model structure on categories of modules over dg categories such as the one given in [111]. We will most leave references to the model structure implicit unless we need to use it explicitly.

Diagonal bimodule

We have a distinguished bimodule $\mathcal{A}_\Delta \in \operatorname{Mod}\text{-}\mathcal{A}^e$ given by morphisms in the category $\mathcal{A}$:

$$\mathcal{A}_\Delta(a, a') = \operatorname{Hom}_{\mathcal{A}}(a, a')$$

Here we denote $\mathcal{A}^e = \mathcal{A} \otimes \mathcal{A}^{op}$. We will alternatively denote by $\mathcal{A}_\Delta$ the equivalent objects in $\mathcal{A}^e\text{-}\operatorname{Mod}$ and $\mathcal{A}\text{-}\operatorname{Mod}\text{-}\mathcal{A}$. Tensoring over $\mathcal{A}$ with the diagonal bimodule also gives the identity morphism on any category of bimodules, ie.

$$- \otimes^L_{\mathcal{A}} \mathcal{A} : \mathcal{B}\text{-}\operatorname{Mod}\text{-}\mathcal{A} \to \mathcal{B}\text{-}\operatorname{Mod}\text{-}\mathcal{A}$$

is canonically equivalent to the identity functor.

Hochschild complex

We write

$$\mathrm{HH}(\mathcal{A}) := \mathcal{A}_\Delta \otimes^L_{\mathcal{A}^e} \mathcal{A}_\Delta \in \mathrm{Vect}$$

for the Hochschild complex of $\mathcal{A}$, where the $\mathcal{A}_\Delta$ are viewed either as elements of $\operatorname{Mod}\text{-}\mathcal{A}^e$ and $\mathcal{A}^e\text{-}\operatorname{Mod}$. We will denote by $\mathrm{HH}_n(\mathcal{A}) = H^n(\mathrm{HH}(\mathcal{A}))$ the Hochschild homology of $\mathcal{A}$.

The derived tensor product above can be computed by taking any projective resolution of bimodules $P_{\mathcal{A}} \to \mathcal{A}$; in particular one can take the bar resolution [77]. Tensoring over $\mathcal{A}^e$ with the diagonal bimodule gives then the *reduced bar complex* $\overline{C}^{\mathrm{bar}}(\mathcal{A})$ representing $\mathrm{HH}(\mathcal{A})$. This complex carries a canonical homotopy S^1 action, whose homotopy orbits $\mathrm{HH}(\mathcal{A})_{S^1}$ and fixed points $\mathrm{HH}(\mathcal{A})^{S^1}$ were classically termed the cyclic and negative cyclic complexes, with homology groups denoted by

$$\mathrm{HC}_n(\mathcal{A}) = H^*(\mathrm{HH}(\mathcal{A})_{S^1}), \quad \mathrm{HC}^-_n(\mathcal{A}) = H^*(\mathrm{HH}(\mathcal{A})^{S^1})$$

Explicit representatives for these complexes can be given by using the formalism of mixed complexes [77]; we refer the reader to [115, Sec. 2.4] for an explicit application of this formalism to the context of dg categories.

One can also consider Hochschild homology with coefficients in an arbitrary $\mathcal{A},\mathcal{A}$-bimodule $\mathcal{M}$, defined by

$$\mathrm{HH}(\mathcal{A}, \mathcal{M}) := \mathcal{A}_\Delta \otimes_{\mathcal{A}^e} \mathcal{M} \in \mathrm{Vect}$$

.

Linear and bimodule duals

Consider the linear duality functor $\mathrm{Hom}_k(-, k)$ on the dg category Vect. For any $\mathcal{M}$ in $\mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B}$, we define its *linear dual*

$$\mathcal{M}^* = \mathrm{Hom}_k(\mathcal{M}, k)$$

which is an object of $\mathcal{B}\text{-}\mathrm{mod}\text{-}\mathcal{A}$. Explicitly, as a functor $\mathcal{M}^* : \mathcal{B} \otimes \mathcal{A}^{op} \to \mathrm{Vect}$, it is given by

$$(b, a) \mapsto \mathrm{Hom}_k(\mathcal{M}(a, b), k)$$

For any bimodule $\mathcal{M}$ there is an (underived) adjunction

$$- \otimes_{\mathcal{A}\otimes\mathcal{B}^{op}} \mathcal{M} : \mathcal{B}\text{-}\mathrm{mod}\text{-}\mathcal{A} \rightleftarrows \mathrm{Vect} : \mathrm{Hom}_k(\mathcal{M}, -)$$

$\mathcal{M}$ is called linear-dualizable (or right-dualizable) if the natural transformation $- \otimes_k \mathcal{M}^* \to \mathrm{Hom}_k(\mathcal{M}, -)$ is an equivalence of functors. Equivalently, $\mathcal{M}$ is linear-dualizable if it always evaluates to a perfect k-complex, i.e. $\mathcal{M}(a, b) \in \mathrm{Perf}_k$ for every $(a, b) \in \mathcal{A} \otimes \mathcal{B}^{op}$. In that case, there is a canonical isomorphism of bimodules $\mathcal{M} \xrightarrow{\simeq} (\mathcal{M}^*)^*$, so we also get another adjunction

$$\mathcal{M}^* \otimes_{\mathcal{A}\otimes\mathcal{B}^{op}} - : \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} \rightleftarrows \mathrm{Vect} : \mathcal{M} \otimes_k -$$

Taking linear duals is naturally contravariant; this defines a functor $(-)^* : \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} \to (\mathcal{B}\text{-}\mathrm{mod}\text{-}\mathcal{A})^{op}$. Since k is an injective object, this functor is exact and therefore defines a functor

$$(-)^* : \mathcal{A}\text{-}\mathrm{Mod}\text{-}\mathcal{B} \to (\mathcal{B}\text{-}\mathrm{Mod}\text{-}\mathcal{A})^{op}$$

between the derived categories. When $\mathcal{M}$ is linear-dualizable, one can check that the adjunctions above are Quillen adjunctions so they also give adjunctions of derived categories

$$- \otimes^L_{\mathcal{A}\otimes\mathcal{B}^{op}} \mathcal{M} : \mathcal{B}\text{-}\mathrm{Mod}\text{-}\mathcal{A} \rightleftarrows \mathrm{Vect} : - \otimes^L_k \mathcal{M}^*$$
$$\mathcal{M}^* \otimes^L_{\mathcal{A}\otimes\mathcal{B}^{op}} - : \mathcal{A}\text{-}\mathrm{Mod}\text{-}\mathcal{B} \rightleftarrows \mathrm{Vect} : \mathcal{M} \otimes^L_k -$$

We will define another kind of dual: consider the *bimodule dual* of $\mathcal{M}$ defined using the internal Hom of bimodules

$$\mathcal{M}^\vee = \mathrm{Hom}_{\mathcal{A}\otimes\mathcal{B}^{op}}(\mathcal{M}, \mathcal{A}_\Delta \otimes_k \mathcal{B}_\Delta)$$

which is an object of $\mathcal{B}\text{-}\mathrm{mod}\text{-}\mathcal{A}$. Explicitly, this is given by

$$\mathcal{M}^\vee(b, a) = \mathrm{Hom}_{\mathcal{A}\otimes\mathcal{B}^{op}}(\mathcal{M}(-, -'), \mathcal{A}_\Delta(-, a) \otimes_k \mathcal{B}_\Delta(-', b))$$

For any bimodule $\mathcal{M}$ there is an adjunction

$$\mathcal{M} \otimes_k - : \mathrm{Vect} \rightleftarrows \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} : \mathrm{Hom}_{\mathcal{A}\otimes\mathcal{B}^{op}}(\mathcal{M}, -)$$

$\mathcal{M}$ is called bimodule-dualizable (or left-dualizable) if the natural transformation $\mathcal{M}^\vee \otimes_{\mathcal{A}\otimes\mathcal{B}^{op}} - \to \mathrm{Hom}_{\mathcal{A}\otimes\mathcal{B}^{op}}(\mathcal{M}, -)$ is an equivalence of functors. Equivalently $\mathcal{M}$ is bimodule-dualizable if it is perfect as a bimodule (ie. a compact object in that category). In that case we get a canonical isomorphism of bimodules $\mathcal{M} \xrightarrow{\simeq} (\mathcal{M}^\vee)^\vee$, and we also have another adjunction

$$- \otimes_k \mathcal{M}^\vee : \mathrm{Vect} \rightleftarrows \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} : - \otimes_{\mathcal{A}\otimes\mathcal{B}^{op}} \mathcal{M}$$

It is easy to see that taking bimodule duals is natually contravariant; this defines a functor $(-)^\vee : \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} \to (\mathcal{B}\text{-}\mathrm{mod}\text{-}\mathcal{A})^{op}$. However, in this case this functor is only *right-exact*; to mark the difference let us then denote by

$$(-)^! : \mathcal{A}\text{-}\mathrm{Mod}\text{-}\mathcal{B} \to (\mathcal{B}\text{-}\mathrm{Mod}\text{-}\mathcal{A})^{op}$$

its left derived functor between the derived categories. Explicitly, this can be computed by taking any projective bimodule resolution $P_\mathcal{M} \to \mathcal{M}$ and taking the object represented by $P_\mathcal{M}^\vee$. The adjunctions above can also be checked to give Quillen adjunctions [20, Sec.3] so in the case that $\mathcal{M}$ is bimodule-dualizable we get adjunctions of derived categories

$$\mathcal{M} \otimes_k^L - : \mathrm{Vect} \rightleftarrows \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} : \mathcal{M}^! \otimes^L_{\mathcal{A}\otimes\mathcal{B}^{op}} -$$
$$- \otimes_k^L \mathcal{M}^! : \mathrm{Vect} \rightleftarrows \mathcal{A}\text{-}\mathrm{mod}\text{-}\mathcal{B} : - \otimes^L_{\mathcal{A}\otimes\mathcal{B}^{op}} \mathcal{M}$$

Remark. These duals are also called in the literature respectively *right dual* and *left dual*, or respectively *proper dual* and *smooth dual*; this comes from looking at $\mathcal{M}$ as an object of $\mathcal{A}\otimes\mathcal{B}^{op}\text{-}\mathrm{Mod}\text{-}k$. However since we will already have too many objects labelled by the words right and left we will use the terminology linear and bimodule dual to avoid confusion.

Properness and smoothness

A dg category $\mathcal{A}$ is said to be *proper* if all the Hom spaces are perfect as k-complexes, which is equivalent to the diagonal bimodule $\mathcal{A}_\Delta$ being linear-dualizable. In this case the dual of Hochschild homology can be computed in terms of the linear dual $\mathcal{A}^*_\Delta$ (by adjunction):

$$\mathrm{RHom}_k(\mathrm{HH}(\mathcal{A}), k) = \mathrm{RHom}_k(\mathcal{A}_\Delta \otimes^L_{\mathcal{A}^e} \mathcal{A}_\Delta, k) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A}_\Delta, \mathcal{A}^*_\Delta)$$

A dg category $\mathcal{A}$ is said to be *smooth* if the diagonal bimodule is perfect as a module over $\mathcal{A}^e$, or equivalently bimodule-dualizable. In this case the Hochschild homology of $\mathcal{A}$ can be calculated in terms of $\mathcal{A}^!_\Delta$:

$$\mathrm{HH}(\mathcal{A}) = \mathrm{RHom}_k(k, \mathcal{A}_\Delta \otimes^L_{\mathcal{A}^e} \mathcal{A}_\Delta) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A}^!_\Delta, \mathcal{A}_\Delta)$$

When the category $\mathcal{A}$ is both proper and smooth, these bimodules give endofunctors on $\mathrm{Perf}\text{-}\mathcal{A}$, and moreover the functors $- \otimes_\mathcal{A} \mathcal{A}^!_\Delta$ and $- \otimes_\mathcal{A} \mathcal{A}^*_\Delta$ are inverse autoequivalences. [20].

The discussion above also holds for coefficients in an arbitrary $\mathcal{A},\mathcal{A}$-bimodule $\mathcal{M}$. That is, if $\mathcal{A}$ is proper there is a canonical isomorphism

$$\mathrm{RHom}_k(\mathrm{HH}(\mathcal{A},\mathcal{M}),k) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{M},\mathcal{A}^*_\Delta)$$

and if $\mathcal{A}$ is smooth there is a canonical isomorphism

$$\mathrm{HH}(\mathcal{A},\mathcal{M}) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A}^!_\Delta,\mathcal{M})$$

.

Calabi-Yau structures

Definition 13. A d-dimensional proper (or right) Calabi-Yau structure on a proper dg category $\mathcal{A}$ is a map of complexes

$$\mathrm{HH}(\mathcal{A})_{S^1} \to k[-d]$$

so that the induced morphism in $\mathrm{RHom}(\mathrm{HH}(\mathcal{A}),k[-d]) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A},\mathcal{A}^*[-d])$ is an isomorphism.

A d-dimensional smooth (or left) Calabi-Yau structure on a smooth dg category $\mathcal{A}$ is a map

$$k[d] \to \mathrm{HH}(\mathcal{A})^{S^1}$$

so that the induced morphism in $\mathrm{RHom}(k[d],\mathrm{HH}(\mathcal{A})) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A}^![d],\mathcal{A})$ is an isomorphism.

Remark. Note that we require these maps to be compatible with the S^1 action; we can equivalently look at the data of a proper CY structure as a dual cyclic homology class $[\xi^*] \in \mathrm{Hom}_k(\mathrm{HC}_d(\mathcal{A}),k)$, and the data of a smooth CY structure as a negative cyclic homology class $[\xi] \in \mathrm{HC}^-_d(\mathcal{A})$. The resulting CY structures can be shown to be independent of the choice of representatives ξ^*,ξ; this is shown explicitly in [115] for the smooth CY structures and the argument there can be easily adapted to the proper case.

Suppose that $\mathcal{A}$ is smooth, and $\mathcal{P}$ is a full dg subcategory spanned by a set of locally proper objects, i.e. for any objects $p \in \mathcal{P}$ and $a \in \mathcal{A}$, the Hom space is a perfect complex: $\mathcal{A}(a,p) \in \mathrm{Perf}_k$. Then it can be proven [20] that a smooth CY structure on $\mathcal{A}$ automatically gives a proper CY structure on $\mathcal{P}$. Explicitly, consider the functor

$$\mathcal{D} : \mathcal{A}\text{-}\mathrm{Mod}\text{-}\mathcal{A} \to (\mathcal{P}\text{-}\mathrm{Mod}\text{-}\mathcal{P})^{op}$$

which to an $\mathcal{A},\mathcal{A}$-bimodule M associates the $\mathcal{P},\mathcal{P}$-bimodule

$$(p,q) \mapsto \mathrm{RHom}_{\mathcal{A}}(M(q,-),\mathcal{A}_\Delta(p,-))$$

Proposition 14. *The functor $\mathcal{D}$ maps $\mathcal{A}_\Delta$ to $\mathcal{P}_\Delta$, $\mathcal{A}^!_\Delta$ to $\mathcal{P}^*_\Delta$, and the induced map*

$$\mathrm{HH}(\mathcal{A}) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A}^!_\Delta, \mathcal{A}_\Delta) \to \mathrm{RHom}_{\mathcal{P}^e}(\mathcal{P}_\Delta, \mathcal{P}^*_\Delta) = \mathrm{RHom}_k(\mathrm{HH}(\mathcal{P}), k)$$

is compatible with the S^1 action and takes smooth Calabi-Yau structures to proper Calabi-Yau structures. Moreover if $\mathcal{P} \cong \mathcal{A}$ is smooth and proper, then the functor $\mathcal{D}$ is an auto-equivalence.

Remark. Note that if $\mathcal{A}$ is not both smooth and proper, in general one cannot reverse the procedure above and get a smooth CY structure on $\mathcal{A}$ from a proper CY structure on some proper subcategory $\mathcal{P}$.

Relative Calabi-Yau structures

Consider a dg functor $f : \mathcal{A} \to \mathcal{B}$ between two dg categories. This gives functors between the module categories: writing $F = f \otimes f^{op} : \mathcal{A}^e \to \mathcal{B}^e$ there is an adjunction of (derived) functors

$$F_! : \mathcal{A}\text{-}\mathrm{Mod}\text{-}\mathcal{A} \rightleftarrows \mathcal{B}\text{-}\mathrm{Mod}\text{-}\mathcal{B} : F^*$$

The functor F^* is easy to calculate explicitly: the *underived* functor given by sending a $(\mathcal{B}, \mathcal{B})$-bimodule $\mathcal{N}$ to the bimodule

$$(a, a') \mapsto \mathcal{N}(f(a), f(a'))$$

is already exact, so the (derived) functor F^* is precisely given by this construction. Its left adjoint $F_!$ is more complicated: for a $\mathcal{A}, \mathcal{A}$-bimodule $\mathcal{M}$ we have

$$F_!\mathcal{M} = \mathcal{M} \otimes^L_{\mathcal{A}^e} F^*(\mathcal{B}^e) = (f^{op})^*(\mathcal{B}_\Delta) \otimes^L_{\mathcal{A}} \mathcal{M} \otimes^L_{\mathcal{A}} f^*(\mathcal{B}_\Delta)$$

In the last term, $(f^{op})^*(\mathcal{B}_\Delta)$ is an object of $\mathcal{B}\text{-}\mathrm{Mod}\text{-}\mathcal{A}$ given by pulling back the diagonal bimodule $\mathcal{B}_\Delta$ on the right, i.e.

$$(f^{op})^*(\mathcal{B}_\Delta)(b, a) = \mathcal{B}_\Delta(b, f(a))$$

and $f^*(\mathcal{B}_\Delta)$ is an object of $\mathcal{A}\text{-}\mathrm{Mod}\text{-}\mathcal{B}$ given by pulling back $\mathcal{B}_\Delta$ on the left, i.e.

$$f^*(\mathcal{B}_\Delta)(a, b) = \mathcal{B}_\Delta(f(a), b)$$

We can compute this derived functor by taking any projective resolution $P_\mathcal{M} \to \mathcal{M}$ and calculating $P_\mathcal{M} \otimes_{\mathcal{A}^e} F^*(\mathcal{B}^e)$ using the underived tensor product.

There is a natural 'unit' morphism of $(\mathcal{A}, \mathcal{A})$-bimodules $u : \mathcal{A}_\Delta \to F^*\mathcal{B}_\Delta$, which on pairs of objects (a, a') maps

$$\mathcal{A}(a, a') \mapsto \mathcal{B}(f(a), f(a'))$$

By adjunction this gives a 'counit' morphism of $(\mathcal{B}, \mathcal{B})$-bimodules $c : F_!\mathcal{A}_\Delta \to \mathcal{B}$

These functors also interact nicely with the linear and bimodule duals we defined above: for any bimodule $\mathcal{N} \in \mathcal{B}\text{-}\operatorname{Mod}\text{-}\mathcal{B}$ we have a canonical equivalence

$$F^*(\mathcal{N}^*) = (F^*\mathcal{N})^*$$

Conversely, for a bimodule-dualizable $\mathcal{M} \in \mathcal{A}\text{-}\operatorname{Mod}\text{-}\mathcal{A}$ one can use the adjunctions above to show that there is a canonical equivalence

$$F_!(\mathcal{M}^!) = (F_!\mathcal{M})^!$$

.

A functor $f : \mathcal{A} \to \mathcal{B}$ induces a map of Hochschild complexes $f_\sharp : \mathrm{HH}(\mathcal{A}) \to \mathrm{HH}(\mathcal{B})$, compatible with the S^1 action. We define the relative Hochschild complex

$$\mathrm{HH}(f) := \mathrm{Cone}(\mathrm{HH}(\mathcal{A}) \to \mathrm{HH}(\mathcal{B}))$$

An explicit representative for this complex can be given by $\mathrm{HH}(\mathcal{A})[1] \oplus \mathrm{HH}(\mathcal{B})$ with differential given by the matrix

$$d_{\mathrm{HH}(f)} = \begin{pmatrix} d_{\mathrm{HH}(\mathcal{A})} & f_\sharp \\ 0 & d_{\mathrm{HH}(\mathcal{B})} \end{pmatrix}$$

Similarly, we define the relative cyclic complex $\mathrm{HH}_{S^1}(f) = \mathrm{Cone}(\mathrm{HH}(\mathcal{A})_{S^1} \to \mathrm{HH}(\mathcal{B})_{S^1})$ and relative negative cyclic complex $\mathrm{HH}^{S^1}(f) = \mathrm{Cone}(\mathrm{HH}(\mathcal{A})^{S^1}(\mathcal{A}) \to \mathrm{HH}(\mathcal{B})^{S^1})$.

Let us first discuss the case where $\mathcal{A}$ and $\mathcal{B}$ are *proper* categories. An element $\phi \in \mathrm{RHom}_k(\mathrm{HH}(f), k[-d+1])$ is determined up to homotopy by the data of dual cycles

$$\phi_\mathcal{B} : \mathrm{HH}(\mathcal{B}) \to k[-d+1], \qquad \phi_\mathcal{A} : \mathrm{HH}(\mathcal{A}) \to k[-d]$$

such that $\phi_\mathcal{A}$ gives a null-homotopy of $\phi_\mathcal{B} \circ f_\sharp$, ie. satisfying the equations

$$\phi_\mathcal{B} \circ d_{\mathrm{HH}(\mathcal{B})} = 0, \qquad \phi_\mathcal{A} \circ d_{\mathrm{HH}(\mathcal{A})} + \phi_\mathcal{B} \circ f_\sharp = 0$$

Properness implies that we have isomorphisms

$$\mathrm{RHom}_k(\mathrm{HH}(\mathcal{A}), k) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A}_\Delta, \mathcal{A}_\Delta^*), \qquad \mathrm{RHom}_k(\mathrm{HH}(\mathcal{B}), k) = \mathrm{RHom}_{\mathcal{B}^e}(\mathcal{B}_\Delta, \mathcal{B}_\Delta^*)$$

Let us fix any two projective resolutions $P_\mathcal{A} \to \mathcal{A}_\Delta$ and $P_\mathcal{B} \to \mathcal{B}_\Delta$. The element ϕ gives then two morphisms of bimodules

$$\phi_\mathcal{A} : P_\mathcal{A} \to \mathcal{A}_\Delta^*[-d], \qquad \phi_\mathcal{B} : P_\mathcal{B} \to \mathcal{B}_\Delta^*[-d+1]$$

We can pull this last map back using F^* and compose it with (appropriate lifts of) the unit map $u : \mathcal{A}_\Delta \to F^*\mathcal{B}_\Delta$ and its dual $u^* : F^*\mathcal{B}_\Delta^* \to \mathcal{A}_\Delta^*$ to get a morphism

$$P_\mathcal{A} \xrightarrow{u} F^*P_\mathcal{B} \xrightarrow{F^*\phi_\mathcal{B}} F^*\mathcal{B}_\Delta[-d+1] \xrightarrow{(-1)^{-d+1}u^*} A_\Delta^*[-d+1]$$

One can check that the condition $\phi_{\mathcal{A}} \circ d_{\mathcal{A}} + \phi_{\mathcal{B}} \circ f_\sharp = 0$ is equivalent to the condition that the corresponding morphism of bimodules $\phi_{\mathcal{A}}$ gives a nullhomotopy to the composition above. By functoriality, this gives a morphism of distinguished triangles in $\mathcal{A}\text{-}\mathrm{Mod}\text{-}\mathcal{A}$:

$$\begin{array}{ccccc} P_{\mathcal{A}} & \xrightarrow{u} & F^*P_{\mathcal{B}} & \longrightarrow & \mathrm{Cone}(u) \\ \downarrow & & \downarrow{\scriptstyle F^*\phi_{\mathcal{B}}} & & \downarrow \\ \mathrm{Cone}((-1)^{-d+1}u^*)[-1] & \longrightarrow & F^*\mathcal{B}^*_\Delta[-d+1] & \xrightarrow{(-1)^{-d+1}u^*} & \mathcal{A}^*_\Delta[-d+1] \end{array}$$

Let us now discuss the case where $\mathcal{A}$ and $\mathcal{B}$ are smooth. A degree d element of $\mathrm{HH}(f)$ is determined up to homotopy by the data of cycles

$$\xi_{\mathcal{A}} : k[d-1] \to \mathrm{HH}(\mathcal{A}), \qquad \xi_{\mathcal{B}} : k[d] \to \mathrm{HH}(\mathcal{B})$$

such that $\xi_{\mathcal{B}}$ gives a null-homotopy of $f_\sharp \circ \xi_{\mathcal{A}}$, ie. satisfying the equations

$$d_{\mathrm{HH}(\mathcal{A})} \circ \xi_{\mathcal{A}} = 0, \quad d_{\mathrm{HH}(\mathcal{B})} \circ \xi_{\mathcal{B}} + f_\sharp \circ \xi_{\mathcal{A}} = 0$$

Smoothness gives us isomorphisms $\mathrm{HH}(\mathcal{A}) = \mathrm{RHom}_{\mathcal{A}^e}(\mathcal{A}^!_\Delta, \mathcal{A}_\Delta)$ and $\mathrm{HH}(\mathcal{B}) = \mathrm{RHom}_{\mathcal{B}^e}(\mathcal{B}^!_\Delta, \mathcal{B}_\Delta)$ and then the data above gives morphisms (after lifting to the projective representatives)

$$\xi_{\mathcal{A}} : P^!_{\mathcal{A}}[d-1] \to P_{\mathcal{A}}, \qquad \xi_{\mathcal{B}} : P^!_{\mathcal{B}}[d] \to P_{\mathcal{B}}$$

Using the derived pushforward $F_!$ and (appropriate representatives of) the counit map $c : F_!\mathcal{A}_\Delta \to B_\Delta$ and its bimodule dual $u^!$ we get a morphism

$$P^!_{\mathcal{B}}[d-1] \xrightarrow{(-1)^{d-1}c^!} F_!P^!_{\mathcal{A}}[d-1] \xrightarrow{F_!\xi_{\mathcal{A}}} F_!P_{\mathcal{A}} \xrightarrow{c} P_{\mathcal{B}}$$

and it can be checked [115, Prop 2.26] that $\xi_{\mathcal{B}}$ gives a nullhomotopy of this composition and that we get a morphism of distinguished triangles

$$\begin{array}{ccccc} P^!_{\mathcal{B}}[d-1] & \xrightarrow{(-1)^{d-1}c^!} & F_!P^!_{\mathcal{A}}[d-1] & \longrightarrow & \mathrm{Cone}((-1)^{d-1}c^!) \\ \downarrow & & \downarrow{\scriptstyle F_!\xi_{\mathcal{A}}} & & \downarrow \\ \mathrm{Cone}(c)[-1] & \longrightarrow & F_!P_{\mathcal{A}} & \xrightarrow{c} & P_{\mathcal{B}} \end{array}$$

It can be checked that these two morphisms of distinguished triangles are independent (up to quasi-isomorphism) of the choice of (dual) cycle ϕ, ξ representing a given (dual) class and projective resolutions.

Definition 14. (Relative Calabi-Yau structures)[20] A d-dimensional proper (or right) relative Calabi-Yau structure on dg functor $f : \mathcal{A} \to \mathcal{B}$ between proper dg categories is a dual cyclic class

$$[\phi] \in (\mathrm{HC}_{d-1}(f))^\vee = H^0(\mathrm{Hom}_k(\mathrm{HH}(f)_{S^1}, k[-d+1]))$$

such that the induced map of distinguished triangles in $\mathcal{A}\text{-}\operatorname{Mod}\text{-}\mathcal{A}$ (first of the two above) is an isomorphism.

A d-dimensional proper (or right) relative Calabi-Yau structure on dg functor $f : \mathcal{A} \to \mathcal{B}$ between smooth dg categories is a negative cyclic class

$$[\xi] \in \mathrm{HC}^{-}_{d}(f)$$

such that the induced map of distinguished triangles in $\mathcal{B}\text{-}\operatorname{Mod}\text{-}\mathcal{B}$ (second of the two above) is an isomorphism.

Sheaves, cosheaves, bimodules and Hochschild homology

Our goal is to use sheaves and cosheaves of categories to construct absolute and relative Calabi-Yau structures on categories and functors of interest; for this we will need to understand how the process of taking Hochschild homology of categories interacts with the (co)sheaf condition.

Evaluating sheaves and cosheaves of categories

The (co)sheaf condition on pre(co)sheaves of categories implies that the Hom spaces between objects also follow gluing conditions. This is made precise by the following lemmata.

Lemma 15. *(The diagonal sheaf of bimodules $\mathcal{F}_\Delta$) Let $\mathcal{F}$ be a constructible sheaf of small categories on $\mathbb{X}$, with restriction functors denoted $\rho_U^V : \mathcal{F}(V) \to \mathcal{F}(U)$. Let us denote $R_U^V = \rho_U^V \otimes (\rho_U^V)^{op}$. Then the $\mathcal{F}(\mathbb{X})\text{-}\operatorname{Mod}\text{-}\mathcal{F}(\mathbb{X})$-valued presheaf that assigns*

$$U \mapsto (R_U^{\mathbb{X}})^* \mathcal{F}(U)_\Delta$$

with restriction maps given by applying $(R_V^{\mathbb{X}})^*$ to the canonical 'unit' maps

$$\mathcal{F}(V)_\Delta \to (R_U^V)^* \mathcal{F}(U)_\Delta$$

is a sheaf.

Proof. For any category $\mathcal{A}$, its category of bimodules $\mathcal{A}\text{-}\operatorname{Mod}\text{-}\mathcal{A} = \underline{\operatorname{Hom}}(\mathcal{A}, \operatorname{Vect})$ is by definition the category of Vect- valued presheaves on $\mathcal{A}^e$; by general facts of category theory, limits and colimits in such categories can be computed pointwise; ie. for any $(x, y) \in \mathcal{A}^e$ taking limits or colimits in $\mathcal{A}\text{-}\operatorname{Mod}\text{-}\mathcal{A}$ commutes with evaluation at (x, y). So for the lemma above it is enough to check that for any two objects $x, y \in \mathcal{F}(\mathbb{X})$, the presheaf

$$U \mapsto (R_U^{\mathbb{X}})^* \mathcal{F}(U)_\Delta(x, y) = \mathcal{F}(U)(\rho_U^V x, \rho_U^V y)$$

is a sheaf. This follows immediately from the fact that $\mathcal{F}$ is a sheaf of categories. □

There is an analogous statement for cosheaves:

Lemma 16. *(The diagonal cosheaf of bimodules $\mathcal{W}_\Delta$) Let $\mathcal{W}$ be a constructible cosheaf of dg categories over $\mathbb{X}$, with corestriction functors denoted $\iota_U^V : \mathcal{W}(U) \to \mathcal{W}(V)$. Let us denote $I_U^V = \iota_U^V \otimes (\iota_U^V)^{op}$. Then the $\mathcal{W}(\mathbb{X})$-Mod-$\mathcal{W}(\mathbb{X})$-valued precosheaf that assigns*

$$U \mapsto (I_U^{\mathbb{X}})_! \mathcal{W}(U)_\Delta$$

with corestriction maps given by applying $(I_V^{\mathbb{X}})_!$ to the canonical 'counit' maps

$$(I_U^V)_! \mathcal{W}(U)_\Delta \to \mathcal{W}(V)_\Delta$$

is a cosheaf.

For the proof, see [100, Sec.2.3]

Sheafified and cosheafified Hochschild homology

Let $\mathbb{X}$ be a topological space and $\mathcal{F}$ a sheaf of dg categories over $\mathbb{X}$. Taking Hochschild complexes gives a covariant functor, so there is a corresponding presheaf of complexes $\mathcal{HH}^{pre}(\mathcal{F})$ given by $\mathcal{HH}^{pre}(\mathcal{F})(U) = \mathrm{HH}(\mathcal{F}(U))$. This is *not* generally a sheaf; we write $\mathcal{HH}(\mathcal{F})$ for its sheafification. All the restriction functors are compatible with the S^1 action so we can also define the negative cyclic complex sheaf $\mathcal{HH}(\mathcal{F})^{S^1}$ and the cyclic complex sheaf $\mathcal{HH}(\mathcal{F})_{S^1}$, with maps of sheaves

$$\mathcal{HH}(\mathcal{F})^{S^1} \to \mathcal{HH}(\mathcal{F}) \to \mathcal{HH}(\mathcal{F})_{S^1}$$

There is naturally a morphism of presheaves $\mathcal{HH}^{pre}(\mathcal{F}) \to \mathcal{HH}(\mathcal{F})$, and in particular a morphism

$$\mathrm{HH}(\mathcal{F}(U)) \to \mathcal{HH}(\mathcal{F})(U)$$

on any open set U, compatible with the S^1 actions.

Example. Let $\mathcal{L}oc$ be the constant sheaf of categories over $\mathbb{X}$ with stalk Perf_k. Then the Hochschild complex presheaf is given by cochains on the loop space; this is obtained by dualizing the statement of [77, Thm 7.3.14]. The corresponding cyclic complex presheaf is

$$\mathcal{HH}^{pre}(\mathcal{L}oc)_{S^1}(\mathbb{X}) = \mathrm{HH}(\mathcal{L}oc(\mathbb{X}))_{S^1} \cong C^\bullet(L\mathbb{X})_{S^1}$$

where we take the homotopy orbits of the S^1 action that rotates the loop. On the other hand the local sections over contractible open sets is $\mathrm{HH}(\mathrm{Loc}(U)) \cong k$, with trivial S^1 action and its sheafification is the constant sheaf $k_{\mathbb{X}}$ and hence its derived global sections is cochains on $\mathbb{X}$ itself.

$$\mathcal{HH}(\mathcal{L}oc)_{S^1}(\mathbb{X}) = C^\bullet(\mathbb{X})$$

The localization morphism $\mathcal{HH}^{pre}(\mathcal{L}oc)_{S^1}(\mathbb{X}) \to \mathcal{HH}(\mathcal{L}oc)_{S^1}(\mathbb{X})$ is the pullback on cochains $C^\bullet(L\mathbb{X})_{S^1} \to C^\bullet(\mathbb{X})$ corresponding to the inclusion of constant loops.

We can do the same for a cosheaf of categories $\mathcal{W}$ over $\mathbb{X}$: we get a precosheaf $\mathcal{HH}^{pre}(\mathcal{W})$, with a cosheafification $\mathcal{HH}(\mathcal{W})$ and maps of cosheaves

$$\mathcal{HH}(\mathcal{W})^{S^1} \to \mathcal{HH}(\mathcal{W}) \to \mathcal{HH}(\mathcal{W})_{S^1}$$

and a natural morphism of precosheaves $\mathcal{HH}(\mathcal{W}) \to \mathcal{HH}^{pre}(\mathcal{W})$ (*from* the cosheafification), which over any open set U gives a map compatible with the S^1 actions

$$\mathcal{HH}(\mathcal{W})(U) \to \mathrm{HH}(\mathcal{W}(U))$$

Example. Let $\mathcal{L}oc^w$ be the constant cosheaf over $\mathbb{X}$ with stalk Perf_k. Over any connected open set U, the sections are perfect modules over the chains on the based loop space of U: $\mathcal{L}oc^w(U) = \mathrm{Perf}\text{-}\Omega U$. The Hochschild homology precosheaf is given by chains on the free loop space $\mathcal{HH}^{pre}(\mathcal{L}oc^w)(U) = \mathrm{HH}(\mathcal{L}oc^w(U)) \cong C(LU)$.

On the other hand the Hochschild homology cosheaf is constant with stalk $\mathrm{HH}(\mathrm{Perf}_k) = k$ and trivial circle action, so the global sections of the cyclic cosheaf are are $\mathcal{HH}(\mathcal{L}oc^w)^{S^1}(\mathbb{X}) = C_*(\mathbb{X})$. The colocalization morphism $\mathcal{HH}(\mathcal{L}oc^w)(\mathbb{X}) \to \mathcal{HH}^{pre}(\mathcal{L}oc^w)(\mathbb{X})$ is given by the pushforward of chains under inclusion of constant loops $C(\mathbb{X}) \to C(L\mathbb{X})$, which naturally factors through the (homotopy) fixed points of the circle action.

Proposition 17. *Let $\mathcal{W}$ be a locally saturated constructible cosheaf of triangulated dg categories on $\mathbb{X}$, and $\mathcal{W}^{pp}$ its sheaf of pseudo-perfect modules. Then the Hochschild homology cosheaf $\mathcal{HH}(\mathcal{W})$ and the Hochschild homology sheaf $\mathcal{HH}(\mathcal{W}^{pp})$ are linear duals. Moreover the S^1 actions are compatible.*

Proof. On sufficiently small open sets U_ϵ there is an isomorphism $\mathcal{W}^{pp}(U_\epsilon) \cong \mathcal{W}(U_\epsilon)$ and the functor $\mathcal{D}$ of Proposition 14 gives an anti-involution of $\mathcal{W}(U_\epsilon)\text{-}\mathrm{Mod}\text{-}\mathcal{W}(U_\epsilon)$ which maps $\mathcal{W}(U_\epsilon)^!$ to $\mathcal{W}(U_\epsilon)^*$ and gives isomorphisms

$$\mathrm{HH}(\mathcal{W}^{pp}(U_\epsilon)) \cong \mathrm{HH}(\mathcal{W}(U_\epsilon)) \xrightarrow{\sim} \mathrm{Hom}_k(\mathrm{HH}(\mathcal{W}(U_\epsilon)), k)$$

Taking a cover of any open set U by such small open sets gives an isomorphism

$$\mathcal{HH}(\mathcal{W}^{pp})(U) \cong \lim_{U_\epsilon} \mathrm{Hom}_k(\mathrm{HH}(\mathcal{W}(U_\epsilon)), k) = \mathrm{Hom}_k(\operatorname{colim}_{U_\epsilon} \mathrm{HH}(\mathcal{W}(U_\epsilon)), k) = \mathrm{Hom}_k(\mathcal{HH}(\mathcal{W})(U), k)$$

However since the local categories are smooth and proper the local Hochschild homologies $\mathrm{HH}(\mathcal{W}(U_\epsilon))$ are perfect complexes and therefore $\mathcal{HH}(\mathcal{W}^{pp})(U)$ is perfect, so we also have an isomorphism

$$\mathcal{HH}(\mathcal{W})(U) \cong \mathrm{Hom}_k(\mathcal{HH}(\mathcal{W}^{pp})(U), k)$$

□

Local orientations on sheaves of categories

Definition 15. Let $\mathbb{X}$ be a stratifiable space of pure dimension d, and let $\mathcal{F}$ be a constructible sheaf of proper dg categories on $\mathbb{X}$. A local orientation on $\mathcal{F}$ is a morphism of sheaves

$$\Theta : \mathcal{HH}(\mathcal{F})_{S^1} \to \omega_{\mathbb{X}}[-d]$$

where $\omega_{\mathbb{X}}$ is the Verdier dualizing complex of $\mathbb{X}$.

Remark. Recall that on an open set U, we have $H^*(U, \omega_{\mathbb{X}}[-d]) = H_*(\overline{U}, \partial\overline{U}; k)$. If $x \in \mathbb{X}$ is a point with a compact conical neighborhood $\overline{U}$, we have

$$\omega_{\mathbb{X}}[-d]|_x = H_*(\overline{U}; \partial\overline{U}) \cong H_*(\overline{U}; \overline{U} \setminus x).$$

Thus a morphism $k_{\mathbb{X}} \to \omega_{\mathbb{X}}[-d]$ at stalks is an element of $H_d(\overline{U}; \overline{U} \setminus x)$, i.e. an orientation on $\mathbb{X}$ in the classical sense. So when $\mathcal{F} = \mathcal{L}oc$ a (non-zero) local orientation on $\mathcal{F}$ is an orientation on $\mathbb{X}$.

Let $pt : \mathbb{X} \to *$. Then there is a canonical "integration" map $\Gamma_c(\mathbb{X}, \omega_{\mathbb{X}}) = pt_! pt^! k \to k$. We shift this map by the dimension d of $\mathbb{X}$ to get it as $\Gamma_c(\mathbb{X}, \omega_{\mathbb{X}}[-d]) \to k[-d]$.

Proposition 18. *Let $(\mathbb{X}, \mathcal{F})$ as above, with $\mathbb{X}$ compact. Then composing the following morphisms gives a map from the cyclic complex*

$$\mathrm{HH}(\mathcal{F}(\mathbb{X}))_{S^1} \to \Gamma(\mathbb{X}, \mathcal{HH}(\mathcal{F})_{S^1}) \cong \Gamma_c(\mathbb{X}, \mathcal{HH}(\mathcal{F})_{S^1}) \to \Gamma_c(\mathbb{X}, \omega_{\mathbb{X}}[-d]) \to k[-d]$$

For any open $V \subseteq \mathbb{X}$ and pair of objects $x, y \in \mathcal{F}(V)$, there is a Vect-valued sheaf $\mathcal{F}_\Delta(x, y)$ on V given by evaluating the diagonal bimodule sheaf; explicitly it assigns

$$U \mapsto \mathcal{F}(U)(\rho_U^V y, \rho_U^V x)$$

for $U \subseteq V$. The trace pairing to Hochschild homology gives us a map of sheaves

$$\mathcal{F}_\Delta(x, y) \otimes_V \mathcal{F}_\Delta(y, x) \to \mathcal{HH}(\mathcal{F})|_V \xrightarrow{\Theta} \omega_V[-d]$$

which by adjunction gives us a map of sheaves

$$\tilde{\Theta} : \mathcal{F}_\Delta(x, y) \to \mathcal{H}om(\mathcal{F}_\Delta(y, x), \omega_V)[-d]$$

Note that the rhs is a shift of the Verdier dual, which we can denote by $\mathbb{D}_V \mathcal{F}_\Delta(y, x)[-d]$

Definition 16. (Local nondegeneracy) A local orientation Θ is nondegenerate if, on any open $V \subset \mathbb{X}$ and two objects x, y of $\mathcal{F}(V)$, the morphism of sheaves on V

$$\tilde{\Theta} : \mathcal{F}_\Delta(x, y) \to \mathbb{D}_V \mathcal{F}_\Delta(y, x)[-d]$$

is an isomorphism.

Proper Calabi-Yau structures and sheaves of categories

Proposition 19. *Assume $\mathbb{X}$ is compact, and let $\Theta : \mathcal{HH}(\mathcal{F})_{S^1} \to \omega_{\mathbb{X}}$ be a nondegenerate local orientation on a sheaf $\mathcal{F}$ of proper categories. Then the induced map from the cyclic complex $\mathrm{HH}(\mathcal{F}(\mathbb{X}))_{S^1} \to k[-d]$ defines a d-dimensional proper Calabi-Yau structure on the global sections $\mathcal{F}(\mathbb{X})$.*

Proof. Note that for any pair of objects x, y in $\mathcal{F}(\mathbb{X})$, we can take global sections of the morphism $\tilde{\Theta}$ and get an isomorphism

$$\mathcal{F}(\mathbb{X})_\Delta(x,y) \xrightarrow{\sim} \Gamma(\mathbb{X}, \mathbb{D}_{\mathbb{X}}\mathcal{F}(\mathbb{X})_\Delta(y,x))[-d] \cong \Gamma_c(\mathbb{X}, \mathbb{D}_{\mathbb{X}}\mathcal{F}(\mathbb{X})_\Delta(y,x))[-d] \cong \mathrm{Hom}_k(\mathcal{F}(\mathbb{X})_\Delta(y,x),k)[-d] = \mathcal{F}(\mathbb{X})^*_\Delta(x,y)$$

where we used the fact that $\mathbb{X}$ is compact in going from sections to compactly supported sections and the definition of the linear dual $\mathcal{F}(\mathbb{X})^*_\Delta$ of the diagonal bimodule. By functoriality this isomorphism between is exactly the one coming from the identification $\mathrm{Hom}(\mathrm{HH}(\mathcal{F}(\mathbb{X})),k) \cong \mathrm{Hom}_{\mathcal{F}(\mathbb{X})\text{-}\mathrm{Mod}\text{-}\mathcal{F}(\mathbb{X})}(\mathcal{F}(\mathbb{X})_\Delta, \mathcal{F}(\mathbb{X})^*_\Delta)$ coming from the same local orientation. □

In the relative case:

Proposition 20. *Assume $(\mathbb{X}, \partial\mathbb{X})$ is a stratified space with compact boundary, such that the boundary is transverse to the stratification, and let $\Theta : \mathcal{HH}(\mathcal{F})_{S^1} \to \omega_X X[-d]$ be a nondegenerate local orientation on a sheaf $\mathcal{F}$ of proper categories. Then the local orientation Θ induces a d-dimensional relative proper Calabi-Yau structure on the restriction functor $\partial : \mathcal{F}(\mathbb{X}) \to \mathcal{F}(\partial\mathbb{X})$ and a $(d-1)$-dimensional (absolute) proper Calabi-Yau structure on the boundary sections $\mathcal{F}(\partial\mathbb{X})$.*

The relative case also follows from a Verdier duality argument, but more complicated; see [100, Sec.2.4] for the proof.

Smooth Calabi-Yau structures and cosheaves of categories

Proposition 21. *As above, let $\mathcal{W}$ be a locally saturated cosheaf of smooth dg categories on a stratified space $\mathbb{X}$ with compact boundary ∂XX, and $\mathcal{W}^{pp}$ its sheaf of pseudo-perfect modules; this is a sheaf of proper dg categories. Then a non-degenerate local orientation $\Theta : \mathcal{HH}(\mathcal{W}^{pp}) \to \omega_{\mathbb{X}}[-d]$ on the sheaf $\mathcal{W}^{pp}$ gives a d-dimensional relative smooth Calabi-Yau structure on the corestriction functor $\mathcal{W}(\partial\mathbb{X}) \to \mathcal{W}(\mathbb{X})$, and a $(d-1)$-dimensional (absolute) smooth Calabi-Yau structure on the boundary cosections $\mathcal{W}(\partial\mathbb{X})$*

The proof of this proposition will proceed similarly to the construction of the cosheaf of compactly supported cochains on a topological space $\mathbb{X}$. Explicitly, consider a inclusion of open sets $U \subset V$ and consider the distinguished triangle coming from the long exact sequence of relative homology with coefficients in k

$$C_*(\partial U) \to C_*(U) \to C_*(U; \partial U)$$

By Poincaré duality $C_c^*(U) \cong \mathrm{Hom}_k(C_*(U;\partial U),k)$, yet there is no obvious restriction map $C_*(V;\partial V) \to C_*(U;\partial U)$ corresponding to the inclusion of compactly supported cochains $C_c^*(U) \to C_c^*(V)$. Indeed, $\partial U \not\subseteq \partial V$ so there is no natural map between $C_*(\partial U)$ and $C_*(\partial V)$. Moreover, if $\partial U \subset \partial V$, the map would go in the wrong direction.

The solution is to replace $C_*(U;\partial U)$ by $C_*(\mathbb{X}|U) := C_*(\mathbb{X}, \mathbb{X} \setminus U^\circ)$. Consider inclusion of pairs $(U;\partial U) \subseteq (\mathbb{X};\mathbb{X}\setminus U^\circ)$ giving a map of distinguished triangles

$$\begin{array}{ccccc} C_*(\partial U) & \longrightarrow & C_*(U) & \longrightarrow & C_*(U;\partial U) \\ \downarrow & & \downarrow & & \downarrow \sim \\ C_*(\mathbb{X}\setminus U^\circ) & \longrightarrow & C_*(\mathbb{X}) & \longrightarrow & C_*(\mathbb{X}|U) \end{array}$$

Because C_* is a cosheaf (i.e. excision holds), the left hand side square is a pushout square and the right vertical map is a quasi-isomorphism. There is also an inclusion of pairs $(\mathbb{X}, \mathbb{X}\setminus V^\circ) \subseteq (\mathbb{X}, \mathbb{X}\setminus U^\circ)$ and a corresponding map $C_*(\mathbb{X}|V) \to C_*(\mathbb{X}|U)$, and standard arguments imply that the assignment

$$U \mapsto C_*(\mathbb{X}|U)$$

is a sheaf, meaning that $U \mapsto \mathrm{Hom}_k(C_*(\mathbb{X}|U),k)$ is a cosheaf, which by the quasi-isomorphism above computes the compactly-supported cohomology. One can then prove Poincaré duality for a non-compact manifold M of dimension d by locally constructing a quasi-isomorphism $C_c^*(U) \cong C_{d-*}(U)$ compatible with the corestriction maps, and using the cosheaf property to globalize it to an isomorphism $C_c^*(M) \cong C_{d-*}(M)$

In our proof of Proposition 21 we use the same argument but replacing the cosheaf of cochains with cosheaves of bimodules; see [100, Sec.2.4].

3.4 A local orientation on an arboreal singularity

We will show that the arboreal singularity $\mathbb{T}$, equipped with an appropriate sheaf of categories $\mathcal{N}$, admit local orientations. More precisely, we will show that a rooting of T induces a canonical isomorphism $\mathcal{HH}(\mathcal{N}) \cong \omega_{\mathbb{T}}[1-|T|]$, which is moreover nondegenerate.

In order to define this sheaf of categories, we will need to first review some standard facts about categories of quiver representations.

Quiver representations

We recall some relevant facts about quiver representation theory, and set notation. Let $\vec{T}$ be a quiver, i.e., a directed graph. We write $k[\vec{T}]$ for the path algebra of the quiver, whose generators are the vertices and arrows, subject to the relations that the vertex generators are idempotent, and $ab = 0$ unless the head of a is the tail of b (and the head and tail of a vertex are itself). That is, we read paths *from left to right* and consequently quiver representations correspond to *right modules* over this algebra.

Example. For the quiver $1 \to 2 \to 3 \to \cdots \to n$, the path algebra can be identified with an algebra of triangular $n \times n$ matrices. The matrix $|i\rangle\langle j|$ corresponds to the unique path from the i'th vertex to the j'th vertex. The composition $|i\rangle\langle j||j\rangle\langle k| = |i\rangle\langle k|$ corresponds to the left-to-right composition rule $(i \to j)(j \to k) = (i \to k)$.

The path algebra $k[\vec{T}]$ can be seen as a dg category that has only one object, with endomorphisms given by the algebra $k[\vec{T}]$ concentrated in degree zero. Let us denote by

$$\mathrm{Mod}(\vec{T}) := \mathrm{Mod}\text{-}k[\vec{T}] = \underline{\mathrm{Hom}}(k[\vec{T}], \mathrm{Vect})$$

the (derived) dg category of right modules over $k[\vec{T}]$. We can also define the category of *ordinary* modules by

$$\mathrm{mod}(\vec{T}) := \underline{\mathrm{Hom}}(k[\vec{T}], \mathrm{mod}_k)$$

taking the underived internal hom in the category of k-linear categories from the path algebra to the category mod_k of k-modules. The dg category $\mathrm{Mod}(\vec{T})$ is a dg enhancement of the derived category of the k-linear category $\mathrm{mod}(\vec{T})$.

If the quiver Q is acyclic, the dg category $k[\vec{T}]$ is triangulated, smooth and proper [112] and thus, the notions of perfect and pseudo-perfect modules agree; a module $M \in \mathrm{Mod}(\vec{T})$ is a representation of the path algebra $k[\vec{T}]$ in complexes of k-modules, and M is (pseudo-)perfect if its underlying k-module is a perfect complex. We will denote by $\mathrm{Perf}(\vec{T})$ the category of such modules full dg subcategory of $\mathrm{Mod}(\vec{T})$ spanned by such objects. Note that this case (Q is acyclic) we also have $\mathrm{Perf}(\vec{T}) = \underline{\mathrm{Hom}}(k[\vec{T}], \mathrm{Perf}_k)$.

For vertices $\alpha, \beta \in \vec{T}$, we write $\alpha \geq \beta$ when there is a path from α to β, and we denote this unique path by $|\alpha\rangle\langle\beta|$. These compose in the usual way, $|\alpha\rangle\langle\beta||\beta\rangle\langle\gamma| = |\alpha\rangle\langle\gamma|$, and all other compositions vanish. We are particularly interested in the case when the edge directions arise from the choice of a fixed root vertex of T, by directing all the edges toward the root. We pronounce $\alpha \geq \beta$ as "alpha is above beta" or "beta is below alpha", so that everything is above the root.

We write $P_\alpha := |\alpha\rangle\langle\alpha|k[\vec{T}]$ for the right module of "paths from α"; since $\vec{T}$ is a tree, it is the representation which assigns k to each vertex admitting a path from α (i.e., each vertex below α in the notation above), and all morphisms isomorphisms. All paths must come from somewhere, so there is an internal direct sum splitting $k[\vec{T}] = \bigoplus_\alpha |\alpha\rangle\langle\alpha|k[\vec{T}]$ as right $k[\vec{T}]$-modules. The modules P_α are in fact the indecomposable projectives of the category $\mathrm{mod}(\vec{T})$, and the category $\mathrm{Perf}(\vec{T})$ is their triangulated hull inside of $\mathrm{Mod}(\vec{T})$

When $\alpha \geq \beta$, i.e., there is a path $|\alpha\rangle\langle\beta|$, then composition with this path gives a morphism

$$\begin{array}{rcl} |\beta\rangle\langle\beta|k[\vec{T}] & \to & |\alpha\rangle\langle\alpha|k[\vec{T}] \\ x & \mapsto & |\alpha\rangle\langle\beta|x \end{array}$$

In fact, up to scalars this is the only morphism in $\mathrm{mod}(\vec{T})$; since these modules are projective, this remains true in $\mathrm{Mod}(\vec{T})$ and $\mathrm{Perf}(\vec{T})$.

Quivers corresponding to the same underlying tree but different arrow orientations have representation categories related by reflection functors, defined in [14]. A source (sink) is a vertex that only has outgoing (ingoing) arrows. Given a source α, let $s_\alpha\vec{T}$ be the quiver obtained by reversing all the arrows at α. There is a reflection functor $R^+_\alpha : \mathrm{mod}(\vec{T}) \to \mathrm{mod}(s_\alpha\vec{T})$, which in fact preserves compact objects and so induces a dg derived equivalence $\mathrm{Perf}(\vec{T}) \to \mathrm{Perf}(s_\alpha\vec{T})$. Likewise, at sinks there are similar reflection functors R^-_α. The quiver structure for a rooted tree has all arrows pointing to the root. Because the underlying graph is acyclic, two such structures corresponding to different roots ρ_1, ρ_2 can be related by a sequence of moves s_α. Thus the derived categories $\mathrm{Mod}(\vec{T})$ and $\mathrm{Perf}(\vec{T})$ depends only on the underlying tree (up to non-canonical equivalence). Choosing a root determines a t-structure, and determines the distinguished set of projective generators $\{P_\alpha\}$.

Correspondence functors

Given a root of T, a correspondence $R \twoheadleftarrow S \hookrightarrow T$ induces root vertices, hence arrow orientations, of S and R – the root of S is the closest vertex in S to the root of T, and the root of R is the image of the root of S.

We can identify $k[\vec{S}]$ with the quotient of $k[\vec{T}]$ by the two-sided ideal generated by all paths that are not contained in $\vec{S}$. This gives a map $k[\vec{T}] \to k[\vec{S}]$ and by *extension of scalars* we get a functor $\mathrm{Mod}(\vec{T}) \to \mathrm{Mod}(\vec{S})$. Tensoring with a perfect module preserves compact objects so this restricts to a functor $\mathrm{Perf}(\vec{T}) \to \mathrm{Perf}(\vec{S})$.

On the other hand, given a quotient $S \overset{q}{\twoheadrightarrow} R$ we can construct the following morphism of k-algebras $k[\vec{R}] \to k[\vec{S}]$. For simplicity, assume the quotient corresponds to collapsing one connected subtree $Q \subset S$; the general case can be deduced by iterated quotients like these. Let ρ be the root in the induced quiver structure on Q, i.e. the lowest vertex in Q. The quotient identifies $q(\alpha) = q(\rho)$ for all $\alpha \in Q$. Consider the function $s : V(R) \to V(S)$ between the sets of vertices given by

$$s(\beta) = \begin{cases} q^{-1}(\beta) & \beta \notin Q \\ \rho & \beta \in Q \end{cases}$$

This is a one-sided inverse to q, since $q \circ s = id$ on $V(R)$. This determines a morphism $k[\vec{R}] \to k[\vec{S}]$, which acts on paths as

$$|\alpha\rangle\langle\beta| \mapsto |s(\alpha)\rangle\langle s(\beta)|$$

That is, this sends a path in $\vec{R}$ to the shortest path in $\vec{S}$ whose start and end-points map to the original start and end-points in $\vec{R}$. One can check that this commutes with compositions, defining a map of algebras $k[\vec{R}] \to k[\vec{S}]$, and moreover that this map presents $k[\vec{S}]$ as a perfect module over $k[\vec{R}]$. So *restriction of scalars* under this map gives a functor $\mathrm{Mod}(\vec{S}) \to \mathrm{Mod}(\vec{R})$ which preserves compact objects and restricts to a functor $\mathrm{Perf}(\vec{S}) \to \mathrm{Perf}(\vec{R})$

We get a 'big functor' $c_{\mathfrak{p}}^{\diamond} : \mathrm{Mod}(\vec{T}) \to \mathrm{Mod}(\vec{R})$ and a 'small functor' $c_{\mathfrak{p}} : \mathrm{Perf}(\vec{T}) \to \mathrm{Perf}(\vec{R})$ by composing the functors above.

Lemma 22. *The big functor $c_{\mathfrak{p}}^{\diamond}$ preserves products and coproducts.*

Proof. Restriction of scalars is a right adjoint and always preserves arbitrary products, moreover coproducts in $\mathrm{Mod}(\vec{R})$ are be calculated as coproducts of the underlying k-module so restriction also preserves coproducts. As for extension of scalars, it is a left adjoint and always preserves coproducts, and also preserves products when the bimodule is of finite presentation, which is the case for $k[\vec{S}]$ seen as a $(k[\vec{T}], k[\vec{S}])$-bimodule. □

Lemma 23. *Let $c_{\mathfrak{p}} : \mathrm{Perf}(\vec{T}) \to \mathrm{Perf}(\vec{R})$ and $c_{\mathfrak{p}}^{\diamond} : \mathrm{Mod}(\vec{T}) \to \mathrm{Mod}(\vec{R})$ be the functors induced by a correspondence $\mathfrak{p} : R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T$. Let $\alpha \in T$ be a vertex. Then*

$$c_{\mathfrak{p}}(P_\alpha) = \begin{cases} P_{q(i^{-1}(\alpha))} & \alpha \in i(S) \\ 0 & \textit{otherwise} \end{cases}$$

The morphism $\mathrm{Hom}_{\vec{T}}(P_\alpha, P_\beta) \to \mathrm{Hom}_{\vec{R}}(c_{\mathfrak{p}}(P_\alpha), c_{\mathfrak{p}}(P_\beta))$ sends $|\beta\rangle\langle\alpha| \to |q(i^{-1}(\beta))\rangle\langle q(i^{-1}(\alpha))|$ (and hence is an isomorphism) when these are defined; otherwise it is zero. The exact same description holds for the functor $c_{\mathfrak{p}}^{\diamond}$

Proof. In general, if we are given a quiver $\vec{T}$ and a right $k[\vec{T}]$-module M, in order to identify which module we have it is sufficient to look at the k-vector spaces $M^{(\alpha)} = M|\alpha\rangle\langle\alpha|$ for each vertex α, and whenever there is an arrow $\mu \to \alpha$, the map $M^{(\mu)} \to M^{(\alpha)}$ given by right multiplication by $|\mu\rangle\langle\alpha|$, since this data determines the module M.

Consider first the functor $\mathrm{Perf}(\vec{T}) \to \mathrm{Perf}(\vec{S})$. The image of $P(\alpha)$ is the $k[\vec{S}]$-module $P_\alpha \otimes_{k[\vec{T}]} k[\vec{S}]$. If α lies outside S, every in path P_α can be expressed as $|\alpha\rangle\langle\alpha|x$, which gets sent to zero the quotient to $k[\vec{S}]$. If $\alpha \in i(S)$, exactly the paths in P_α exiting $\vec{S}$ are sent to zero, and $P_\alpha \otimes_{k[\vec{T}]} k[\vec{S}]$ is spanned by all paths in $\vec{S}$ starting at $i^{-1}(\alpha)$ so this module is $P_{i^{-1}(\alpha)}$.

Now for the functor $\mathrm{Perf}(\vec{S}) \to \mathrm{Perf}(\vec{R})$, for some vertex β of S, let M be the image of P_β under this functor. Remembering that this functor is induced by a map of k-algebras $f : k[\vec{R}] \to k[\vec{S}]$, for any vertex λ of R we have an isomorphism of k-vector spaces

$$M^{(\lambda)} = M|\lambda\rangle\langle\lambda| \cong P_\beta f(|\lambda\rangle\langle\lambda|) = P_\beta|s(\lambda)\rangle\langle s(\lambda)|$$

which is k exactly when $s(\lambda) \leq \beta$ or equivalently $\lambda \leq q(\beta)$. The morphisms between the $M^{(\lambda)}$ are given by multiplication by $|\mu\rangle\langle\lambda|$; we need to check that these are isomorphisms whenever there is an arrow $\mu \to \lambda$ and $\mu \leq q(\beta)$. As maps of vector spaces,

$$M^{(\mu)} \xrightarrow{|\mu\rangle\langle\lambda|} M^{(\lambda)}$$

is the same as the map

$$P_\beta|s(\mu)\rangle\langle s(\mu)| \xrightarrow{|s(\mu)\rangle\langle s(\lambda)|} P_\beta|s(\lambda)\rangle\langle s(\lambda)|$$

which is an isomorphism since s respects the partial ordering $\leq$. This identifies the module M with the indecomposable projective $P_{q(\beta)}$ in $\mathrm{Perf}(\vec{R})$.

Putting the two functors together we get the first half of the result. For the morphisms, the nontrivial case to check is when $\alpha \leq \beta$ and $\alpha, \beta \in i(S)$. At the level of paths, the map $|\beta\rangle\langle\alpha| \in \mathrm{Hom}_{\vec{T}}(P_\alpha, P_\beta)$ is given by pre-concatenation with the path $|\beta\rangle\langle\alpha|$. After applying the functor $c_\mathfrak{p}$ this becomes a map in $\mathrm{Hom}_{\vec{R}}(c_\mathfrak{p}(P_\alpha), c_\mathfrak{p}(P_\beta))$ given by concatenation with the path $|q(i^{-1}(\beta))\rangle\langle q(i^{-1}(\alpha))|$, which is nonzero since S is connected. By definition the functor $c_\mathfrak{p}$ is the restriction of $c_\mathfrak{p}^\diamond$ to compact objects so the exact same calculation holds for the big categories. □

Definition 17. Fix a rooted tree $\vec{T}$. We define a functor into small categories $N : \mathrm{Arb}_T \to \mathrm{dgst}_k$ at the level of objects by

$$(R \twoheadleftarrow S \hookrightarrow T) \mapsto \mathrm{Perf}(\vec{R})$$

and a big functor $N^\diamond : \mathrm{Arb}_T \to \mathrm{dgSt}_k^R$ at the level of objects by A morphism $(R \twoheadleftarrow S \hookrightarrow T) \to (R' \twoheadleftarrow S' \hookrightarrow T)$ is by definition a correspondence $\mathfrak{p} = (R' \twoheadleftarrow S'' \hookrightarrow R)$; we send this respectively to the functors $c_\mathfrak{p} : \mathrm{Perf}(\vec{R}) \to \mathrm{Perf}(\vec{R}')$ and $c_\mathfrak{p}^\diamond : \mathrm{Mod}(\vec{R}) \to \mathrm{Mod}(\vec{R}')$ determined by this correspondence.

Note that this prescription does give a functor to dgSt_k^R because the functors $c_\mathfrak{p}^\diamond$ preserve limits. Therefore, by taking left adjoints we also get a functor $(N^\diamond)^L : (\mathrm{Arb}_T)^{op} \to \mathrm{dgSt}_k^L$ which sends the correspondence $\mathfrak{p}$ to the left adjoint $(c_\mathfrak{p}^\diamond)^L : \mathrm{Mod}(\vec{R}') \to \mathrm{Mod}(\vec{R})$. Moreover since the functor $N^\diamond$ also preserves coproducts its left adjoint preserves compact objects, so we can restrict it and get a functor $N^w : \mathrm{Perf}(\vec{R}') \to \mathrm{Perf}(\vec{R})$.

The sheaf and cosheaf associated to an arboreal singularity

Constructible sheaves on simplicial complexes

We briefly recall how to describe constructible sheaves on a simplicial complex. For a simplex σ in a simplicial complex X, we write $\mathrm{Star}(\sigma)$ for the union of open simplices whose closure contains σ. To give a sheaf $\mathcal{F}$ on X, constructible with respect to the stratification by simplices, it suffices to give the values of $\mathcal{F}$ on the open sets $\mathrm{Star}(\sigma)$, and the corresponding restriction maps $\mathcal{F}(\mathrm{Star}(\sigma)) \to \mathcal{F}(\mathrm{Star}(\tau))$ when $\mathrm{Star}(\tau) \subset \mathrm{Star}(\sigma)$, i.e., when σ lies in the closure of τ. The appropriate diagrams should commute. Our definition of simplicial complex demands that the closure of an open simplex is a closed simplex, so there are no non-trivial overlaps, hence no descent conditions.

The restriction $\mathcal{F}(\mathrm{Star}(\sigma)) \to \mathcal{F}_\sigma$ is then necessarily an isomorphism, so one could instead discuss "generization maps" $\mathcal{F}_\sigma \to \mathcal{F}_\tau$ when σ lies in the closure of τ; this is the so-called "exit path" description of a constructible sheaf. A similar description works for any sufficiently fine stratification.

Functors give sheaves and cosheaves on the nerve

Recall that a functor $F : \mathcal{X} \to \mathcal{Y}$ determines a $\mathcal{Y}$-valued constructible sheaf Nerve(F) on Nerve($\mathcal{X}$) as follows. On objects, we set

$$Nerve(F)([x_1 \to x_2 \to \cdots \to x_n]) = F(x_n)$$

and the generization maps are given by

$$Nerve(F)([x_m \to \cdots \to x_{m'}] \to [x_1 \to x_2 \to \ldots \to x_n]) = F(x'_m \to x_n)$$

where the map $x'_m \to x_n$ comes from the fact that $x_m \to \cdots \to x_{m'}$ was a subsequence of $x_1 \to x_1 \to \ldots \to x_n$. The fact that these restriction maps satisfy the appropriate conditions to determine a sheaf is immediate from the fact that F is a functor. Note this sheaf is constructible on a much coarser stratification than the stratification by all simplices.

Nadler's sheaf and cosheaf

Let $\vec{T}$ be a rooted tree. Above we constructed functors

$$N : \mathrm{Arb}_T \to \mathrm{dgst}_k \text{ and } N^\diamond : \mathrm{Arb}_T \to \mathrm{dgSt}_k^R$$

Definition 18. The 'small' Nadler sheaf $\mathcal{N}$ (or $\mathcal{N}_\mathbb{T}$ when we want to specify $\mathbb{T}$) is the sheaf of categories on $\mathbb{T} = \mathrm{Nerve}(\mathrm{Arb}_T)$ given by the nerve of the functor $N : \mathrm{Arb}_T \to \mathrm{dgst}_k$. The 'big' Nadler sheaf $\mathcal{N}^\diamond$ is the sheaf given by the nerve of the functor $N^\diamond$.

As in section 3.2, the sheaf $\mathcal{N}^\diamond$ valued in dgSt_k^R gives a cosheaf valued in dgSt_k^L upon taking left adjoints (in fact the data of these two objects is the same). Moreover, the restriction maps of the sheaf $\mathcal{N}^\diamond$ preserve coproducts [85, Prop. 3.16], so their left adjoints, the corestriction maps of the cosheaf $\mathcal{N}^\diamond$, preserve compact objects.

Definition 19. The *wrapped cosheaf* $\mathcal{W}$ is the dgst_k-valued cosheaf obtained from the cosheaf $\mathcal{N}^\diamond$ by restriction to the full subcategories of compact objects.

By definition, the stalks of $\mathcal{N}$ or $\mathcal{W}$ over any stratum is the saturated (smooth and proper) Perf($\vec{R}$) for some finite acyclic quiver $\vec{R}$. Thus, by finiteness of the stratification, $\mathcal{N}(U)$ is proper for any open set $U \subseteq \mathbb{T}$ since a finite limit of proper categories is proper. Analogously, $\mathcal{W}(U)$ is smooth and finite type for any open U since a finite colimit of smooth and finite type categories is smooth and finite type. From the fact that the cosheaf $\mathcal{W}$ is locally saturated in the sense of Section 3.3 it follows that:

Lemma 24. *[85, Prop. 3.16] The Nadler sheaf $\mathcal{N}$ is equivalent to the sheaf $\mathcal{W}^{pp}$ of pseudo-perfect modules over the wrapped cosheaf $\mathcal{W}$.*

By construction the generization maps between stalks of the cosheaf $\mathcal{W}$ are left adjoint to the generization maps of the sheaf $\mathcal{N}$. Note that over bigger open sets U, there is still a map $\mathcal{N}(U) \hookrightarrow \mathcal{W}(U)$ equivalent to the inclusion of a full dg subcategory, but in general this will not be essentially surjective; $\mathcal{W}(U)$ can be a bigger category.

Remark. For the rest of this section, we will keep writing the sheaf as $\mathcal{N}$ for conciseness, but in fact the cosheaf $\mathcal{W}$ is the more fundamental object; one can obtain $\mathcal{N} \cong \mathcal{W}^{pp}$ by taking pseudo-perfect modules, and the 'big' co/sheaf $\mathcal{N}^{\diamond}$ by cocompletion, but $\mathcal{W}$ cannot be obtained from $\mathcal{N}$ if one doesn't have knowledge of $\mathcal{N}^{\diamond}$. Thus in our definition of locally arboreal space 8 we take $\mathcal{W}$ as part of the data.

Let $\mathfrak{p}$ be a correspondence $R \twoheadleftarrow S \hookrightarrow T$. We write $\mathbb{T}(\mathfrak{p})$ for the union of all simplices $[\mathfrak{p}_1 \to \mathfrak{p}_2 \cdots \to \mathfrak{p}]$. Then $\mathbb{T} = \coprod \mathbb{T}(\mathfrak{p})$, and, by definition, any sheaf associated to a functor from Arb_T is constant on the $\mathbb{T}(\mathfrak{p})$. In fact, each $\mathbb{T}(\mathfrak{p})$ is topologically an open cell of dimension $|T| - |R|$ [81, Prop 2.14].[1] In particular, $\mathbb{T}(\mathfrak{p}_T)$ is the unique zero dimensional cell. Moreover, $\overline{\mathbb{T}(\mathfrak{p})} = \coprod_{\mathfrak{p}' \leq \mathfrak{p}} \mathbb{T}(\mathfrak{p}')$.

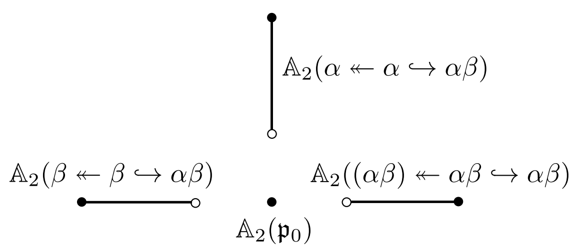

Figure 3.1: Coarser stratification of the arboreal singularity $\mathbb{A}_2$, by the strata $\mathbb{A}_2(\mathfrak{p})$.

Fix $\mathbb{T}$ and denote $\mathcal{N} = \mathcal{N}_{\mathbb{T}}$. Because $\mathcal{N}$ is constructible with respect to a stratification by a union of cells which all adjoin $\mathfrak{p}_T$, the restriction map $\Gamma(\mathbb{T}, \mathcal{N}) \to \mathcal{N}_{\mathfrak{p}_T}$ is an isomorphism. In particular, $\Gamma(\mathbb{T}, \mathcal{N}) \cong \mathrm{Perf}(\vec{T})$. Given an object $X \in \Gamma(\mathbb{T}, \mathcal{N}) = \mathrm{Perf}(\vec{T})$, the germ at a point in $\mathbb{T}(R \twoheadleftarrow S \hookrightarrow T)$ is an element of the category $\mathrm{Perf}(\vec{R})$. The desired object is produced by applying the correspondence functor $c_{\mathfrak{p}}$ obtained from $(R \twoheadleftarrow S \hookrightarrow T)$ to X.

Hom sheaves

The correspondence functors also give a coherent choice of maps between Hom spaces. Suppose we are given elements $x, y \in \Gamma(\mathbb{T}, \mathcal{N}) = \mathrm{Perf}(\vec{T})$. Because $\mathcal{N}$ is a sheaf of dg categories on $\mathbb{T}$, we can evaluate the corresponding diagonal sheaf of bimodules and get the Vect-valued sheaf $\mathcal{N}_{\Delta}(x, y)$, on $\mathbb{T}$. This is just the hom sheaf; it is the nerve of the functor

$$(R \twoheadleftarrow S \hookrightarrow T) \mapsto Hom_{\vec{R}}(c_{\mathfrak{p}}(n), c_{\mathfrak{p}}(n'))$$

[1] Nadler writes L_T for our $\mathbb{T}$ and $\mathrm{L}_T(\mathfrak{p})$ for our $\mathbb{T}(\mathfrak{p})$ in [81]. Our notation is chosen to emphasize that no symplectic geometry or microlocal sheaf theory is directly needed to understand the essentially combinatorial definitions and proofs.

By definition, the generization maps between the stalks of the Hom sheaf are induced by the correspondence functors $c_\mathfrak{p}$ $c_\mathfrak{p} : \mathrm{Hom}_{\vec{T}}(x, y) \to \mathrm{Hom}_{\vec{R}}(c_\mathfrak{p}x, c_\mathfrak{p}y)$

We will need explicit descriptions of the Hom sheaves between the generating projectives P_α. Since we know what the functors $c_\mathfrak{p}$ do to the projective objects from Lemma 23, it is just a matter of assembling the sheaf of the morphisms between Hom spaces.

Definition 20. For α a vertex of T, we write

$$\mathbb{T}(\alpha) := \coprod_{\alpha \in S} \mathbb{T}(R \twoheadleftarrow S \hookrightarrow T)$$

Remark. Nadler gives an explicit construction of the arboreal singularities: for each vertex α of T, take a copy of $\mathbb{R}^{|T|-1}$ with coordinates $x^\gamma(\alpha), \gamma \neq \alpha$. The interior of $\mathbb{T}$ is recovered by gluing these spaces: for each edge $\{\alpha, \beta\} \in E(T)$, identify points with coordinates $x_\gamma(\alpha)$ and $x_\gamma(\beta)$ whenever $x_\beta(\alpha) = x_\alpha(\beta) \geq 0$ and $x_\gamma(\alpha) = x_\gamma(\beta)$ for $\gamma \neq \alpha, \beta$. Comparing this construction with the combinatorial definition [81, Sec. 2] it is proven that the strata $\mathbb{T}(R \twoheadleftarrow S \hookrightarrow T)$ sit in the closure of the Euclidean space corresponding to α exactly when $\alpha \in S$. Thus $\mathbb{T}(\alpha)$ is homeomorphic to a closed ball of dimension $|T| - 1$.

The following calculations are new.

Proposition 25. *The sheaf $\mathcal{N}_\Delta(P_\alpha, P_\alpha)$ is the constant rank one sheaf on $\mathbb{T}(\alpha)$.*

Proof. Let us describe the functor on Arb_T whose nerve is the sheaf $\mathcal{N}_\Delta(P_\alpha, P_\alpha)$. By Lemma 23, on objects this functor is:

$$(R \twoheadleftarrow S \hookrightarrow T) \mapsto \mathrm{Hom}_{\vec{R}}(P_{q(i^{-1}(\alpha))}, P_{q(i^{-1}(\alpha))}) = \begin{cases} k, & \alpha \in i(S) \\ 0, & \text{otherwise} \end{cases}$$

These Hom spaces have the identity as a basis element, which must be preserved by the functorial structure, hence gives a global section trivializing the sheaf hom. □

To describe other Hom sheaves we have to worry about the orientation of the arrows in the quiver. More generally let

$$\mathbb{T}(\alpha, \beta) := \coprod_{\substack{\alpha, \beta \in S \\ q(i^{-1}(\alpha)) \leq q(i^{-1}(\beta))}} \mathbb{T}(R \twoheadleftarrow S \hookrightarrow T)$$

Proposition 26. *The sheaf $\mathcal{N}_\Delta(P_\alpha, P_\beta)$ is the constant rank one sheaf on $\mathbb{T}(\alpha, \beta)$.*

This can be deduced by a similar proof, see [100, Sec.3].

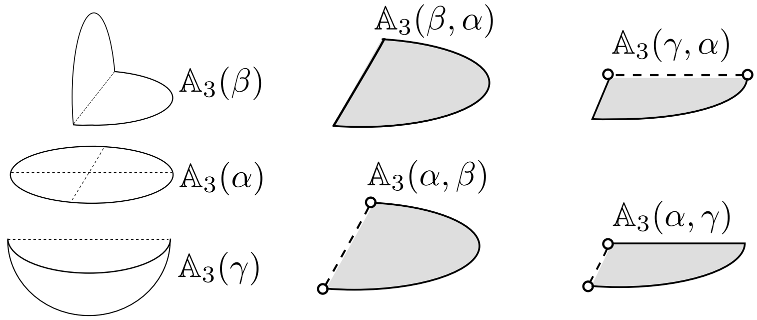

Figure 3.2: The subsets $\mathbb{A}_3(\bullet,\bullet)$ where $\mathbb{A}_3$ has the quiver structure $\alpha \to \beta \leftarrow \gamma$. Note that the subsets $\mathbb{T}(\lambda_1,\lambda_2)$ depend on the directions of the arrows in T, and moreover as in the proof above for any vertices λ_1,λ_2, the difference between $\mathbb{T}(\lambda_1,\lambda_2)$ and $\mathbb{T}(\lambda_1)\cap\mathbb{T}(\lambda_2)$ is at most deletion of some boundary strata.

The dualizing complex of an arboreal singularity

Verdier's dualizing complex on a space X is usually defined as $\omega_X = pt^! k$, where pt is the map to a point. An explicit representative is given by the "sheaf of local singular chains". That is, let $\mathbf{C}^{-d}$ be the sheaf which on sufficiently small open sets is given by $\mathbf{C}^{-d}(U) = C_d(X, X\setminus U; k)$, where C_d is the singular d-chains, and the sheaf structure is defined by the evident restriction maps. The singular chain differential collects these into a complex of sheaves, which is quasi-isomorphic to the dualizing complex.

We can restrict the stratification of $\mathbb{T}$ given by the $\mathbb{T}(\mathfrak{p})$ to $\mathbb{T}_\circ$; this stratification of $\mathbb{T}_\circ$ by open simplices $\mathbb{T}_\circ(\mathfrak{p})$ agrees exactly with the stratification originally presented in [81]. Since each stratum is an open simplex, the neighborhoods of every point along each strata are all homeomorphic, so the dualizing complex of $\mathbb{T}_\circ$ is constructible with respect to this stratification We will only need to calculate the dualizing complex on the open arboreal singularity $\mathbb{T}_\circ$, and moreover $\mathbb{T}_\circ \hookrightarrow \mathbb{T}$ is an open inclusion, so we can identify $\omega_{\mathbb{T}_\circ}$ with the restriction of $\omega_{\mathbb{T}}$. To give a complete description of the dualizing complex, it suffices to identify each stalk $\omega_{\mathbb{T}}^{-(n-1)}(\mathbb{T}_\circ(\mathfrak{p}))$ over each stratum, together with the necessary generization maps.

Proposition 27. *With notation as above, the stalk of $\omega_{\mathbb{T}}$ at a stratum labelled by a correspondence $\mathfrak{p} = (R \twoheadleftarrow S \hookrightarrow T)$ is concentrated in degree $-(n-1)$, where it is given by a direct sum decomposition*

$$\omega_{\mathbb{T}}^{-(n-1)}(\mathbb{T}_\circ(\mathfrak{p})) \cong \bigoplus_{\alpha\in R} k_\alpha \cong k^{|R|}$$

where each $k_\alpha \cong k$. Now suppose we have correspondences $\mathfrak{p}' = \mathfrak{q}\circ\mathfrak{p}$, where $\mathfrak{p} = (R' \twoheadleftarrow S' \hookrightarrow T)$ and $\mathfrak{q} = (R' \overset{q}{\twoheadleftarrow} Q \overset{i}{\hookrightarrow} R)$. Then the simplex $[\mathfrak{p}]$ is in the closure of $[\mathfrak{p}\to\mathfrak{p}']$ and the *generization map* $\omega_{\mathbb{T}}(\mathbb{T}(\mathfrak{p})) \to \omega_{\mathbb{T}}(\mathfrak{p}')$ *is given, in the decomposition above, by*

$$\bigoplus_{\alpha \in R} k_\alpha \to \bigoplus_{\beta \in R'} k_\beta$$

where $1 \in k_\alpha$ *gets sent to* $1 \in k_{q(\alpha)}$ *if* $\alpha \in S$ *and* 0 *otherwise. In other words, the map adds all the factors corresponding to vertices that get identified by the quotient* q*. Moreover, if one picks an orientation of a top stratum of* $\mathbb{T}$*, there is a canonical choice of isomorphisms above.*

The proof follows from a direct computation of the dualizing complex $\omega_{\mathbb{T}}$, for which we use a decomposition of the singularity into discs 3.3 and use Mayer-Vietoris for relative homology. See [100, Sec.4] for the full proof.

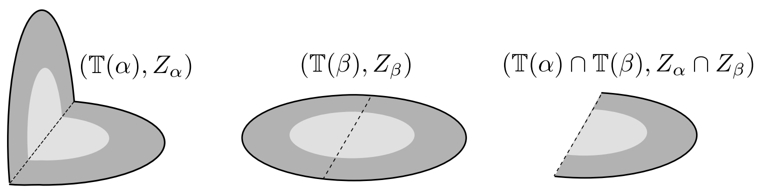

Figure 3.3: A pair of discs appearing in the Mayer-Vietoris decomposition of $\omega_{\mathbb{A}_3}$. The relative homology of their intersection on the right vanishes relative to the outer boundary

The Hochschild homology and the cyclic homology sheaf

Recall that the Hochschild homology of an algebra A is calculated by the Hochschild chain complex

$$\mathrm{HH}(A) : \cdots \to A \otimes A \otimes A \to A \otimes A \to A$$

where $A^{\otimes n}$ is placed in degree $-n$, and the differential $d_{-n} : C_{-n}(A) \to C_{-n+1}(A)$ given by

$$d(x_n \otimes \cdots \otimes x_0) = x_{n-1} \otimes \cdots \otimes x_0 x_n + \sum_{i=0}^{n-1} (-1)^i x_n \otimes \cdots \otimes x_{i+1} x_i \otimes \cdots \otimes x_0$$

Remark. Recall that even though we denote HH we are still using a cohomological (increasing) grading for all our complexes; in most definitions of Hochschild homology this is given in positive degrees and graded homologically but we invert the degrees since this is more natural in the context of Hochschild homology as a left derived tensor product.

We are interested in the stalks of the Hochschild homology sheaf $\mathcal{HH}(\mathcal{N})$ and of the cyclic homology sheaf $\mathcal{HH}(\mathcal{N})_{S^1}$, which are the Hochschild/cyclic homology of the stalks of $\mathcal{N}$, i.e., of the categories $\mathrm{Perf}(\vec{T})$. To calculate these, recall that Hochschild/cyclic homology is invariant under dg Morita equivalences, and for any dg-algebra $\mathcal{A}$, there is a quasi-isomorphism of Hochschild complexes $\mathrm{HH}(\mathcal{A}) \cong \mathrm{HH}(\mathrm{Perf}\text{-}\mathcal{A})$ with the Hochschild homology of the category of perfect $\mathcal{A}$-modules, with compatible S^1 actions Thus we need only know the Hochschild homology $\mathrm{HH}_*(k[\vec{T}]) = H^*(\mathrm{HH}(k[\vec{T}]))$ and the circle action on it. But this is a well known result in the case of an acyclic quiver:

Proposition 28. *[27] Let $\vec{Q}$ be an acyclic quiver. Then $\mathrm{HH}_0(k[\vec{Q}]) = k^{|Q|}$ and all higher Hochschild homologies vanish. Moreover the S^1 action on the Hochschild complex is trivial, i.e. the cyclic complex is given by*

$$\mathrm{HH}(k[\vec{Q}])_{S^1} = H^*(BS^1) \otimes \mathrm{HH}(k[\vec{Q}]) = k[u] \otimes \mathrm{HH}(k[\vec{Q}])$$

with u being the canonical generator of $H^*(BS^1)$ in degree 2, and the map $\mathrm{HH}(k[\vec{Q}]) \to \mathrm{HH}(k[\vec{Q}])_{S^1}$ sends $x \mapsto 1 \otimes x$

Proof. For convenience of the reader, we indicate the proof. Since $\vec{Q}$ is a tree, $k[\vec{Q}]$ has a basis whose elements are the paths from vertex α to vertex β. These include the idempotents $|\alpha\rangle\langle\alpha|$. Consider the subspace Δ_n of $k[\vec{Q}]^{\otimes n}$ spanned by the powers of the idempotents $|\alpha\rangle\langle\alpha|^{\otimes n}$, and its complement L_n spanned by all other tensor products of paths.

The Hochschild chain complex then splits as $C_*(k[\vec{Q}]) = \Delta_* \oplus L_*$. The complex L_* is acyclic, ultimately because $\vec{Q}$ has no cycles. The diagonal subcomplex Δ_* is $|Q|$ copies of the Hochschild chain complex for the base field k.

As for the cyclic homology, consider Connes' long exact sequence connecting the Hochschild homology $\mathrm{HH}_*(k[\vec{T}]) = H^*(\mathrm{HH}(k[\vec{T}]))$ and cyclic homology $\mathrm{HC}_*(k[\vec{T}]) = H^*(\mathrm{HH}(k[\vec{T}])_{S^1})$:

$$\cdots \to \mathrm{HH}_n(k[\vec{T}]) \to \mathrm{HC}_n(k[\vec{T}]) \to \mathrm{HC}_{n-2}(k[\vec{T}]) \to \mathrm{HH}_{n-1}(k[\vec{T}]) \to \cdots$$

The result then follows immediately. □

Since all the S^1 actions we consider will be trivial, we will ignore it from now on; every map out of $\mathrm{HH}(k[\vec{T}])$ can be factored through the map $\mathrm{HH}(k[\vec{T}]) \to \mathrm{HH}(k[\vec{T}])_{S^1}$ by sending $u \mapsto 0$, so for all our applications we can just construct maps out of/into the Hochschild homology HH itself.

There is a natural basis on $\mathrm{HH}_0(\mathrm{Perf}(\vec{Q})) = \mathrm{HH}_0(k[\vec{Q}]) \cong k^{|Q|}$, given by the images of the idempotents $|\alpha\rangle\langle\alpha| \in k[\vec{Q}]$, or equivalently, of the modules $P_\alpha \in \mathrm{Perf}(\vec{Q})$. Note that this basis depends on $\vec{Q}$ and not just the underlying graph. In terms of these bases, Lemma 23 gives the generization functors of the Hochschild homology sheaf.

Proposition 29. *Let $\mathfrak{p} = (R \twoheadleftarrow S \hookrightarrow T)$ be a correspondence inducing a functor $\mathrm{Perf}(\vec{T}) \to \mathrm{Perf}(\vec{R})$. The induced map between Hochschild homologies is given by*

$$|\alpha\rangle\langle\alpha| \mapsto \begin{cases} |q(i^{-1}(\alpha))\rangle\langle q(i^{-1}(\alpha))| & \text{if } \alpha \in i(S) \\ 0 & \text{otherwise} \end{cases}$$

Comparison

Comparing Proposition 27 with Propositions 28 and 29 gives an abstract isomorphism $\mathcal{HH}(\mathcal{N}) \cong \omega_{\mathbb{T}}[1-n]$.

Remark. Again, technically this is an isomorphism of sheaves only on the open arboreal singularity $\mathbb{T}_\circ$, but this difference will be of no effect to our calculations.

The choice of this isomorphism is not unique, but as we saw above, upon fixing the decomposition of $\mathbb{T}$ as the union of discs $\mathbb{T}(\alpha)$ and an orientation of one of these discs, we get distinguished bases for the stalks of $\omega_{\mathbb{T}}[1-n]$. In addition if we pick a root in T this induces choices of roots in all R, and we get sets of distinguished elements $|\alpha\rangle\langle\alpha|$ in all the stalks of $\mathcal{HH}(\mathcal{N})$. We can then make a *canonical* choice of isomorphism $\mathcal{HH}(\mathcal{N}) \xrightarrow{\sim} \omega_{\mathbb{T}}[1-n]$, which on a stalk over the stratum $\mathbb{T}(R \twoheadleftarrow S \hookrightarrow T)$ gives the isomorphism

$$\mathrm{HH}_0(k[\vec{R}]) \xrightarrow{\sim} \bigoplus_{\alpha \in V(R)} k_\alpha$$

sending $|\alpha\rangle\langle\alpha|$ to $1 \in k_\alpha \cong k$ in the direct sum decomposition of proposition 27.

Nondegeneracy

Recall we have constructed an isomorphism $\mathcal{HH}(\mathcal{N}) \xrightarrow{\sim} \omega_{\mathbb{T}}[-d]$, with $d = |T| - 1$; moreover since the S^1 action on $\mathcal{HH}(\mathcal{N})$ is trivial, this naturally descends to an isomorphism $\mathcal{HH}(\mathcal{N})_{S^1} \xrightarrow{\sim} \omega_{\mathbb{T}}[-d]$

Theorem 30. *The local orientation given by the map $\mathcal{HH}(\mathcal{N})_{S^1} \xrightarrow{\sim} \omega_{\mathbb{T}}[-d]$ constructed above is nondegenerate in the sense of Definition 16.*

The proof follows from a Verdier duality argument; see [100, Sec.4.1] for the full argument.

3.5 Global orientations

We have constructed above local orientations on the local models $(\mathbb{T}, \mathcal{W}_{\mathbb{T}})$, by giving a nondegenerate orientation on the sheaf $\mathcal{N}_{\mathbb{T}} = \mathcal{W}_{\mathbb{T}}^{pp}$ of pseudo-perfect modules. For an arbitrary locally arboreal space $(\mathbb{X}, \mathcal{W})$, it follows that the Hochschild homology sheaf $\mathcal{HH}(\mathcal{W}^{pp})$ is locally isomorphic to the dualizing complex, and has trivial S^1 action. In this section we study the obstruction to global orientability, and note a class of examples in which it vanishes.

The obstruction to orientability

Note, by Verdier duality, $\mathscr{H}om(\omega_{\mathbb{X}}, \omega_{\mathbb{X}}) = k_X$. It follows that on a locally arboreal space $(\mathbb{X}, \mathcal{W})$, we have $\mathcal{HH}(\mathcal{W}^{pp}) \cong \omega_X \otimes \mathcal{L}$, for some locally constant rank one sheaf $\mathcal{L}$. Such sheaves are classified by $H^1(\mathbb{X}, k^*)$.

Theorem 31. *The obstruction to orientability is the image of a class $w_1(\mathbb{X}, \mathcal{W}) \in H^1(\mathbb{X}, \pm 1)$.*

Proof. Rather than work with sheafified Hochschild homology, we could have worked with a sheafified Grothendieck group $\mathcal{K}_0$. The Dennis trace map from K-theory to Hochschild homology induces a morphism $\mathcal{K}_0(\mathcal{W}^{pp}) \to \mathcal{HH}(\mathcal{W}^{pp})$. Evidently $K_0(\mathrm{Perf}(\vec{T})) = \mathbb{Z}^{|T|}$. This morphism becomes an isomorphism after tensoring with k. The above matching with the dualizing sheaf would have all worked just as well for $\mathcal{K}_0$ as $\mathcal{HH}$, except now we can work over $\mathbb{Z}$. In particular we see that $\mathcal{K}_0(\mathcal{W}^{pp}) \cong \omega_X \otimes \mathcal{L}$ over $\mathbb{Z}$. Thus $\mathcal{L}$ is classified by an element of $H^1(\mathbb{X}, \mathbb{Z}^*) = H^1(\mathbb{X}, \pm 1)$. □

Remark. The space of choices of possible $\mathcal{W}$ over a given locally arboreal space $\mathbb{X}$ is a torsor over

$$H^1(\mathbb{X}, \mathcal{A}ut(\mathcal{W})) \ltimes H^2(\mathbb{X}, \mathcal{A}ut(1_{\mathcal{W}})) = H^1(\mathbb{X}, \mathbb{Z}) \ltimes H^2(\mathbb{X}, k^*)$$

There are no other terms because there is no higher local automorphisms of the cosheaf $\mathcal{W}$, since Hochschild cohomology of the tree quivers is just k in degree zero and nothing else [111]. Here, the fact that the connected components of the local automorphisms of the cosheaf of categories $\mathcal{W}$ are just the shift functor can be seen by observing that, for an arboreal singularity $\mathbb{T}$, the restriction from $\mathbb{T}$ to the smooth locus of $\mathbb{T}$ remembers the subcategories generated by every indecomposable.

Our $w_1(\mathbb{X}, \mathcal{W})$ is the reduction mod 2 of the above H^1 information; which takes values in a vector space rather than a torsor because we have now the basepoint given by comparison with the dualizing sheaf.

Global orientations from immersed front projections

The Kashiwara-Schapira sheaf

One way in which locally arboreal spaces $(\mathbb{X}, \mathcal{W})$ can arise is by taking, inside the cotangent bundle T^*M of an ambient manifold M, the union of the zero section and a cone over a general position Legendrian which itself has arboreal singularities. The sheaf $\mathcal{W}^{pp}$ then arises as the restriction of the so-called Kashiwara-Schapira sheaf of categories [81], and the cosheaf $\mathcal{W}$ is the cosheaf of "wrapped microlocal sheaves" described in [85].

The resulting categories are already quite rich. On the one hand, they provide powerful invariants in symplectic and contact geometry which are closely related to, but conceptually simpler than, the holomorphic curve invariants. On the other hand, many spaces of interest for other reasons can be constructed as moduli of objects in these categories, e.g., positroid varieties, cluster algebras from surfaces, and wild character varieties [103, 102]. We will recall some explicit examples in Section 3.6.

Let us briefly recall the notions of microlocalization and of the Kashiwara-Schapira sheaf of categories. Given a sheaf $\mathcal{F}$ on a manifold M, the locus of codirections in which the sections fail to propagate is called the microsupport of $\mathcal{F}$. The properties of the microsupport are developed in [67], where in particular it is shown that the microsupport is a conical co-isotropic subset of T^*M, which is Lagrangian if and only if $\mathcal{F}$ is constructible.

Let M be a manifold and $L \subset T^*M$ a conical Lagrangian. We write $sh_L(M)$ for the category of sheaves with microsupport in L. A fundamental result is that the category of sheaves with microsupport in L localizes, not only over M, but in fact over L.

Theorem 32. *[67, Chap. 6] The category $sh_L(M)$ of sheaves on M with microsupport in L is the global sections of a constructible sheaf of categories on L, obtained by sheafifying the presheaf*

$$U \mapsto \{sh_L(M)\}/\{sh_{L \cup T^*M \setminus U}(M)\}$$

We call this sheaf the Kashiwara-Schapira sheaf on L, and write it as μloc.

In the neighborhood of a smooth point of L, there is a non-canonical isomorphism from μloc to the category of derived local systems on L.

In particular, a (possibly singular) Legendrian $\Lambda \subset T^\infty M$ carries a Kashiwara-Schapira sheaf, given by restricting the Kashiwara-Schapira sheaf from the union of M with the positive cone over Λ. It follows from the theory of contact transformations developed in [67, Sec. 7] that the stalk of this sheaf at a point depends only on the local contact geometry. A more global version of this statement appears in [54].

The relation to the sheaves of categories on the locally arboreal spaces is the following:

Theorem 33. *[81] For each a rooted tree $\vec{T}$, there is a Legendrian embedding $\mathbb{T} \hookrightarrow T^\infty \mathbb{R}^T$ and a canonical isomorphism $\mathcal{N} \cong \mu loc$.*

Remark. It follows from the theory of contact transformations that for *any* Legendrian embedding $\mathbb{T} \hookrightarrow T^\infty \mathbb{R}^T$ (with behavior at the singularities constrained in a sense clarified in [84]), there is again an isomorphism $\mathcal{N} \cong \mu loc$. The union of the cone over $\mathbb{T}$ and the zero section $\mathbb{R}^T$ is again an arboreal singularity, corresponding now to the graph obtained by attaching one vertex below the root of $\vec{T}$ and making this the new root. Again the sheaf μloc on this larger space is identified with the sheaf $\mathcal{N}$.

Immersed front projections

We say that a Legendrian $\Lambda \subset T^\infty M$ has an immersed front projection when the projection $\Lambda \to M$ is an immersion. In this case, there is a natural identification of the Kashiwara-Schapira sheaf with the category of local systems on Λ, i.e., $\mu loc(\Lambda) = loc(\Lambda)$.

We say that the front projection has *normal crossings* when it is locally diffeomorphic to a union of coordinate hyperplanes.

Lemma 34. *Let M be a manifold, $\Lambda \subset T^\infty M$ a smooth Legendrian with normal crossings projection Let $m \in M$ be a point where the front projection of Λ is immersed with normal crossings image. Let U be a conical neighborhood of m in $\mathbb{X} = M \cup \mathbb{R}_+\Lambda$. Then there is a canonical isomorphism of $(U, \mu loc)$ with (a trivial factor times) an arboreal singularity corresponding to a star quiver, given by a single root vertex with as many leaves as there points of Λ over m.*

Proof. Let $\pi : T^\infty M \to M$, and $d = \dim M$. Around a smooth point of $\pi(\Lambda)$ *on* the base M, U is homeomorphic to $\mathbb{R}^d$ with a d-dimensional half-space glued to a coordinate hyperplane, which is homeomorphic to the local model $\mathbb{A}_2 \times \mathbb{R}^{d-1}$.

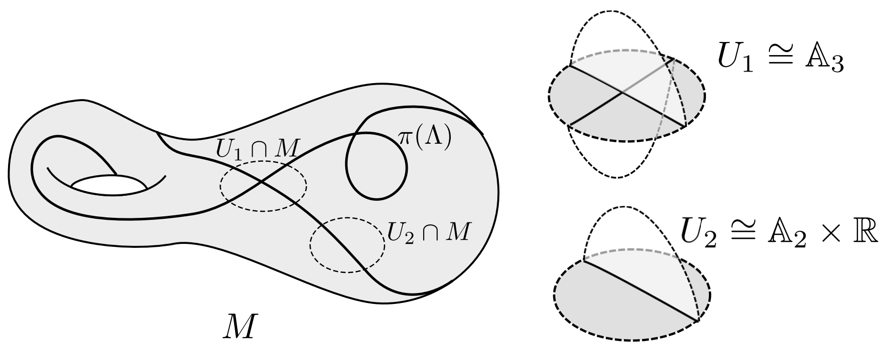

Figure 3.4: The case $d = 2$. The neighborhood of a point in the front projection of the Legendrian Λ is homeomorphic to either $\mathbb{A}_2 \times \mathbb{R} = \mathbb{S}tar_1 \times \mathbb{R}$ or $\mathbb{A}_3 = \mathbb{S}tar_2$. Note that in general it is the star quivers $Star_k$ that appear, and not the A_k series

Similarly, around a singular point of $\pi(\Lambda)$ by the normal crossings condition, U is homeomorphic to $\mathbb{R}^d$ with k half-spaces glued along k coordinate hyperplanes. Consider now the arboreal singularity $\mathbb{S}tar_k$ corresponding to the star quiver with k leaves; this is homeomorphic to $\mathbb{R}^k$ with k half-spaces glued along coordinate hyperplanes. By comparison we have a homeomorphism $U \cong \mathbb{S}tar_k \times \mathbb{R}^{d-k}$. □

Consider now the Kashiwara-Schapira sheaf $\mu loc_{\mathbb{X}}$. By the inspection above and results of Nadler [81, 84], on a neighborhood homeomorphic to some arboreal singularity model $U \cong \mathbb{S}tar_k \times \mathbb{R}^{d-k}$, this sheaf is locally isomorphic to the Nadler sheaf of categories $\mathcal{N}_{\mathbb{S}tar_k}$, and there is a corresponding cosheaf $\mu loc^w_{\mathbb{X}}$ of wrapped microlocal constructible sheaves. So $(\mathbb{X}, \mu loc^w_{\mathbb{X}})$ is a locally arboreal space, with a sheaf of categories $\mu loc_{\mathbb{X}} = (\mu loc^w_{\mathbb{X}})^{pp}$, and by the following result a sufficient criterion for its orientability is the orientability of the base manifold M.

Theorem 35. *Let M be a compact oriented manifold of dimension d, with a given orientation $\underline{k}_M \xrightarrow{\sim} \omega_M[-d]$. Let $\Lambda \subset T^\infty M$ be a smooth Legendrian whose projection to M has normal crossings, and $\mathbb{X} = M \cup \mathbb{R}_+\Lambda$. Then the locally arboreal space $(\mathbb{X}, \mu loc^w_{\mathbb{X}})$ admits a local orientation, given by an isomorphism $\mathcal{HH}(\mu loc_{\mathbb{X}}) \xrightarrow{\sim} \omega_{\mathbb{X}}[-d]$, extending the orientation on M, i.e. which agrees with the given orientation on M when restricted to the smooth locus on M.*

This follows from the fact that the coorientation of the front projection of the Legendrian is globally defined in a coherent way; see [100, Sec.5] for the proof.

3.6 Examples

In this section we recall various categories which arise as the global sections of a sheaf or cosheaf of categories on an oriented locally arboreal space. As a consequence of our results, we conclude the existence of certain absolute and relative Calabi-Yau structures.

To relate this discussion to the subject of shifted symplectic geometry, in this section we will use the main theorem of [21]. This theorem determines, from an absolute d-dimensional smooth Calabi-Yau structure on a category $\mathcal{A}$, a $(2-d)$-shifted symplectic structure on the the moduli $\mathcal{M}_{\mathcal{A}}$ of objects in $\mathcal{A}^{pp}$, and from a relative d-dimensional smooth Calabi-Yau structure on $f : \mathcal{A} \to \mathcal{B}$, a Lagrangian structure on the morphism $\mathcal{M}_{\mathcal{B}} \to \mathcal{M}_{\mathcal{A}}$ into the $(3-d)$-shifted symplectic space $\mathcal{M}_{\mathcal{A}}$.

Assuming the result above will give rise to shifted symplectic structures, Poisson structures, quantizations, etc. in the appropriate circumstances [91, 25]. Direct application of our methods yield constructions in the "type A" cases; e.g. the moduli space for a point is $\mathcal{M}_{\mathrm{Perf}}$ inside which one can find the various BGL_n. We expect that the desired symplectic structures for the analogous moduli spaces for other groups can be constructed via Tannakian considerations as in [104, Sec. 6], but do not develop this in detail here.

Remark. In this section to ease notation we will denote the moduli stack of objects using a blackboard bold capital to distinguish it from the category, e.g. the category of perfect complexes is denoted Perf and the moduli stack of perfect complexes is $\mathbb{P}\mathrm{erf} = \mathcal{M}_{\mathrm{Perf}}$, same for local systems Loc and $\mathbb{L}\mathrm{oc}$.

The associated graded of a filtration

Let $\mathbb{X}$ be a comb: a one-dimensional space formed as the union of $\mathbb{R}$ and the positive cone on some n points $\{p_i\}$ at positive contact infinity. The category $Sh_{\{p_i\}}(\mathbb{R})$ is equivalent to the category $\mathbb{F}\mathrm{ilt}_n$ of n-step filtered perfect complexes, which just means sequences of perfect complexes

$$F_0 \to F_1 \to \ldots \to F_n$$

Corollary 36. *The functor*

$$\begin{array}{rcl} \mathrm{Filt}_n & \to & \mathrm{Perf}_k^{\otimes(n+1)} \times \mathrm{Perf}_k \\ F_0 \to \ldots \to F_n & \mapsto & (F_0, Cone(F_0 \to F_1), Cone(F_1 \to F_2), \ldots, Cone(F_{n-1} \to F_n)),\ F_n \end{array}$$

has a 1*-dimensional relative proper Calabi-Yau structure, and so assuming the theorem in [20] the corresponding map of moduli spaces of objects*

$$\mathbb{F}\mathrm{ilt} \to \mathbb{P}\mathrm{erf}^{\times(n+1)} \times \mathbb{P}\mathrm{erf}$$

is a Lagrangian mapping to a 2*-shifted symplectic space.*

Proof. This follows immediately from Theorem 35. Alternatively, the fact that the comb is orientable follows from its being contractible. □

Remark. Note that following the description in 35, the decomposition into "discs" (here intervals) is such that the boundary of the ith interval is given by the endpoints $(+\infty)-(p_i)$, if we pick the positive orientation on the base manifold $\mathbb{R}$. So in the Lagrangian map of moduli stacks the first factor $\mathbb{P}\mathrm{erf}^{\times(n+1)}$ is endowed with the opposite 2-shifted symplectic structure, where the last factor $\mathbb{P}\mathrm{erf}$ has the usual 2-shifted symplectic structure.

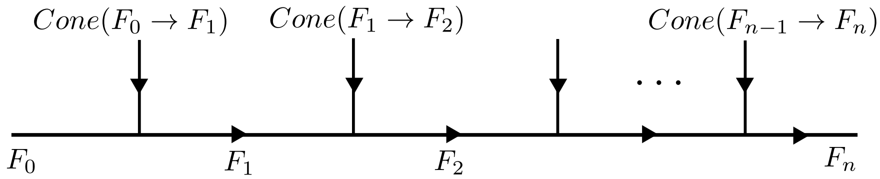

Figure 3.5: The comb $\mathbb{X}$. We pick a decomposition of $\mathbb{X}$ into overlapping intervals; in this case the ith interval has endpoints p_i and $+\infty$, and extending the positive orientation on $\mathbb{R}$, its boundary is $(+\infty)-(p_i)$.

Restricting to the open substack $\mathbb{F}\mathrm{ilt}_n^\circ$ where all F_i and all $Cone(F_i \to F_{i+1})$ can be represented by vector spaces in degree zero, we see that the image of the morphism above in each factor $\mathbb{P}\mathrm{erf}$ lands in some substack $\mathbb{B}GL_{m_i} \subset \mathbb{P}\mathrm{erf}$, and moreover that each map $F_i \to F_{i+1}$ is injective. Denoting $m_0 = \dim F_0, m_i = \dim Cone(F_i \to F_{i+1})$ and $m = \dim F_n$, this implies that $m_0 + \cdots + m_n = m$. So $\mathbb{F}\mathrm{ilt}_n^\circ$ splits into components labelled by collections of integers $m, \{m_i\}$, and each component has a 2-shifted Lagrangian morphism

$$\mathbb{F}\mathrm{ilt}_n^\circ(m, \{m_i\}) \to (\mathbb{B}GL_{m_0} \times \cdots \times \mathbb{B}GL_{m_n}) \times \mathbb{B}GL_m$$

which can be interpreted as a 2-shifted Lagrangian correspondence between $\mathbb{B}GL_m$ and $\mathbb{B}L$ for the Levi subgroup corresponding to the partition $\{m_i\}$. This was originally shown in [94], where a Lagrangian correspondence

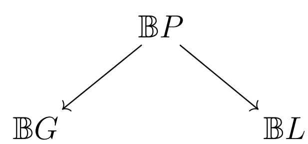

is used to define the "partial group-valued symplectic implosion". Note that once we fix the decomposition $m = m_0 + \cdots + m_n$, we can identify the substack $\mathbb{F}\mathrm{ilt}_n^\circ(m, \{m_i\})$ as the classifying space $\mathbb{B}P$ for the parabolic P corresponding to L: a map from some other space $X \to \mathbb{F}\mathrm{ilt}_n^\circ(m, \{m_i\})$ determines an invariant filtration of the vector space k^m, therefore up to equivalence it is the data of a P-bundle over X.

Invariant filtrations near punctures on surfaces

Let Σ be an oriented surface with boundary consisting of n circles; draw a collection of m_i concentric circles at each boundary component, and choose co-orientations and therefore

Legendrian lifts. Let Λ denote the union of these lifts. Just as in the previous subsection, the corresponding category $Sh_\Lambda(\Sigma)$ amounts to the category of sheaves on Σ with invariant filtrations at the punctures. Note that $\Lambda \cup \partial\Sigma$ is just a union of $m = \sum m_i$ circles.

Corollary 37. *The morphism* $\mathbb{S}\mathrm{h}_\Lambda(\Sigma) \to \mathbb{L}\mathrm{oc}(\Lambda \cup \partial\Sigma) = \mathbb{L}\mathrm{oc}(S^1)^{\times m}$ *is* 1*-shifted Lagrangian.*

Proof. This follows immediately from Theorem 35. □

Again, we can take the substack $\mathbb{S}\mathrm{h}_\Lambda(\Sigma)^\circ$ on which the restriction of the sheaf to Σ has cohomology concentrated in degree zero; the resulting space is an Artin stack is the classical sense, and as in the comb example above, it has components labelled by the microlocal ranks along the boundary components; the morphism splits into components

$$\mathbb{S}\mathrm{h}_\Lambda(\Sigma)^\circ(\{m_i\}) \to [\frac{GL(m_1)}{GL(m_1)}] \times \cdots \times [\frac{GL(m_n)}{GL(m_n)}]$$

where the stacky quotient is taken with respect to the adjoint action.

To get a space with a symplectic structure, one can choose correspondingly another 1-shifted Lagrangian morphism to the moduli space of local systems around the boundary, and performing Lagrangian intersection between the two 1-shifted Lagrangians as in [93] this gives a 0-shifted symplectic space.

Corollary 38. *The moduli space of local systems on a surface equipped with invariant filtrations at the punctures, of which the conjugacy classes C_i of the associated graded holonomies are pre-specified, carries a* 0*-shifted symplectic structure.*

Proof. As in Safronov [93], this can be deduced from the above corollary (which has a different proof in that paper) by observing that fixing a conjugacy class C in a reductive group G determines a Lagrangian morphism $[\frac{C}{G}] \to [\frac{G}{G}]$. Performing Lagrangian intersection between $\mathbb{S}\mathrm{h}_\Lambda(\Sigma)^\circ(\{m_i\})$ and $[\frac{C_1}{GL(m_1)}] \times \cdots \times [\frac{C_n}{GL(m_n)}]$ gives the symplectic structure on the moduli space with prescribed holonomies. □

Example. Consider Σ an oriented surface with boundary $\partial\Sigma$ = union of n circles, without any Legendrians. The Artin stack $\mathbb{S}\mathrm{h}(\Sigma)^\circ$ has disjoint components $\mathbb{L}\mathrm{oc}_{GL(m)}(\Sigma)$ labelled by the rank m of the local system, and so we have 1-shifted Lagrangian morphisms

$$\mathbb{L}\mathrm{oc}_{GL(m)}(\Sigma) \to (\mathbb{L}\mathrm{oc}_{GL(m)}(S^1))^{\times n} = \left[\frac{GL(m)}{GL(m)}\right]^{\times n}$$

and picking n conjugacy classes C_i in G, we can perform the intersection and get the so-called "tame" character variety of Σ.

Example. Take Σ to be the open cylinder with n concentric circles around one of the boundary components, and no circles around the other. Fixing the ranks $m = m_1 + \cdots + m_n$ at each boundary, we get a 1-shifted Lagrangian morphism

$$\mathbb{S}\mathrm{h}_\Lambda(\Sigma)^\circ \to \mathbb{L}\mathrm{oc}_{GL_m}(S^1) \times \mathbb{L}\mathrm{oc}_{GL_{m_1}}(S^1) \times \cdots \times \mathbb{L}\mathrm{oc}_{GL_{m_n}}(S^1) = \left[\frac{G}{G}\right] \times \left[\frac{L}{L}\right]$$

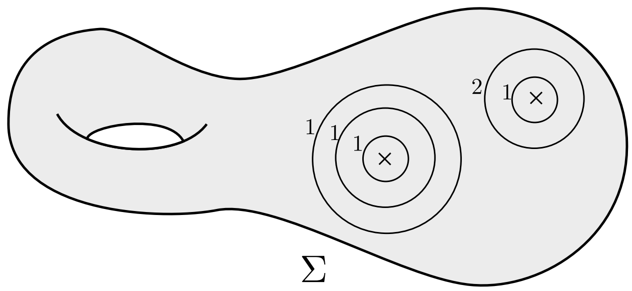

Figure 3.6: On the surface Σ with two punctures, and with fixed microlocal rank along the concentric circles of Λ. If we look at the substack of $Sh_\Lambda(\Sigma)$ of objects with rank 0 at the punctures, the microlocal rank conditions mean we have rank 3 local systems equipped with invariant filtrations near each puncture; in this particular case we have two filtrations respectively of the form $0 \subset k \subset k^2 \subset k^3$ and $0 \subset k \subset k^3$

where L is the Levi subgroup corresponding to the partition $\{m_i\}$. This example also appears in [93], as a 1-shifted Lagrangian correspondence

$$\begin{array}{ccc} & \left[\frac{P}{P}\right] & \\ \swarrow & & \searrow \\ \left[\frac{G}{G}\right] & & \left[\frac{L}{L}\right] \end{array}$$

where the identification of our space with $[\frac{P}{P}]$ comes from the observation that an invariant filtration on Σ is the same data as a P-local system on S^1.

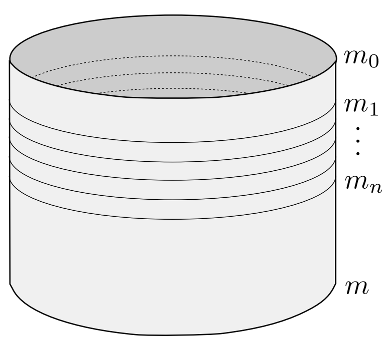

Figure 3.7: The restriction to the upper boundary components can be assembled into a map to $[\frac{L}{L}]$ where L is a Levi subgroup of $G = GL_m$ given by the integers $\{m_i\}$. The restriction to the lower boundary components is a map to $[\frac{G}{G}]$ giving the monodromy of the G-local system

Another way of deducing this particular case is by writing Σ as the product of a circle and a comb (as in 3.6). Therefore we have an equivalence of derived stacks

$$\mathbb{S}\mathrm{h}_\Lambda(\Sigma) = \mathbb{R}Map(S^1, \mathbb{F}\mathrm{ilt}_n)$$

and the Lagrangian correspondence is obtained from the correspondence in section 3.6 by applying the mapping stack functor $\mathbb{R}Map(S^1, -)$, which shifts the degree down to the 1-shifted Lagrangian correspondence above.

Remark. In the rank one case, when $k_i = 1$, the correspondence $[\frac{B}{B}] \to [\frac{G}{G}] \times [\frac{T}{T}]$ is a group version of the Grothendieck-Springer correspondence, and we have $[\frac{B}{B}] \cong [\frac{\tilde{G}}{G}]$ where G acts on the Springer resolution $\tilde{G}$ by conjugation.

Stokes filtrations near punctures on surfaces

Rather than take an invariant filtration around a puncture, one can allow the filtration itself to undergo monodromy. The resulting notion generalizes the notion of Stokes structure; we will call it a Stokes filtration. In most works, this was presented as suggested above: in terms of a sheaf on the boundary circle equipped with a filtration that itself varies. This notion can be found e.g. in [80]. Defining what precisely it means for a filtration to vary along a circle is nontrivial and somewhat mysterious at the points where the steps in the filtration cross.

We prefer to turn this notion sideways: rather than a filtered sheaf on S^1 with varying filtration, we take a sheaf on $S^1 \times \mathbb{R}$ with microsupport in a prescribed Legendrian braid closure. This determines a filtration in the $\mathbb{R}$ direction, just as in the previous examples; as it happens, the above notion exactly captures at the crossings the notion in [80] of Stokes filtration. This idea seems to have been known to the experts, but we have not found any systematic exposition of it in the classical literature.[2]

We made some attempt in this direction in [103, Sec. 3.3]. Here we simply recall that the Deligne-Malgrange account of the irregular Riemann-Hilbert correspondence on Riemann surfaces can be formulated as follows. Suppose we are given a Riemann surface Σ with marked points p_i, and a specification of a (possibly ramified) irregular type τ_i, i.e., formal equivalence class of irregular singularity, at each. Then there is an associated Legendrian link $\Lambda = \coprod \Lambda(\tau_i)$, a union of links localized near the p_i, and an equivalence of categories between the irregular connections with these singularities, and the full subcategory of $Sh_\Lambda(\Sigma)$ on objects which have cohomology concentrated in degree zero, and appropriate rank stalks and microstalks. The corresponding component of the moduli space is the moduli space of Stokes data. Finally, the microlocal restriction morphism $Sh_\Lambda(\Sigma) \to Loc(\Lambda)$ is what would have classically been called "taking the formal monodromies".

[2] More precisely, we only know the following other occurrences of this picture. In [34], there is a letter from Deligne in which the Stokes sheaf is viewed as a sheaf on an annulus rather than a filtered sheaf on a line. In [68], the idea that a Legendrian knot can be associated to a Stokes filtration appears as a remark. Finally, the drawing of at least the projection of a knot already appears in the original work of Stokes [106].

Corollary 39. *The morphism from a moduli space of Stokes data to the moduli space of formal monodromies is 1-shifted-Lagrangian.*

Corollary 40. *A moduli space of Stokes data with formal monodromies taking values in prescribed conjugacy classes is 0-shifted-symplectic. In particular, any open substack which happens to be a scheme is symplectic in the usual sense.*

This recovers and generalizes all constructions of symplectic structures in e.g. [16, 17, 18] for GL_n connections.

Example. (Wild character variety) Consider a disc Σ punctured at the origin and a trivial rank n vector bundle $E \to \Sigma$, where the origin is marked with the irregular type

$$Q(z) = \frac{A}{z^r}, \qquad r \in \mathbb{Z}, A \in \mathfrak{t}^{reg} \subset \mathfrak{gl}_n$$

i.e. A has all distinct eigenvalues. A meromorphic connection ∇ on E has this irregular type if it can be brought by a local analytic gauge transformation to the connection defined by the connection one-form dQ.

By the irregular Riemann-Hilbert correspondence, the category of meromorphic connections on the trivial vector bundle with this irregular type is equivalent to the category of Stokes data. From the Stokes data, one can recover the monodromy of ∇ and the formal monodromy; this latter is an element of the centralizer $Z_G(A)$ (in this case, the maximal torus T) up to conjugation.

The moduli of Stokes data with fixed monodromies and formal monodromies is commonly known as the *wild character variety*. Upon fixing conjugacy classes C_G, C_T in G and T, the wild character variety can be described as a quasi-Hamiltonian quotient [16]

$$(G \times T \times (U_+ \times U_-)^r) /\!\!/_{C_G, C_T} (G \times T)$$

where $U_\pm$ are the unipotent subgroups corresponding to the maximal torus T. The moment map to G is taking the monodromy around the singularity, and the map to T is taking the formal monodromy.

In our description, this category of Stokes data becomes a full subcategory of the category $Sh_\Lambda(\Sigma)$ of microlocal sheaves, for a corresponding Legendrian link $\Lambda \subset T^\infty$ around the singularity. The monodromy and formal monodromy then become literal monodromies of the local systems one gets by restriction to the boundary. To explicitly construct Λ, one can follow the prescriptions in [103, Sec. 3.3]. This can be heuristically stated in terms of the asymptotics of flat sections, i.e. the growth behavior of the solutions to

$$\frac{df}{dz} = \frac{dQ}{dz} f(z)$$

In this case, the solutions are spanned by n different solutions $f_i \sim \exp(\lambda_i z^{-r})$, where λ_i are the eigenvalues of A, and we only keep the exponential part of the asymptotics. The Stokes phenomenon refers to the fact that these corresponds to asymptotics of solutions in different sectors; as we go from one sector to the other, the growth of these solutions changes. On each sector, we draw concentric strands for the f_i, ordering them by growth: the faster-growing ones further from the origin. Whenever we cross a Stokes ray, where the solutions f_i, f_j switch growth asymptotics, we introduce a crossing between the i and j strands.

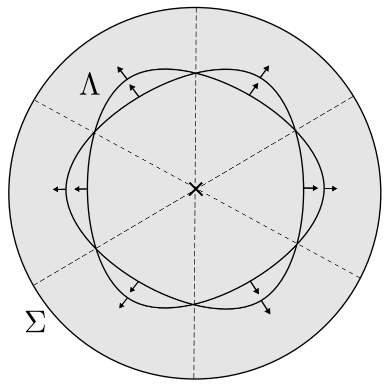

Figure 3.8: The Legendrian link for the irregular type $1/z^3$. The Legendrian link $\Lambda \subset T^\infty\Sigma$ is obtained by lifting the projection using the outward coorientation. The dashed lines are the Stokes rays, where the asymptotics of formal solution changes. Note that each component of the link is unknotted with itself: this is true of all the examples of this form.

In the case we described above ($A \in \mathfrak{t}^{reg}$), the corresponding Legendrian link Λ is the closure of a $(n, 2r)$ braid , cooriented outward, where we enforce the condition that the rank of the stalk inside Λ is zero, and the microlocal ranks on each component of Λ is one. What we call the moduli of Stokes data $\mathcal{M}_\Lambda$ is the moduli of objects in the full subcategory of $Sh_\Lambda(\Sigma)^\circ$ with those rank conditions. The maps given by restriction to the boundary components can be assembled into a 1-shifted Lagrangian map

$$\mathcal{M}_\Lambda \to \left[\frac{G}{G}\right] \times \left[\frac{T}{T}\right]$$

where $[\frac{G}{G}]$ is G-local systems on the boundary of the disc, and $[\frac{T}{T}]$ is rank one local systems on Λ. In this description, taking quasi-Hamiltonian quotient corresponds to taking intersection with another Lagrangian $[\frac{C_G}{G}] \times [\frac{C_T}{T}]$. The explicit description of the moduli space $\mathcal{M}_\Lambda$ can be obtained by following the prescriptions in [102]; one can check that this stack can in fact be expressed as the quotient $[(G \times T \times (U_+ \times U_-)^r)/(G \times T)]$, agreeing with the previously existing description.

Example. With the same notation of the previous example, consider the irregular type

$$Q = \frac{A}{z^{r/2}}, \qquad A \in \mathfrak{t}^{reg} \subset \mathfrak{gl}_n$$

where r is some odd number. Following the discussion in the last example, we get n solutions $f_i \sim \exp(\lambda_i z^{-r})$ The strands i and j will cross whenever f_i and f_j "switch" growth asymptotics. Suppose for instance that $\lambda_i, \lambda_j \in \mathbb{R}$. Writing $z = Re^{i\theta}$ the asymptotics will switch whenever $Re(z^{-r/2}) = 0$, i.e. on the rays

$$\theta = \frac{\pi}{r} + \frac{2n\pi}{r}$$

There are r such rays between any pair i, j even if $\lambda_i, \lambda_j \notin \mathbb{R}$: the expression for the rays is more complicated by the number of rays doesn't change. Therefore Λ is the closure of a (n, r) braid.

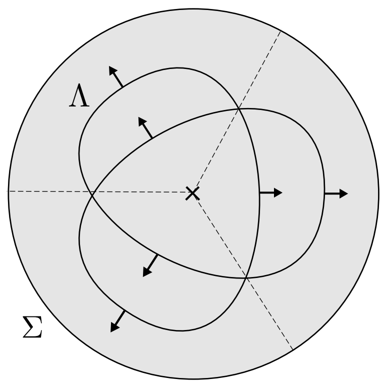

Figure 3.9: The Legendrian link for the irregular type $1/z^{3/2}$. The microlocal rank on the link Λ is 1.

The case where $n = 2, r = 3$ gets us a trefoil and appears in Stokes' discussion of the Airy equation [106]. Note that the trefoil has only one component, so the formal monodromy map lands in $Loc_{k^*}(S^1) \cong [\frac{k^*}{k^*}]$ which doesn't stand for $\frac{T}{T}$ for any maximal torus $T \subset GL_2$. This explains why this more general case cannot be described just in terms of moment maps into subgroups of G.

Remark. One can play many variations on the theme of Stokes filtrations and irregular singularities. Considering connections with matrices A that are not in the regular locus of $\mathfrak{g}$, it becomes necessary to look at further less singular terms in the expression for dQ. To find the corresponding Legendrians one can still follow the prescriptions in [102]. One obtains cablings of torus knots by other torus knots and cablings of torus knots by such cablings and so on.

But there are many other knots that one can consider: pick any positive braid and close it around the origin into a Legendrian Λ. Picking rank conditions and monodromies around the components of Λ, this gives a symplectic space $\mathcal{M}_\Lambda$ that doesn't necessarily come from an irregular meromorphic singularity. We can expect these spaces to carry some of the same structures as the tame and wild character varieties; whether this is true remains a topic of future research.

Remark. For moduli of Stokes data for connections on higher dimensional complex varieties, the corresponding Riemann-Hilbert theorem recently been proven [33]. We expect that it can be reformulated into an analogous "microsupport in certain smooth Legendrians" version, from which we would be able to immediately deduce the existence of the shifted symplectic structure.

Positroid varieties, multiplicative Nakajima varieties, and other cluster structures

The combinatorics of cluster algebras arising from surfaces was originally organized around data given variously as a graph on a surface, a triangulation on a surface, etc [92, 43, 52, 44]. One presentation of this data is in terms of the so-called "alternating strand diagram", the manipulation of which by combinatorial topology [110] underlies various theorems of the cluster algebra.

The authors in [103] present the perspective that the alternating strand diagram should be viewed as a Legendrian knot, that triangulations of the surface give rise to Lagrangian fillings of it, and that all the corresponding cluster algebraic formulas are computing the Floer homology between such fillings. In particular, the corresponding cluster X-variety was identified as a moduli space of "rank one" objects in $Sh_\Lambda(\Sigma)$, where Λ is the Legendrian lift of the alternating strand diagram, and Σ is the base curve.

In [101], a slightly different perspective is taken: rather than work from Σ, Λ, we began with a Legendrian L — one could view it as one of the above-mentioned fillings of Λ — and attached Weinstein handles to its cotangent bundle along Legendrians which project to simple closed curves. This perspective is yet more general than the previous.

It includes as a special case the multiplicative Nakajima quiver varieties of [32, 114, 15]; this being the case where the attaching circles are contractible. Indeed, this case is very close to the presentation in [15]. In that reference, rather than locally arboreal singularities, they consider the spaces which are locally either a smooth surface or modeled on the Lagrangian singularity given by the union of the zero section and the conormal to point. However, this local model admits a noncharacteristic deformation to the union of the zero section and the positive conormal to a circle. The deformation is just given by the contact isotopy induced by the Reeb flow; the fact that it is noncharacteristic follows then immediately from [54].

Thus, the present work recovers all constructions of symplectic and Poisson structures on such spaces. It will be interesting to investigate how the deformation quantization formalism of [25] interacts with these notions.

The augmentation variety of knot contact homology

Consider a knot or link $K \subset S^3$. Naturally associated to this is the category of sheaves constructible with respect to the stratification $S^3 = K \cup S^3 \setminus K$. This study of this category led recently to a proof that the Legendrian isotopy type of the conormal torus to a knot determines the knot [99]. In an appropriate sense, it is equivalent to the a category of augmentations of knot contact homology [40].

The union of the conormal to the knot with the zero section is not arboreal, but as the above discussion of [15], this can be remedied by perturbing the conormal torus by the Reeb flow, resulting in a skeleton given by the union of the zero section and the positive conormal to the inward (or outward) co-oriented boundary of a tubular neighborhood of the knot.

In any case, we can study the space of objects in this category. It has a map to the category of local systems on T^2, which by the results here, becomes a 0-shifted Lagrangian morphism on moduli spaces. We note that the study of this moduli space was also suggested in [15].

To select a connected component (indeed, the connected component corresponding to what is usually called the augmentation variety), we can pick those sheaves whose microsupport on the conormal torus is rank one in degree zero, i.e., what are called *simple sheaves* in [67]. We restrict further to the open locus on which objects which have no global sections, the main point of which is to eliminant constant summands. Let us write $A_1(K)$ for this component.

Restriction to the microlocal boundary gives a map $A_1(K) \to Loc_{k^*}(T^2) = (k^*)^2$. There is an analogous map in knot contact homology, described in [5, 41]. Note that objects in $A_1(K)$ are easy to understand: they are a local system in the complement of K, which is extended by a codimension one subspace of meridian invariants along K; or possibly a nontrivial rank one local system supported on K.

From this point of view it is clear both why $A_1(K)$ contains the classical A-polynomial curve, and also what are the other components: the A polynomial curve has to do with SL_2 representations of the fundamental group; any such becomes, after rescaling by an eigenvalue of the meridian, a GL_2 representation with a meridian invariant subspace. Similarly it is clear what the other components of $A_1(K)$ are. Thus we have shown that the morphism $A_1(K) \to k^* \times k^*$ is (0-shifted) Lagrangian. Quantization of this morphism features prominently in the conjectures of [5].

Chapter 4

Stability and Fukaya categories

It has been understood since the work of Douglas and others [39, 38] that understanding the appearance of derived categories in mirror symmetry requires thinking about the problem of stability of branes in the relevant SCFTs. This investigation arises from attempting to reconcile the multitude of objects in the derived category with their description as physical boundary conditions.

One of the most evident problems is that the derived category seems to have too many objects. Moreover, this category has auto-equivalences given by nontrivial monodromies, for example around nontrivial loops in the moduli space of complex structures of the underlying variety. Under these autoequivalences, objects with a simple physical description such as (submanifolds with) vector bundles get sent to more exotic objects. Therefore, if the use of the derived category is consistent, one must correctly interpret these exotic boundary conditions as particular combinations of branes. Following these and other clues, Douglas suggested the notion of Π-stability on SCFTs which is partially inspired by slope stability (or 'μ-stability') of vector bundles.

Inspired by the notion of Π-stability in physics, Bridgeland [22] proposed a notion of stability for general triangulated categories, which can be seen as a refinement of the notion of t-structures. One of the features of Bridgeland stability is that the space of stability conditions itself can be given the structure of a complex manifold, with interesting actions on it coming from automorphisms of the category. Calculating this space for categories of geometric interest, such as derived categories of algebraic varieties, has been a very active line of research [79, 13, 90, 89, 58]. A recent advance in this direction of research is the proof [74] that a candidate family of stability conditions on the quintic threefold inspired by Douglas' initial proposal in fact does give a family of Bridgeland stability conditions.

On the other hand, a similar understanding on the other (A-side) of mirror symmetry is still lacking. It has been conjectured for a while that Bridgeland stability conditions on A-side categories associated to Calabi-Yau varieties should reflect something about the geometry of *special Lagrangians*. In the literature this is often referred to as the Thomas-Yau conjecture, after [109], though strictly speaking Thomas and Yau in that reference make no mention of Bridgeland stability conditions, and instead use different notions of stability.

In this chapter we will present the definition and basic properties of Bridgeland stability conditions, and mention some recent developments in understanding them on the A-side of mirror symmetry. In particular, we will focus on the work of Haiden, Katzarkov and Kontsevich [56], which constructs Bridgeland stability conditions on Fukaya categories of surfaces. The constructions in that paper establish a precise relation between Bridgeland stability and an analog of special Lagrangian geometry, given by *quadratic differentials*.

In preparation for the following chapter, where we define a relative notion of stability conditions, we prove here some general lemmas about stability in such Fukaya categories, following the author's own work in [107].

4.1 Bridgeland stability conditions and the Thomas-Yau conjecture

Bridgeland stability conditions

Let us introduce the definition of a Bridgeland stability condition. Let us fix a triangulated category $\mathcal{D}$ and a lattice (finitely generated abelian group over $\mathbb{Z}$) Λ, with a fixed homomorphism $K_0(\mathcal{D}) \to \Lambda$.

Before we define a Bridgeland stability condition, let us define the notion of a slicing.

Definition 21. [23] A slicing $\mathcal{P}$ on $\mathcal{D}$ is a collection of full additive subcategories $\mathcal{P}_\phi$ for $\phi \in \mathbb{R}$, such that

- $\mathcal{P}_\phi[1] = \mathcal{P}_{\phi+1}$
- If $X \in \mathcal{P}_\phi$ and $Y \in \mathcal{P}_\psi$ with $\phi > \psi$, then $\mathrm{Hom}_{\mathcal{D}}(X, Y) = 0$
- For every $0 \neq X \in \mathcal{D}$ there is a sequence of real numbers $\phi_1 > \cdots > \phi_n$

$$\begin{array}{ccccccc} 0 & \longrightarrow & X_1 & \to \dots \to & X_{n-1} & \longrightarrow & X_n = X \\ & \nwarrow \swarrow & & & & \nwarrow \swarrow & \\ & A_1 & & & & A_n & \end{array}$$

where $A_i \in \mathcal{P}_{\phi_i}$ and each triangle is distinguished.

We will refer to objects in $\mathcal{P}_\phi$ as *semistable objects of phase* ϕ. A slicing should be understood as a refinement of a t-structure; for any interval $I \subseteq \mathbb{R}$, let us denote $\mathcal{P}_I = \bigsqcup_{\phi \in I} \mathcal{P}_\phi$. Then one can prove that setting $\tau_{>0}\mathcal{D} = \mathcal{P}_{(0,+\infty)}$ gives a bounded t-structure whose heart is $\mathcal{A} = \mathcal{P}_{(0,1]}$.

Definition 22. [23] A Bridgeland stability condition on the category $\mathcal{D}$ is a pair $(Z, \mathcal{P})$, where $Z \in \mathrm{Hom}(\Lambda, \mathbb{C})$ and $\mathcal{P}$ is a slicing on $\mathcal{D}$ compatible with Z in the sense that for every $0 \neq X \in \mathcal{P}_\phi$, $Z(X) = m(X)e^{i\pi\phi}$ for some $m(X) > 0$.

Note that for simplicity we denoted $Z(X)$ for the value of Z at the image of $[X]$ under the map $K_0(\mathcal{D}) \to \Lambda$.

In order to state the important deformation result of Bridgeland which gives the space of stability conditions the structure of a complex manifold, we need to restrict attention to stability conditions which obey the following finiteness property, known in the literature as the *support property* [73, 12]:

Definition 23. A stability condition $\sigma = (Z, \mathcal{P})$ satisfies the support property if

$$\inf_{0 \neq X \text{semistable}} \frac{|Z(X)|}{\|[X]\|} = C > 0$$

, where $\|\cdot\|$ is a norm on $\Lambda \otimes \mathbb{R}$.

From now on, we will only consider stability conditions satisfying the support condition above; in [23] these are referred to as 'locally-finite' stability conditions. Let $\mathrm{Stab}(\mathcal{D})$ denote the set of such stability conditions.

Theorem 41. *[23] The set* $\mathrm{Stab}(\mathcal{D})$ *naturally has the structure of a complex manifold, and the map* $\mathrm{Stab}(\mathcal{D}) \to \mathrm{Hom}_{\mathbb{Z}}(\Lambda, \mathbb{C})$ *is a local homeomorphism of complex manifolds.*

Special Lagrangians and the Thomas-Yau conjecture

While a lot is known about Bridgeland stability conditions on categories of coherent sheaves, it has been difficult to make similar progress on A-side categories. However there are many expectations about what (some) Bridgeland stability conditions on the Fukaya category of a Calabi-Yau manifold should look like. From the physics of branes in the SCFT, it is evident that stability in the A-model should be roughly given by minimality condition, which for Lagrangian branes takes the form of the special Lagrangian condition.

Let Y be a Calabi-Yau m-fold with Calabi-Yau form $\Omega \in \Omega^{m,0}(Y)$. A *special Lagrangian* submanifold is an oriented Lagrangian submanifold L such that there is a real number $\xi(L)$, called the phase of L, such that we have

$$\Im(e^{-i\pi\xi(L)}\Omega|_L) = 0$$

ie. the phase of Ω is the constant $\xi(L)$ when restricted to L.

We denote by $Z(L) = \int_L \Omega$ the BPS central charge of the brane. Note that we have $\xi(L) = \arg(Z(L))/\pi$, and from this perspective the phase is only defined mod 2; as usual in Floer theory if we want to lift our gradings to $\mathbb{Z}$ we must have some vanishing conditions and make choices in grading the Lagrangian submanifolds; in the case we will present in the rest of this chapter these choices come with the grading of a surface.

Now we will describe a theorem due to Joyce [61] motivating the relation between special Lagrangian geometry and stability conditions, following the exhibition in [8, Sec.5.2.1].

Consider a family of CY m-folds Y_z for z in some small disk $D \in \mathbb{C}$ centered at zero, and suppose that Y_0 has two transversely intersecting special Lagrangians L_1, L_2 with equal phases. Then in the neighborhood of zero we have the following result about existence of special Lagrangians

Theorem 42. *Inside of Y_z, there is a special Lagrangian close to the connected sum of L_1 and L_2 if and only if $\xi(L_2) \leq \xi(L_1)$.*

This shows that the special Lagrangian geometry behaves with respect to deformations of complex moduli very much like the slicings of semistable objects in a Bridgeland stability condition under deformations of stability conditions; for example one has walls of real codimension one giving a wall-and-chamber structure. Based on this type of comparison, it has been conjectured by Thomas-Yau [109], Joyce [60], Bridgeland [22] and others that the choice of CY form determines a Bridgeland stability condition on the derived Fukaya category $D^\pi\mathcal{F}(Y)$. More specifically, in terms of stability data, we have the following conjecture:

Conjecture 1. Let (Y, Ω) be a CY m-fold as above. Then there is a natural Bridgeland stability condition $(Z, \mathcal{P})$ on $D^\pi\mathcal{F}(Y)$ such that the central charge Z is given by the composition

$$K_0(D^\pi\mathcal{F}(Y)) \to H_m(Y, \mathbb{Z}) \xrightarrow{\int \Omega} \mathbb{C}$$

of the natural map to homology with the integral of the CY form, and if L is a special Lagrangian with phase ϕ, then there is a semistable object in $\mathcal{P}_\phi$ supported on L

There are many possible enhancements/modifications of this conjecture; for a recent review of the technical aspects and progress on the subject the reader can consult eg. [60].

4.2 Fukaya categories of marked surfaces

Though the remainder of this chapter and the next one, we will discuss one context in which the relation between Bridgeland stability conditions and a version of special Lagrangian geometry is fully realized. This is the case of stability conditions on partially wrapped Fukaya categories of marked surfaces. It has been understood in recent years that these definitions fit in the framework of partially wrapped Fukaya categories of Liouville manifolds, which is an analog of the fully wrapped case that we discussed in 2.4.

In [56], Haiden, Katzarkov and Kontsevich (which we will refer from now on by the acronym HKK) construct Bridgeland stability conditions on these categories by using geometric data on the surface; specifically they use a quadratic differential with prescribed singularities. An equivalent description is that one starts from certain flat metrics on the surface. This perspective makes the connection to the Thomas-Yau conjecture mentioned above more explicit; the choice of flat metric should be seen as analogous to the choice of Calabi-Yau form.

Here we will first present the Fukaya category of a marked surface and then some lemmas appearing in the author's own work [107] that will be important for the main results of the next chapter. The exposition of the Fukaya category of a marked surface will mostly follow [56], which the reader can consult for a more complete account.

Marked surfaces

A graded marked surface (Σ, M, η) is a topological surface Σ with boundary $\partial\Sigma$, a set of *marked boundary intervals* $M \subset \partial\Sigma$ and a grading or line field η, ie. a section of the projectivized tangent bundle: $\eta \in \Gamma(\Sigma, \mathbb{P}T\Sigma)$. A graded surface has a shift automorphism which we will denote [1] given by the generator of $\pi_1(\mathbb{P}T_p\Sigma)$ at every point.

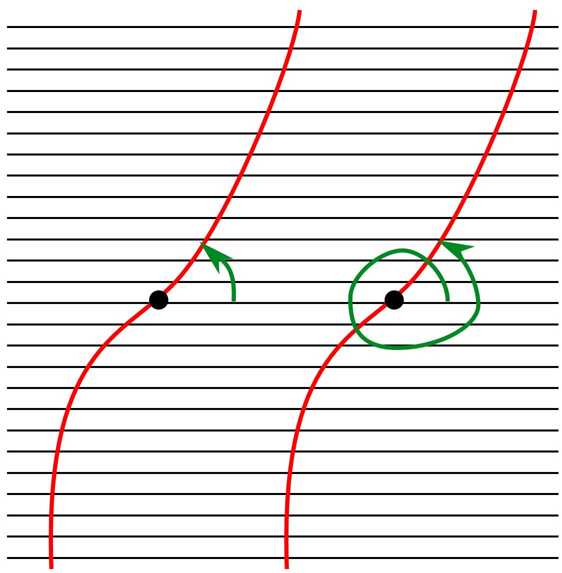

Figure 4.1: Schematic depiction of the data defining a graded curve. Here the line field of Σ is given by the parallel black lines, the red line is a curve in Σ and the green arrow depicts the paths $\tilde{c}$, which can be seen as choices of paths between the line field and the tangent to the curve. The two curves depicted here differ by a shift [2]

The choice of grading η allows us to give a grading to curves in Σ, which will be important for giving a $\mathbb{Z}$-grading on the Fukaya category. A graded curve is given by a triple $\gamma = (I, c, \tilde{c})$ where I is a 1-manifold, $c : I \to \Sigma$ is an immersion and $\tilde{c}$ is a choice of homotopy class of path from the tangent vector field of $c(I)$ to the line field. We will set the requirement that $c(\partial I) \subset M$, that is if I is an interval (as opposed to a circle) its ends must be mapped to the marked part of the boundary of $\partial\Sigma$.

This data allows us to associate an integer degree to each transverse intersection point between graded curves. Let p be a point of transverse intersection between immersed curves γ_1, γ_2. Then one can associate a degree of intersection $i_p(\gamma_1, \gamma_2)$ by comparing the homotopy classes of paths $\tilde{c}_1, \tilde{c}_2$. These degrees satisfy the following relation:

$$i_p(\gamma_1, \gamma_2) + i_p(\gamma_2, \gamma_1) = 1$$

and also behaves well under the shift automorphism, ie. $i_p(\gamma_1[m], \gamma_2[n]) = i_p(\gamma_1, \gamma_2) + m - n$.

The Fukaya category

While there are different possible strategies for defining the Fukaya category of a marked surface, here we will follow the conventions and definitions of [56], and present a geometricity result from that reference that relate that definition to other approaches.

Let (Σ, M) be a marked surface as above. An arc is an embedded interval with ends in M, not isotopic to an interval inside M itself. We call an arc a boundary arc if it is isotopic to a connected component of $\partial\Sigma \setminus M$. We call a collection of pairwise disjoint, non-isotopic arcs an *arc system*. An arc system $\mathcal{A}$ is full if $\mathcal{A}$ includes all the boundary arcs and $\Sigma \setminus \mathcal{A}$ is a disjoint union of polygons.

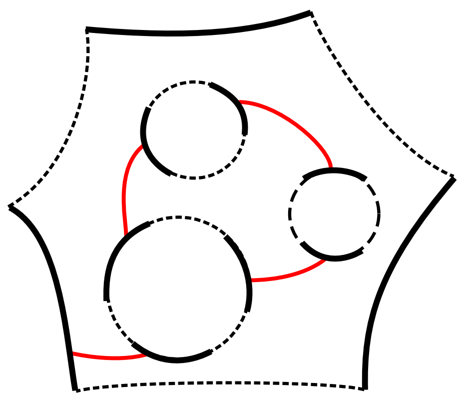

Figure 4.2: A marked surface with a system of arcs in red. The marked boundary intervals are denoted by solid black lines and the unmarked ones by dotted black lines. Note that if we were to count the red arcs and add arcs isotopic to the dotted black lines, we would have a full sysyem of arcs

Let us fix a full arc system $\mathcal{A}$ on Σ. Let us define a *boundary path* between arcs X and Y to be an oriented path in Σ, isotopic to an interval in M, starting at an endpoint of X and ending at an endpoint of Y. We require this path to follow the reverse orientation of $\partial\Sigma$, ie. keep the interior of Σ to its right hand side.

Then [56, Prop.3.1] there is a strictly unital A_∞-category $\mathcal{F}_{\mathcal{A}}(\Sigma)$ with:

- Set of objects $\mathcal{A}$.

- Morphisms between X, Y given by homotopy classes of boundary paths b from X to Y, with degree given by
$$|b| = i_p(X, b) - i_q(Y, b)$$
where p, q are the start and end of b.

- Composition given by concatenation of boundary paths when possible, and zero otherwise.

- Higher operations given by polygons; ie. if $b_1, \ldots, b_n$ is a sequence of paths bounding a polygon, then for any path b beginning at the end of b_1 we have
$$\mu^n(b_n, \ldots, b_1 b) = (-1)^{|b|} b$$
and for every path b ending at the beginning of a_n
$$\mu^n(b b_n, \ldots, b_1) = b$$

The A_∞ category $\mathcal{F}_{\mathcal{A}}(\Sigma)$ clearly depends on $\mathcal{A}$, in a strict sense, but one can show that its Morita equivalence class is independent of the choice of full arc system $\mathcal{A}$. Consider then the category of twisted complexes (in the sense of [19]) over $\mathcal{F}_{\mathcal{A}}(\Sigma)$; up to quasi-equivalence this triangulated category is independent of the choice of arc system. We will take this category $\mathcal{F}(\Sigma) := \mathrm{Tw}(\mathcal{F}_{\mathcal{A}}(\Sigma))$ to be our model for the topological Fukaya category of (Σ, M).

It will be important for our calculations to have explicit descriptions of the indecomposable objects of $\mathcal{F}(\Sigma)$; a priori, since we are taking twisted complexes, one would expect a generic object not to have a geometric interpretation. Fortunately we have the following result establishing the geometricity of objects in this category.

Theorem 43. *[56, Theorem 4.3] Every isomorphism class of indecomposable objects in $\mathcal{F}(\Sigma)$ can be represented by an admissible graded curve with indecomposable local system, unique up to graded isotopy.*

An admissible graded curve is either an immersed interval ending at marked intervals or an immersed circle, which does not bound a teardrop. An important role will be played by objects that can be represented by embedded curves. Let us from now on call an object an *(embedded) interval object* if it can be represented by an (embedded) interval, and a *(embedded) circle object* if it can be represented by an (embedded) circle. Note that every local system on an embedded interval is trivial so an indecomposable embedded interval object necessarily has a rank one local system.

Another result of [56] is a description of $K_0(\mathcal{F}(\Sigma))$ for surfaces Σ without unmarked boundary circles (which is the case that we are considering here). The grading on Σ gives a double cover τ by the orientation of the foliation lines; consider the local system of abelian groups $\mathbb{Z}_\tau = \mathbb{Z} \otimes_{\mathbb{Z}/2} \tau$.

Theorem 44. *[56, Theorem 5.1] There is a natural isomorphism of abelian groups $K_0(\mathcal{F}(\Sigma)) \cong H_1(\Sigma, \mathcal{M}; \mathbb{Z}_\tau)$.*

4.3 Lemmas about stability conditions

In this section we collect some lemmas about stability conditions in general, and also about the specific case where $\mathcal{D} = \mathcal{F}(\Sigma)$ is the Fukaya category of a marked surface Σ. Throughout the rest of this document we will only deal with the "fully stopped" case, ie. the case where each boundary circle in $\partial\Sigma$ has at least one marked interval.

Stability conditions and genericity

We will make use of genericity assumptions, which will play an important role in later proofs. To express genericity we need to define walls in this space, following [24]. Let us fix a class $\gamma \in \Lambda$, and consider other classes α such that α and γ are not both multiples of the same class in Λ.

Definition 24. The wall $W_\gamma(\alpha) \subset \mathrm{Stab}(\mathcal{D})$ is the subset of stability conditions such that there is a phase $\phi \in \mathbb{R}$ and objects A, G with respective classes α, γ such that $A \subset G$ in the abelian category $\mathcal{P}_\phi$.

Each wall $W_\gamma(\alpha)$ is contained within a codimension one subset of $\mathrm{Stab}(\mathcal{D})$ where $Z(\alpha)/Z(\gamma)$ is real, and we have the following local finiteness result:

Lemma 45. *[24, Lemma 7.7] If $B \subset \mathrm{Stab}(\mathcal{D})$ is compact then for a fixed γ only finitely many walls $W_\gamma(\alpha)$ intersect B.*

Note that this is not true if we consider the whole collection of walls for all γ; the union of all walls can be dense in $\mathrm{Stab}(\mathcal{D})$. So we will have to be specific when discussing genericity.

Definition 25. Let $\Xi \subset \Lambda$ be a finite subset of classes. Take

$$W_\Xi = \bigcup_{\gamma,\alpha\in\Lambda} W_\gamma(\alpha)$$

ie. the union of all closures walls for classes in Λ; we will say a stability condition σ is Ξ-generic if $\sigma \in \mathrm{Stab}(\mathcal{D}) \setminus \bar{W}_\Xi$.

By local finiteness, $\bar{W}_\Xi$ is a locally-finite union of closed subsets so Ξ-genericity is an open condition. The connected components of $\mathrm{Stab}(\mathcal{D}) \setminus \bar{W}_\Xi$ will be called the Ξ-chambers.

We will later make use of the following simple fact, which holds for any stability condition, generic or not.

Lemma 46. *If $X = E \oplus F$ then $\mathrm{HNLen}(X)$ is equal to the total number of distinct phases appearing among the HN decomposition of E and F. In particular, $\max(\mathrm{HNLen}(E), \mathrm{HNLen}(F)) \leq \mathrm{HNLen}(X) \leq \mathrm{HNLen}(E) + \mathrm{HNLen}(F)$.*

Proof. Follows from uniqueness of the HN decomposition, and the fact that given a HN decomposition of E and F one can algorithmically produce an HN decomposition of $E \oplus F$. □

We will also need prove the following proposition, which constrains the type of objects that can be stable under some stability condition. This will play an important role in our later theorems.

Proposition 47. *For any stability condition* $\sigma \in \mathrm{Stab}(\mathcal{F}(\Sigma))$*, every stable object is either an embedded interval object or an embedded circle object.*

Proof. Since L is indecomposable its support cannot have more than one connected component. Thus the only objects we have to rule out are objects whose representatives all have self-intersections; we will call these *truly immersed objects*.

A stable object L must have $\mathrm{Ext}^i(L, L) = 0$ for $i < 0$. Let L be a truly immersed objects and pick a representative of L with minimal number of self-intersections, supported on an immersed curve γ_L. Perturbing L to calculate endomorphisms, we see that a self-intersection point p of γ_L contributes classes to $\mathrm{Ext}^*(L, L)$ in degrees i_p and $1 - i_p$, where i_p is the degree of intersection at p. These classes are nonzero by minimality of self-intersections, so if there is a self-intersection point with $i_p \neq 0, 1$, one of these degrees is negative and therefore L cannot be semistable.

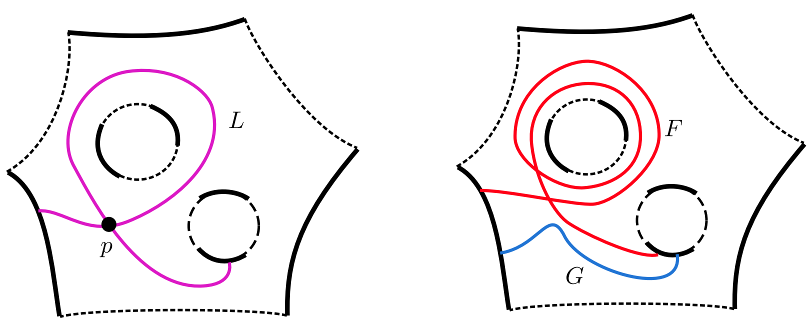

Figure 4.3: A truly immersed Lagrangian L. The self-extension $L \to E \to L$ at the self-intersection point p splits as a direct sum $E = F \oplus G$.

The only case left to consider is when γ_L only has self-intersection points of degree 0 and 1; each one of these points gives nonzero classes in $\mathrm{Hom}(L, L)$ and $\mathrm{Ext}^1(L, L)$. Let us pick one of these points p, and consider the corresponding nontrivial extension $L \to E \to L$. Note that the support of E is given by two superimposed curves so we have a direct sum decomposition $E = F \oplus G$. But by assumption L is stable of phase ϕ_L, so E, F and G are also all semistable of the same phase. Consider now the abelian category $\mathcal{P}_{\phi_L}$ of semistable objects of that phase. Since the stability condition is locally finite, this category is finite length; therefore the Jordan-Hölder theorem applies [59]. Since the length of E is 2, F and

G are length one, and by uniqueness of the simple objects in the Jordan-Hölder filtration (up to permutations) we must have $F \cong G \cong L$. But this is impossible because E is a nontrivial extension so $E \neq L \oplus L$. □

Remark. Note that the proof above does not preclude a self-intersecting object L from being *semistable*; it just cannot be simple in $\mathcal{P}_{\phi_L}$. In fact this even happens generically: take Σ to be the annulus with one marked interval on each boundary circle and grading such that the nontrivial embedded circle is gradable; by mirror symmetry the category $\mathcal{F}(\Sigma)$ is equivalent to $D^b(\mathrm{Coh}(\mathbb{P}^1))$. Under this equivalence, the rank one circle object with monodromy $z \in \mathbb{C}^\times$ gets mapped to the skyscraper sheaf $\mathbb{C}_z$ on $\mathbb{P}^1$, and the interval object I with both ends on the outer boundary, wrapping the annulus once, gets mapped to the skyscraper sheaf $\mathbb{C}_\infty$ on $\mathbb{P}^1$.

The space of stability conditions on this category is known to be isomorphic to $\mathbb{C}^2$ as a complex manifold [90], and there is a geometric (top dimensional) chamber in $\mathrm{Stab}(\mathbb{P}^1)$ where all the rank one skyscraper sheaves are stable. In particular, the nontrivial extension $I \to L \to I$, represented by an immersed Lagrangian with one self-intersection as in Figure 4.4, is semistable. So self-intersecting objects do appear generically, but they always have Jordan-Hölder decompositions into embedded objects.

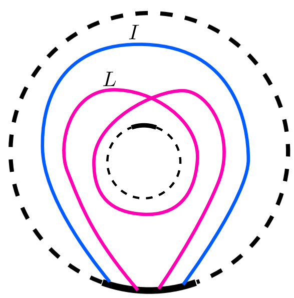

Figure 4.4: The annulus mirror to $D^b(\mathrm{Coh}(\mathbb{P}^1))$. For a geometric stability condition on $\mathbb{P}^1$, the truly immersed object L (corresponding to an irreducible rank 2 skyscraper sheaf $\mathcal{O}_{x^2}$) is semistable.

The result above characterizes which objects can be stable, namely embedded intervals and embedded circles with indecomposable local systems. It turns out that similar index computations also allows us to constrain the form of the HN decompositions of objects.

Definition 26. (Chain of stable intervals) Let us fix a stability condition $\sigma \in \mathrm{Stab}(\mathcal{F}(\Sigma))$ and consider an indecomposable object X in $\mathcal{F}(\Sigma)$. We say that X has a chain of stable intervals decomposition (COSI decomposition) under σ if there is

- A sequence of stable (therefore embedded) interval objects $X_1, \ldots, X_N$ and a sequence of marked boundary intervals $M_0, \ldots, M_N$, where the support γ_i of the object X_i has ends on M_{i-1} and M_i,

- Extension morphisms $\eta_i \in \mathrm{Ext}^1(X_i, X_{i+1})$ or $\eta_i \in \mathrm{Ext}^1(X_{i+1}, X_i)$ corresponding to the shared M_i marked boundary (including an extension at $M_0 = M_N$ if X is a circle object),

such that the iterated extension by all the η_i is isomorphic to X.

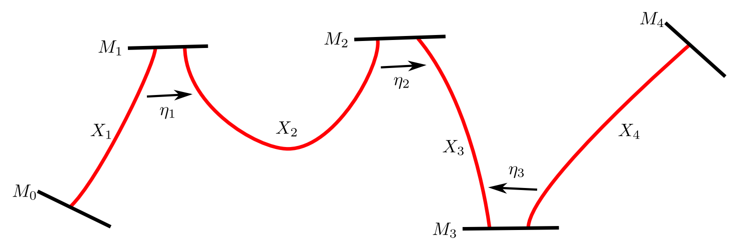

Figure 4.5: A chain of stable intervals with N=4.

Remark. Note that the order $X_1, \dots, X_N$ here is not directly related to the ordering of semistable objects in the HN decomposition of X; in particular there is no constraint on the phases of the X_i, and the extension maps can go either way.

Note that if X has a COSI decomposition then its HN decomposition can be produced from it by grouping together all stable interval objects of the same phase.

Lemma 48. *If X has a COSI decomposition under σ, then it is essentially unique, ie. the sets $\{X_i\}$ and $\{M_i\}$ are uniquely defined up to isomorphism.*

Proof. Follows from the uniqueness of the HN filtration and the uniqueness (up to permutation) of the Jordan-Hölder filtration on each finite-length abelian category $\mathcal{P}_\phi$. □

This decomposition also captures the isotopy class of the object X. Let us produce an immersed curve γ from this data as follows: for each i, if the extension map η_i belongs to $\mathrm{Ext}^1(X_i, X_{i+1})$ we connect γ_i to γ_{i+1} counterclockwise (ie. by a boundary path following M_i and keeping Σ to the right), and if $\eta_i \in \mathrm{Ext}^1(X_{i+1}, X_i)$ we use the corresponding clockwise path from γ_i to γ_{i+1}. From the geometricity result in Theorem 43 we can deduce that:

Lemma 49. *The curve γ is isotopic to the support γ_X of the object X.*

The following lemma will be central to our proofs later, and essentially means that COSI decompositions are not allowed to cross each other. From now on, we will leave the extension morphisms implicit and denote a COSI decomposition by its stable intervals.

Lemma 50. *Let X and Y be two objects with respective COSI decompositions $(X_1, \dots, X_m)$ and $(Y_1, \dots, Y_n)$. We choose representatives for all the stable intervals such that the number of crossings between these two chains of intervals is minimal. Then on the surface Σ there are none of the following polygons*

1. *Polygons bounded by the two chains and two transversal crossings between stable intervals.*
2. *Polygons bounded by the two chains and two common marked boundary intervals (with boundary paths inside the polygon).*
3. *Polygons bounded by the two chains, one transversal crossing and one common marked boundary interval (with a boundary path inside the polygon).*

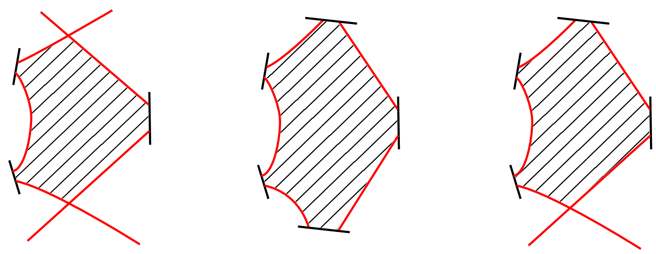

Figure 4.6: The three kinds of polygons of stable intervals that cannot appear by Lemma 50. Here we have polygons with $k = 3$ sides on the left and $l = 2$ sides on the right. The shaded interior means that these polygons bound disks inside of Σ.

Remark. In case (2), we exclude the trivial bigon with isomorphic sides. This case is obviously allowed, and happens whenever X and Y share a same interval in their COSI decompositions. From all cases, we exclude the degenerate configuration where all the objects around the polygon are multiples of the same class in $K_0(\mathcal{F}(\Sigma))$. For cases (2) and (3), the parenthetical condition is there because the chains could meet at some marked boundary interval 'on the other side of the polygon'. For instance again in the case of the annulus 4.3, for an 'algebraic' stability condition on $\mathbb{P}^1$ the two intervals (corresponding to line bundles on $\mathbb{P}^1$) are stable and we can have the following bigon of stable objects; but the boundary path giving the extension runs outside the polygon.

Proof. Let us first prove that it is sufficient to prove the statement for adequately generic σ. By standard arguments, the locus of $\mathrm{Stab}(\mathcal{D})$ in which the all the objects X_i, Y_i are stable is open. Consider now the collection $\Xi \subset \Lambda$ containing all the classes of these objects; the corresponding union of walls $\bar{W}_\Xi$ is a locally-finite union of closed subsets of positive codimension. So we can find some other stability condition σ', arbitrarily close to σ, where X_i, Y_i still give COSI decompositions of X, Y, and where the phases of any X_i and Y_j are pairwise distinct when $[X_i]$ and $[Y_j]$ are not proportional. If the noncrossing statement of the lemma is true for σ' it is also true for σ since it makes no further reference to the stability condition.

Let us start with the first type of polygon. Assume the polygon has k edges on the right and l edges on the left, and for ease of notation we label the intervals in this polygon

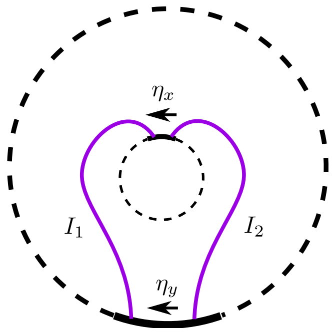

Figure 4.7: Again, the annulus $\Delta^*_{(1,1)}$ mirror to $\mathbb{P}^1$. Under this correspondence $\mathcal{F}(\Delta^*_{(1,1)}) \cong D^b(\mathrm{Coh}(\mathbb{P}^1))$ we have $I_1 \cong \mathcal{O}(1)$ and $I_2 \cong \mathcal{O}$ with $\mathrm{Hom}(I_2, I_1)$ spanned by η_x, η_y. Note that is not a counterexample to case (2) of Lemma 50 since the polygon doesn't bound a disk.

starting by 1 on both sides. Without loss of generality shift the grading of X such that the intersection point p has index $i_p(X_1, Y_1) = 1$. By minimality of crossings p contributes nonzero classes in $\mathrm{Ext}^1(X_1, Y_1)$ and in $\mathrm{Hom}(Y_1, X_1)$. Since both are stable objects, this implies that

$$\mathrm{phase}(Y_1) \leq \mathrm{phase}(X_1) \leq \mathrm{phase}(Y_1) + 1.$$

Smoothing out each one of the chains of intervals separately, one gets a bigon with vertices at p and q; the existence of the embedded bigon constrains the index of q to be $i_q(X_k, Y_l) = 0$, and by the same argument we have

$$\mathrm{phase}(X_k) \leq \mathrm{phase}(Y_l) \leq \mathrm{phase}(X_k) + 1.$$

By assumption, all the other vertices of this polygon give, on the left hand side, extension maps $X_i \xrightarrow{+1} X_{i+1}$, and on the right hand side, extension maps $Y_{i+1} \xrightarrow{+1} Y_i$. Since all these maps appear in HN decompositions we must have the following inequality between phases

$$\mathrm{phase}(X_i) \leq \mathrm{phase}(X_{i+1} \text{ for all } 1 \leq i \leq k-1, \qquad \mathrm{phase}(Y_j) \geq \mathrm{phase}(Y_{j+1}) \text{ for all } 1 \leq j \leq l-1$$

, which together with the previous inequality gives that the phases are all equal. But since we excluded the degenerate polygons, at least two of the K_0 classes of this object these objects are not multiples of the same class so by Ξ-genericity of σ' they have distinct phases. The three other cases are proven by small variations of this same argument. □

Remark. Note that the two chains might still share a common stable interval; this is not ruled out by the argument above and in fact happens generically. Similarly, note that our definition of chain-of-intervals decomposition above does not exclude the possibility that the chain of intervals overlaps with itself. Again, in the annulus example consider some algebraic stability condition such that the stable objects are two intervals I_1, I_2 connecting the outer and inner boundary, and consider the embedded interval object also connecting the two boundaries but wrapping around more times; this object has a COSI decomposition given by multiple copies of I_1 and I_2.

Self-overlapping chains of intervals will pose some serious technical difficulties later on, so we will rule them out with the following criterion. Let X be an indecomposable object with a COSI decomposition $(X_1, \dots, X_N)$, with X_i supported on γ_i.

Definition 27. This is a *simple* COSI decomposition if all the γ_i are in pairwise distinct isotopy classes, all the marked boundary intervals $M_1, \dots, M_N$ are pairwise distinct and also distinct from the ends M_0, M_{N+1} of X.

This condition implies that among the stable objects X_i, one does not find more than one copy of any given isomorphism class, or any of its shifts more than once. Moreover, only successive intervals share marked boundary components, so among these objects the only nontrivial degree zero homs are the self-homs and the only non-trivial extension homs are between adjacent intervals.

Lemma 51. *If X has a simple* COSI *decomposition as above, then its HN envelope* $\mathrm{HNEnv}(X)$ *is equivalent to either:*

- *The Fukaya category of the disk Δ_{N+1} with $N+1$ marked boundary intervals, or equivalently the derived category of the A_N Dynkin quiver, if X is an interval object with ends on distinct marked boundary intervals, or*
- *The Fukaya category of the annulus $\Delta^*_{p,q}$ with p and q inner and outer boundary intervals for some $p+q = N+1$ and grading of index zero around the circle, or equivalently the derived category of the $\tilde{A}_N$ quiver, if X is a circle object.*

Proof. We can prove this constructively by giving a map of arc systems. Consider the (non-full) arc system given by all the intervals γ_i; this defines an A_∞-category $\mathcal{A}$. Since this is a chain of arcs there are no polygons so all the higher structure maps μ^i between them are trivial. Note that $\mathrm{HNEnv}(X)$ is obtained by taking the triangulated closure of $\mathcal{A}$.

If X is an interval object, let us denote by m the number of indices i such that the extension map at M_i it 'on the left' ie. given by an extension map in $\mathrm{Ext}^1(X_{i+1}, X_i)$. Similarly we denote by n the number of extensions 'on the right' ie. given by an extension map in $\mathrm{Ext}^1(X_i, X_{i+1})$; we have $m + n = N - 1$. Consider the disk Δ_{N+1} with the following arc system: position m of the marked boundary intervals on the left and n on the right, with the remaining two on the top and bottom. There is then a unique chain of arcs α_i starting from the bottom and ending at the top such that α_i and α_{i+1} meet on the left if the extension is in $\mathrm{Ext}^1(X_i, X_{i+1})$ and on the right if the extension is in $\mathrm{Ext}^1(X_{i+1}, X_i)$.

This arc system gives an A_∞-category equivalent to $\mathcal{A}$, since the morphisms all agree and all the higher structure maps are zero. The argument for the circle case is similar, except we put m of the marked boundary components on the inner boundary circle and n on the outside (considering also the extension given by $M_0 = M_N$) □

In general, objects will not have a simple COSI decomposition, but the following topological condition is sufficient.

Lemma 52. *Let X be an object with a* COSI *decomposition, supported on an embedded interval γ separating the surface Σ into two connected components, such that the two ends of γ belong to distinct marked boundary intervals. Then X has a simple* COSI *decomposition.*

Proof. Let us write as before $\gamma_1, \ldots, \gamma_N$ for the intervals and $M_1, \ldots, M_{N-1}$ for the marked boundary intervals between them. We would like to rule out the possibility of having repeated intervals or marked boundary intervals.

Suppose that the subsequence

$$M_i, \gamma_{i+1}, M_{i+1}, \ldots, M_{i+k-1}, \gamma_{i+k}, M_{i+k}$$

repeats itself, ie. all those intervals and marked boundary components are isomorphic to

$$M_j, \gamma_{j+1}, M_{j+1}, \ldots, M_{j+k-1}, \gamma_{j+k}, M_{j+k}$$

for some other j. For simplicity assume that $j > i + k$ so there's no overlap; and let us assume that k is maximal. Let us also assume that $i > 0$ and $j + k < N$ so that we are in the middle of the chain and not at the ends, and that j is the smallest index possible with these properties (because this sequence could in principle repeat many times).

There are then four possibilities for the extension maps at M_i and M_{i+k}, as below:

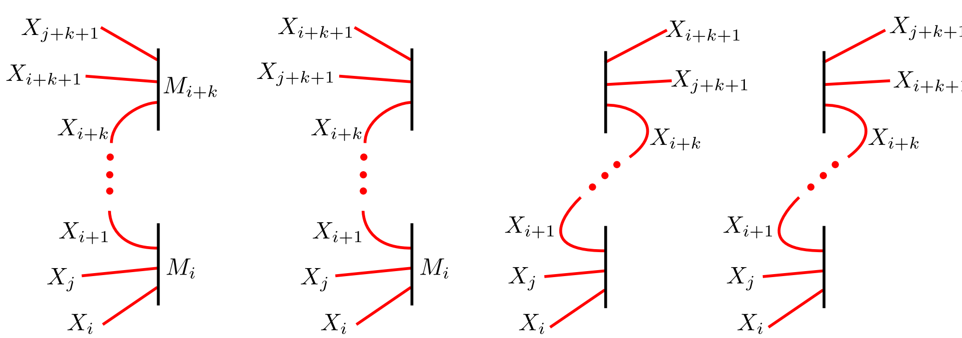

Figure 4.8: Four possible cases for extensions within a self-overlapping chain.

If we are in the first case or third case, note that concatenating the chain by those boundary walks leads to a self-crossing of γ_X. This self-crossing cannot be eliminated by isotopy, because due to Lemma 50 there are no polygons of stable intervals bound by the chain. Since we assumed that X is an embedded interval object this is impossible.

As for the second case and fourth case, note that concatenating the chain by those boundary walks leads to an embedded interval that does not separate the surface into two parts, contradicting the topological condition.

The special cases to be dealt with are when this repeated sequence is at one end of the chain; in this case it is easy to see that the concatenation is always non-trivially self-intersecting, unless the overlap is just a single boundary component $M_0 = M_N$ which we also excluded by assumption. The more general case of repeated intersections, nested intersections etc. poses no essential difficulties and can be argued by repeating the argument above recursively. □

With these lemmas, we prove the following proposition constraining the form of the HN decomposition of an object.

Proposition 53. *Let X be an rank one indecomposable object of $\mathcal{D} = \mathcal{F}(\Sigma)$ and $\sigma \in \mathrm{Stab}(\mathcal{D})$ any stability condition. Then X is either a stable circle or has a chain of stable intervals decomposition under σ.*

Proof. Suppose that X is not a semistable circle. Consider the HN decomposition of X under σ and further decompose each semistable factor of phase ϕ using the Jordan-Hölder filtration on the abelian category $\mathcal{P}_\phi$. We get then a total filtration

$$\begin{array}{ccccccccc} 0 & \longrightarrow & X_1 & \longrightarrow & X_2 \to \dots \to & X_{N-1} & \longrightarrow & X_N = X \\ {\scriptstyle \epsilon_1=0} \nwarrow & & \swarrow {\scriptstyle \pi_1} \; {\scriptstyle \epsilon_2} \nwarrow & & \swarrow {\scriptstyle \pi_2} & {\scriptstyle \epsilon_n} \nwarrow & & \swarrow {\scriptstyle \pi_n} \\ & A_1 & & A_2 & & & A_N & \end{array}$$

where each factor A_i is stable but the phases ϕ_i might repeat.

We will prove by induction on the total length N. The case $N = 1$ is obvious. Assume now that the statement is true for any object of total length $N - 1$, and take an object X as above.

Consider the extension $X_{N-1} \to X_N \to A_N$. Since the object A_N is stable, by Lemma 47 it is either representable either by an embedded interval or an embedded circle. We will treat these cases separately.

If A_N is an interval object supported on a embedded interval α_N, and X_{N-1} is supported on some collection of immersed curves γ_{N-1}. Note that we can also express X_{N-1} as an extension

$$A_N[-1] \to X_{N-1} \to X_N$$

, so we conclude that X_{N-1} is either supported on a single immersed curve (interval or circle) or a direct sum of two intervals.

We choose α_N and γ_{N-1} to have minimal intersections with each other. The extension map $\eta \in \mathrm{Ext}^1(A_N, X_{N-1})$ comes from a linear combination of classes corresponding to transverse intersection points between α_N and γ_{N-1}, and shared marked boundary intervals; let us write

$$\eta = c_1 M_1 + c_2 M_2 + \sum_p c_p p$$

where M_1, M_2 are extension maps given by the marked boundary intervals at the end of A_N and p labels extension maps coming from intersection points. Note that the coefficients c_1, c_2, c_p are not uniquely defined.

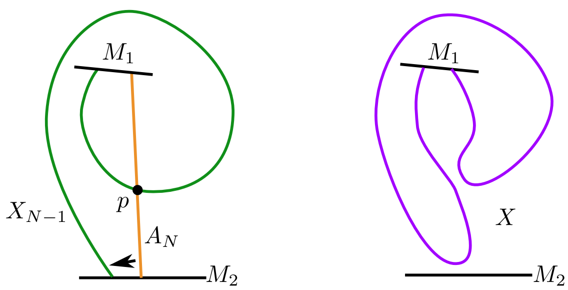

Figure 4.9: One example where A_N extends X_{N-1} with an extension map $c_2 M_2 + c_p p$. Using only the extension at p we obtain X' which is the sum of two interval objects (each of smaller total length), which can be extended at M_2 to give X. In this case X_{N-1} and A_N shared the other boundary too; this does not have to be the case in general

We see that it is impossible to have $c_1 = c_2 = 0$. If the extension happens only at transverse intersection points, then this extension is supported on two (or more) superimposed curves which is impossible since we assumed $X_N = X$ was indecomposable.

Consider then the modified extension map

$$\eta' = \sum_p c_p p$$

and the corresponding extension $X_{N-1} \to X' \to A_N$. This is supported on a set of curves that share the marked boundary intervals M_1 and/or M_2 and moreover can be extended at those to obtain the original object X. This topologically constrains X' to be of one of three types:

1. $X' = I_1 \oplus I_2$, two intervals which can be extended at a common boundary to form the interval object X,

2. $X' = I_1 \oplus I_2 \oplus I_3$, three intervals which can be extended at two common boundaries to form the interval object X,

3. $X' = I_1 \oplus I_2$, two intervals which can be extended at both common boundaries to form a circle object X.

Whichever case we are in, since total length is additive, the indecomposable factors I_1, I_2, I_3 are all of length $\leq N - 1$ so by the induction hypothesis they have COSI decompositions, which can then be composed at the shared marked boundaries to give a COSI decomposition for X.

It remains to deal with the case where A_N is a circle object. Since there is no boundary, the extension map $\eta \in \mathrm{Ext}^1(A_N, X_{N-1})$ must be given by a linear combination

$$\eta = \sum_p c_p p$$

of the classes given by transverse intersections p between α_N and γ_{N-1}. Assume first that $N \geq 3$; then $N-1 \geq 2$ and therefore X_{N-1} is not a semistable circle so by the induction hypothesis it has a COSI decomposition coming from concatenating intervals $\alpha_1, \ldots, \alpha_{N-1}$.

We see that every transverse intersection of index 1 between α_N and γ_{N-1} must come from one or more transverse intersections of index 1 between α_N and another α_i. However this gives a nonzero class in $\mathrm{Hom}(A_i, A_N)$ which cannot happen if $\phi_{A_i} \geq \phi_{A_N}$, so the only possibility is that these have the same phase (ie. appear together in the HN filtration). But this is also impossible: since A_i and A_N are both simple objects in the abelian category $\mathcal{P}_{\phi_{A_i}}$, the existence of this nonzero morphism implies that $A_i \cong A_N$ which cannot happen since one is a circle object and another is an interval object.

The only last case to deal with is when $N = 2$ and X is an extension of two stable circle objects A_1, A_2; by the same argument as above this can only happen if the two circles are isomorphic but then X cannot be rank one. □

One easy consequence of this result is that the monodromy of the local system carried by the immersed curve does not matter for its stability.

Corollary 54. *Fix any stability condition σ as above, and X any rank one object supported on a curve γ. If X is stable under σ, then any other rank one object X' supported on γ is also stable under σ.*

Proof. Suppose otherwise; then X' has a COSI decomposition. But the same chain of intervals can be concatenated to give X as well, by taking different multiples of the extension classes between the intervals in the chain, contradicting the assumption. □

The only indecomposable objects not covered by Theorem 53 are circle objects with higher rank local systems, but this will cause no further problems:

Lemma 55. *Let X be an indecomposable object supported on a circle γ with higher-rank local system. Then there are two possibilities for X:*

1. *X is a semistable interval whose stable components are all rank one objects supported on γ,*

2. *X has a decomposition as as chain of semistable intervals, ie. similar to a COSI decomposition except that every piece is a direct sum of stable intervals instead of a single stable interval.*

Proof. Suppose X carries a rank r indecomposable local system $\mathcal{L}$. If the rank one objects supported on γ are stable, then we pick r such objects with monodromies given by the eigenvalues of $\mathcal{L}$; using the self-extension of the circle we can present X as an iterated extension of these objects, proving that X is semistable, so we are in case (1). Otherwise, these rank one objects have a COSI decomposition; again we take r copies of this chain of stable intervals and extend them appropriately to construct the local system $\mathcal{L}$, and we are in case (2). □

Combining the results above, we conclude that certain kinds of embedded intervals always have simple COSI decompositions.

Corollary 56. *Let X be an object of $\mathcal{F}(\Sigma)$ represented by an embedded interval γ_X with trivial rank one local system, such that γ_X cuts the surface into two, and has ends on distinct marked boundary intervals. Then X has a simple* COSI *decomposition under any stability condition, and thus there is an abstract equivalence of triangulated categories* $\mathrm{HNEnv}(X) \cong D^b(A_N)$.

4.4 The HKK construction

In this section we will present a construction by Haiden, Katzarkov and Kontsevich [56] of stability conditions on the Fukaya categories we defined in Section 4.2. This construction gives stability conditions starting from some geometric data on the surface; one starts with a quadratic differential on a compactification of Σ with singularities of prescribed type.

To define such quadratic differentials, one must pick a complex structure on this compactification; this is an indication that the choice of quadratic differential here plays an analogous role as the choice of Calabi-Yau form in the Thomas-Yau conjecture. In Chapter 5 we will introduce some formalism that will allow us to prove in some cases that every Bridgeland stability condition in fact comes from such geometric data.

Flat surfaces

Let us first define the structure that will be used to define such stability conditions.

Definition 28. A *flat surface* is a (possibly non-compact) topological surface X with a choice of complex structure and a section of the square of the holomorphic tangent bundle of X.

Equivalently, this structure can also be presented as a flat Riemannian metric on an oriented surface, together with a covariantly constant foliation η (the *horizontal foliation*). We consider only surfaces whose metric completions have conical singularities; these can be finite-angle singularities, where a neighborhood of the singularity is homeomorphic to finitely many half-planes glued together, or infinite-angle singularities, where the neighborhood is homeomorphic to an union of half-planes indexed by $\mathbb{Z}$.

Let us fix now a surface with boundary $(\Sigma, \partial\Sigma)$. We will define the stucture of a real blowup of a flat surface with underlying surface Σ to be the structure of a flat surface of $\Sigma \setminus \partial\Sigma$, such that every neighborhood of a component of $\partial\Sigma$ is homeomorphic to the real blow-up of a conical singularity. Contracting $\partial\Sigma$ one gets a flat surface X with conical singularities at a collection of points P.

If moreover Σ comes with a grading η (ie. a line field) one can ask that there be a homeomorphism of graded surfaces between $\Sigma \setminus \partial\Sigma$ and $X \setminus P$; we will denote by $\mathcal{M}(\Sigma)$ the moduli space of such flat surfaces X.

Fix now such a flat surface X. A saddle connection on X is a geodesic whose ends converge to points in P; the phase of such a saddle connection will be the angle it makes with the horizontal foliation η divided by π; this is a real number defined mod 2. Consider now the collection of saddle connections that are horizontal, ie. have phase zero. In a generic situation, cutting along the horizontal saddle connections gives a nice decomposition of the flat surface. More specifically, we have the following result:

Proposition 57. *[56, Prop.2.4] Let X be a connected marked surface with finitely many boundary components and with infinite area. Then, possibly after a arbitrarily small rotation of the flat structure, the horizontal saddle connections divide X into finitely many horizontal strips of finite height and (possibly infinitely many) horizontal strips of infinite height.*

A saddle connection that is entirely contained inside a single horizontal strip of finite height will be called a simple saddle connection.

HKK stability conditions

Consider now a real blowup $(\Sigma, \partial\Sigma)$ of a flat surface X with quadratic differential φ. The horizontal foliation of X gives rise to a grading η of Σ. Moreover, for every finite-angle singularity of φ we mark the whole corresponding component of $\partial\Sigma$, and for every infinite-angle cone singularity we put a marked interval in the corresponding component of $\partial\Sigma$.

For example, following the classification of (possibly exponential-type) singularities of quadratic differential in [56, Sec.2], this is equivalent to assigning the following markings:

- Every simple pole or zero of φ gives a marked circle in the corresponding component of $\partial\Sigma$
- Every exponential-type singularity of the type $\varphi \sim \exp(p(z))$ gives $\deg(p)$ many marked intervals in the corresponding component of $\partial\Sigma$

Consider the Fukaya category $\mathcal{F}(\Sigma)$ of this graded marked surface $(\Sigma, \partial\Sigma, \eta)$. The integral of $\sqrt{\varphi}$ gives a map

$$\int \sqrt{\varphi} : H_1(\Sigma, \partial\Sigma; \mathbb{Z}_\tau) \to \mathbb{C}$$

which composed with the natural map of Theorem 44 gives a central charge $Z : K_0(\mathcal{F}(\Sigma)) \to \mathbb{C}$. Now let us define full subcategories $\mathcal{P}_\phi$ for each $\phi \in \mathbb{R}$ to be spanned by objects represented by simple saddle connections of phase ϕ.

Theorem 58. *[56, Thm.5.1] The data $(Z, \mathcal{P})$ define a stability condition on $\mathcal{F}(\Sigma)$, and moreover the map*

$$\mathcal{M}(\Sigma) \to \mathrm{Stab}(\mathcal{F}(\Sigma))$$

is a local homeomorphism to its image.

This implies that the image of $\mathcal{M}(\Sigma)$ is an union of connected components of $\mathrm{Stab}(\mathcal{F}(\Sigma))$. Some simple explicit examples are also computed in [56], namely the disk and the annulus, where one can prove that there are no other components of $\mathrm{Stab}(\mathcal{F}(\Sigma))$; we will recall some of these calculations later in Section 5.3.

In the following chapter we will address this question by presenting a definition of a relative stability conditions on these categories. This definition will satisfy cutting and gluing relations, which allow us to prove that in the 'fully stopped' case these HKK stability conditions do in fact cover all of $\mathrm{Stab}(\mathcal{F}(\Sigma))$.

Chapter 5

Relative stability conditions

In this chapter, we will present the main parts of the author's work [107]. The starting point of that paper is the observation that HKK stability conditions, which are given by quadratic differentials as we have seen in the previous section, can be constructed by a local-to-global principle.

This is quite unusual for Bridgeland stability conditions. For example, suppose that one is given a category $\mathcal{C}_i$ presented as a colimit of a diagram of categories $\mathcal{C}_i, i \in I$ (as it is the case in the cosheaf description of the Fukaya category). Checking whether a pair $(Z, \mathcal{P})$ is a stability condition requires knowing the morphism spaces in $\mathcal{C}$ and checking the necessary conditions for giving a slicing, which turns out to be quite complicated in terms of the local categories $\mathcal{C}_i$.

Another argument for this is that the two parts of a stability condition, the central charge and the slicing, have opposite functorialities. Given a colimit $\mathcal{C} = \operatorname{colim}_{i \in I} \mathcal{C}_i$ with maps $\mathcal{C}_i \to \mathcal{C}$, the central charge naturally pulls back from $\mathcal{C}$ and the slicing naturally pushes forward to $\mathcal{C}$. While there are ways of getting around this in certain examples, for example using semiorthogonal decompositions such as in [29], those techniques are not directly applicable to our context.

Summary of results

The initial motivation for [107] is the observation that [56] provides an enticing counterexample to this trend, since it builds stability conditions on $\mathcal{F}(\Sigma)$ from geometric objects with nice functorial properties, namely flat structures, which glue along nicely under a decomposition of the surface. For example, given a decomposition of Σ into two pieces Σ_1 and Σ_2 mutually overlapping along a rectangular strip R, and a flat structure on Σ, restricting the flat structure to each side gives a flat structure (with the new boundary 'at infinity'). Moreover, once one defines the appropriate notion of compatibility between flat structures along the strip, one can glue compatible flat structures on Σ_1 and Σ_2 into a flat structure on Σ.

This work can be seen as an effort towards abstracting this idea of cutting and gluing into a construction that only makes reference to the stability conditions themselves. The appropriate local pieces of this construction will be given by the definition of *relative stability conditions* on a marked surface. A relative stability condition on Σ with respect to some unmarked boundary arc γ is an ordinary stability condition on another surface $\tilde{\Sigma}$, obtained from Σ by an appropriate modification along γ.

This definition behaves well under certain decompositions of surfaces. Let $\operatorname{RelStab}(\Sigma, \gamma)$ denote the set of relative stability conditions on Σ relative to γ. we prove that this set is naturally a Hausdorff space, with a topology inherited from the topology of the spaces of (ordinary) stability conditions. Consider a decomposition $\Sigma = \Sigma_L \cup_\gamma \Sigma_R$ into two surfaces glued along boundary arcs. Our main technical result is about the existence of cutting and gluing maps relating stability conditions on Σ and relative stability conditions on Σ_L and Σ_R.

Theorem 59. *There is a relation of compatibility along γ defining a subset $\Gamma \subset \operatorname{RelStab}(\Sigma_L, \gamma) \times \operatorname{RelStab}(\Sigma_R, \gamma)$ and continuous maps*

$$\operatorname{Stab}(\mathcal{F}(\Sigma)) \xrightarrow{cut} \Gamma \xrightarrow{glue} \operatorname{Stab}(\mathcal{F}(\Sigma))$$

which compose to the identity.

Consider now any marked graded surface Σ that is 'fully stopped', ie. every boundary circle has at least one marked interval. Assume also that at least one boundary circle has at least two marked intervals. In Section 5.3, we define a procedure for reducing the calculation of $\operatorname{Stab} \mathcal{F}(\Sigma)$ to the calculation of (ordinary) stability conditions on three base cases: the disk, the annulus and the punctured torus.

In all of these cases it can be shown that every stability condition is an HKK stability condition, ie. the map $\mathcal{M}(-) \to \operatorname{Stab}(\mathcal{F}(-))$ is an isomorphism. The cases of the disk and of the annulus are dealt with in [56], but the calculation for the case of the punctured torus is new. Theorem 59 implies that the gluing map $\Gamma \to \operatorname{Stab}(\mathcal{F}(\Sigma))$ is surjective, so knowing that all the base cases are fully described by HKK stability conditions we deduce the same for the surface Σ.

Theorem 60. *Every stability condition on $\mathcal{F}(\Sigma)$ is an HKK stability condition, ie. given by a flat structure on Σ.*

As mentioned above, this author believes that the value of this construction is not necessarily in its specific application to the case of Fukaya categories, but rather in its use for constructing and analyzing stability conditions sheaf-theoretically. It would be very fortunate if these tools could be rephrased in purely categorical terms, without direct reference to the geometry of Σ. In general terms, the idea is to define relative stability conditions on fully faithful functors $\mathcal{A} \to \mathcal{B}$ that can be glued to give stability conditions on pushouts of the form $\mathcal{B} \cup_{\mathcal{A}} \mathcal{B}'$.

For that purpose, we have tried to make the definitions of relative stability conditions as functorial (ie. independent of the explicit description of the surface) as we could, but it has not yet been possible to rephrase the relevant definitions and lemmas in such terms. In particular the theorems involving the cutting and gluing maps of Section 5.2 still depend on the underlying topological structure of the surfaces; one of the main questions to face before generalizing them to other types categories is to find equivalents of the 'non-crossing' Lemma 50.

It is likely that this kind of construction could be extended beyond Fukaya categories of surfaces; this motivates many possible directions of future study. One obvious such direction is towards extending the definition of relative stability conditions to wrapped Fukaya categories of higher-dimensional symplectic manifolds, using the microlocal model discussed previously in Chapter 2. Due to the sheaf-theoretic properties of the microlocal sheaf model, and the relatively simple local nature of the categories involved (ie. quiver representation categories), it appears that this model would be very suitable to the application of relative stability conditions, since the study of stability conditions on quiver representation categories is in general much simpler than on 'more geometric' categories.

The notion of relative stability conditions also opens up the possibility of using these sheaf-theoretic techniques to address some questions about dynamics on surfaces; the work of Dimitrov, Haiden, Katzarkov and Kontsevich [37, 35, 36] investigates the relation between dynamical systems on surfaces and stability conditions on their Fukaya category. The relation between Teichmüller theory and stability conditions was already noted in [24, 47], and in particular there is a close relation between the set of stable phases Φ (which we analyze for some cases in Section 5.3) and measures of dynamical entropy for categories. For now, the possible applications of our methods to such questions remain topic of current and future investigations.

5.1 Relative stability conditions on Fukaya categories of surfaces

In this section, we present a notion of stability conditions on a surface Σ *relative* to part of its boundary. This construction will exhibit functorial behavior and satisfy cutting and gluing relations. First we will give some presentations of the category $\mathcal{F}(\Sigma)$ that will be useful in stating that definition.

Pushouts

In [56], it is shown that given a full system of arcs on Σ, one can define a graph G dual to it and a constructible cosheaf $\mathcal{E}$ of A_∞-categories on G such that:

Theorem 61. *[56, Theorem 3.1] The category $\mathcal{F}(\Sigma)$ represents global sections of the cosheaf $\mathcal{E}$, ie. is the homotopy colimit of the corresponding diagram of A_∞-categories.*

We will describe how to use this result to express $\mathcal{F}(\Sigma)$ as certain useful homotopy colimits. Let γ be some embedded interval dividing Σ into two surfaces, Σ_L and Σ_R. Suppose that we have a chain of intervals $\gamma_1, \ldots, \gamma_N$ in distinct isotopy classes connecting $n+1$ distinct marked boundary intervals $M_0, \ldots, M_n$, such that their concatenation gives the interval γ.

Lemma 62. *Σ admits a full system of arcs $\mathcal{A} = \mathcal{A}_L \sqcup \mathcal{A}_\gamma \sqcup \mathcal{A}_R$ such that every arc in $\mathcal{A}_L$ has a representative contained in Σ_L, every arc in $\mathcal{A}_R$ has a representative contained in Σ_R, and $\mathcal{A}_\gamma = \{\gamma_1, \ldots, \gamma_N\}$.*

Proof. Consider a (non-full) system of arcs $\overline{\mathcal{A}}_\gamma$ given by the 'closure' of $\mathcal{A}_\gamma = \{\gamma_1, \ldots, \gamma_N\}$; that is containing also a chain of arcs connecting all the marked boundary intervals to the left of the chain γ, and the analogous chain to the right of it.

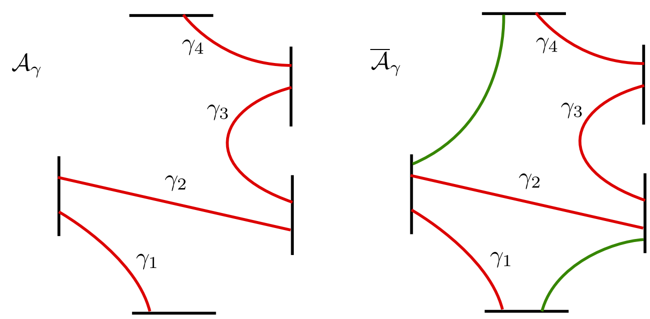

Figure 5.1: The (non-full) system of arcs $\mathcal{A}_\gamma$ and its closure $\overline{\mathcal{A}}_\gamma$. The green arcs are elements of $\overline{\mathcal{A}}_\gamma \setminus \mathcal{A}_\gamma$.

Since all the intervals in $\overline{\mathcal{A}}_\gamma$ are non-intersecting and not pairwise isotopic there is some full arc system $\mathcal{A}$ of Σ containing them; and since γ (and therefore the chain made by the γ_i) cuts the surface into two we can partition the arcs $\mathcal{A}$ that are not among the γ_i into left and right subsets $\mathcal{A}_L$ and $\mathcal{A}_R$. By construction every arc in $\mathcal{A}_L$ is contained in Σ_L and every arc in $\mathcal{A}_R$ is contained in Σ_R. $\square$

Consider this arc system $\mathcal{A}$. Let us define $\tilde{\Sigma}_L$ to be the smallest marked surface with an inclusion into Σ that contains all the arcs in $\mathcal{A}_L \sqcup \mathcal{A}_\gamma$; we define $\tilde{\Sigma}_R$ analogously.

We see that topologically, $\tilde{\Sigma}_L, \tilde{\Sigma}_R$ can being constructed from Σ_L, Σ_R by attaching a disk along γ, that is

$$\tilde{\Sigma}_L = \Sigma_L \cup_\gamma \Delta_m, \qquad \tilde{\Sigma}_R = \Sigma_R \cup_\gamma \Delta_n$$

where Δ_k is the disk with k marked boundary intervals. By minimality of these surfaces, we must have $(m-2) + (n-2) = N - 1$.

Let us denote the triangulated closure of the object represented in an arc system by $\langle \mathcal{A} \rangle$. Then we have $\mathcal{F}(\tilde{\Sigma}_L) = \langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ and $\mathcal{F}(\tilde{\Sigma}_R) = \langle \mathcal{A}_R \sqcup \mathcal{A}_\gamma \rangle$. Using the cosheaf description

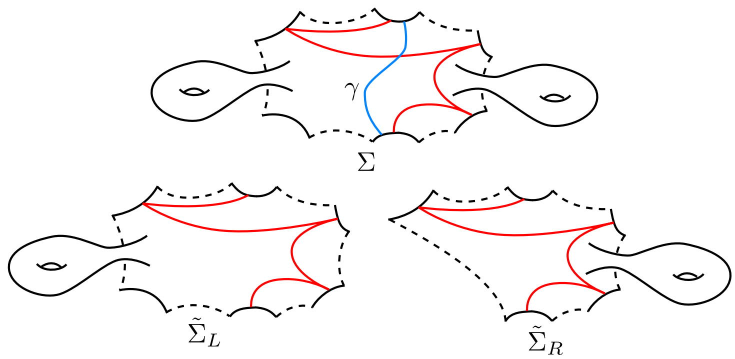

Figure 5.2: The two 'modified' surfaces $\tilde{\Sigma}_L$ and $\tilde{\Sigma}_R$. Each one is obtained from Σ_L, Σ_R respectively by adding more marked intervals (m and n of them) along the boundary according to the chain. In this example $N = 4, m = 2, n = 1$.

above we can assemble all these categories into the following cube diagram:

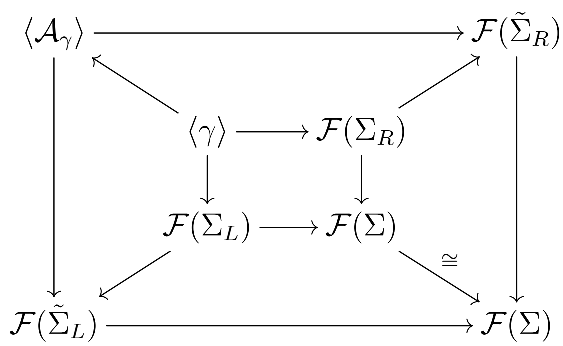

where the inner and outer squares, and the top and left sides are all pushouts (ie. homotopy colimits).

Main definitions

Consider now some surface S with an embedded interval γ which connects two adjacent marked boundary intervals M, M', and runs parallel to the unmarked boundary interval between them (for example we can take $(S, \gamma) = (\Sigma_L, \gamma)$ as above).

Definition 29. A *relative stability condition* on the pair (S, γ) is the data of:

- A surface $\tilde{S} = S \cup_\gamma \Delta_n$ where Δ_n is a disk with n marked boundary intervals, with a given inclusion map $S \hookrightarrow \tilde{S}$,
- A stability condition $\tilde{\sigma} \in \mathrm{Stab}(\mathcal{F}(\tilde{S}))$.

Note that the first condition implies that the embedded interval $\gamma \subset \tilde{S}$ cuts the surface into two, so by Lemma 52 any indecomposable object C supported on γ has a simple COSI decomposition under $\tilde{\sigma}$.

Fix a relative stability condition $\sigma = (Z, \mathcal{P})$ and let us denote by $C_1, \ldots, C_N$ the corresponding chain of stable intervals in the decomposition of C, supported on arcs $\gamma_1, \ldots, \gamma_N$. As in the previous subsection, we can take $(\Sigma_L, \Sigma_R) = (S, \Delta_n)$; this defines an arc system $\mathcal{A}_L \sqcup \mathcal{A}_\gamma \sqcup \mathcal{A}_R$ on $\tilde{S}$.

Restricting stability conditions and minimality

Consider now the central charges

$$Z_L = Z|_{\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle}, \qquad Z_R = Z|_{\langle \mathcal{A}_\gamma \sqcup \mathcal{A}_R \rangle}$$

and the 'candidates for slicings' $\mathcal{P}_L, \mathcal{P}_R$, given by intersecting the full triangulated subcategories $\mathcal{P}_\phi$ with the full triangulated subcategories $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle, \langle \mathcal{A}_\gamma \sqcup \mathcal{A}_R \rangle$, respectively.

Lemma 63. *$\sigma|_L = (Z_L, \mathcal{P}_L)$ and $\sigma|_R = (Z_R, \mathcal{P}_R)$ give stability conditions on the subcategories $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ and $\langle \mathcal{A}_\gamma \sqcup \mathcal{A}_R \rangle$.*

Proof. The compatibility between the central charges and filtrations is obvious by construction; we only need to check that $\mathcal{P}_L, \mathcal{P}_R$ do in fact give slicings, ie. that every object in either category has an HN decomposition by objects in each restricted slicing. This can be checked on indecomposable objects and follows from Lemma 50; every indecomposable object on either side can be represented by some immersed curve keeping to the same side of the chain γ, so therefore its HN decomposition under the original stability condition σ cannot cross to the other side. □

Note that this construction $\sigma \to (\sigma|_L, \sigma|_R)$ does not give a map from the entire stability space $\mathrm{Stab}(\mathcal{F}(\tilde{S}))$ to any other stability space; as σ varies, the target categories $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ change since the decomposition of the interval object C changes as we cross a wall. However, this only happens across some specific kinds of walls, defined by the following condition:

Definition 30. The relative stability condition σ is *non-reduced* if there are two interval objects C_i, C_{i+1} extended on the right (ie. by an extension map $C_{i+1} \xrightarrow{+1} C_i$), with the same phase.

By standard results [24], the subset of non-reduced stability conditions is contained in a locally finite union of walls of $\mathrm{Stab}(\mathcal{F}(\tilde{S}))$ walls, so the subset of reduced stability conditions is composed of open chambers.

Lemma 64. *Within each chamber $\mathcal{C}$ of reduced relative stability conditions, the target subcategory $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ is constant and the map $\mathrm{Stab}(\mathcal{F}(\tilde{S})) \to \mathrm{Stab}(\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle)$ is continuous.*

Proof. Within each reduced chamber $\mathcal{C}$, the chain γ is constant except for the (internal) walls on which two (or more) adjacent interval objects of the same phase C_i, C_{i+1} are extended on the left (ie. by an extension map $C_i \xrightarrow{+1} C_{i+1}$). However, though the chain $\mathcal{A}_\gamma$ changes across such a wall, by construction of $\mathcal{A}_L$ we see that $\langle \mathcal{A}_\gamma \sqcup \mathcal{A}_L \rangle$ stays constant. Continuity follows from the fact that a small enough neighborhood of every stability condition on some category $\mathcal{D}$ is isomorphic to $(K_0(\mathcal{D}))^\vee = \mathrm{Hom}_{\mathbb{Z}}(K_0(\mathcal{D}), \mathbb{C})$ and in that neighborhood the map $\mathrm{Stab}(\mathcal{F}(\tilde{S})) \to \mathrm{Stab}(\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle)$ is described by the projection dual to the inclusion $K_0(\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle) \to K_0(\mathcal{F}(\tilde{S}))$. □

For our later uses, we would like to define a notion of minimality, in the sense that the integer n of marked boundary intervals of Δ_n is as small as possible.

Definition 31. A relative stability condition σ on (S, γ) *minimal* if every marked boundary interval of Δ_n appears in the simple chain of stable intervals decomposition of C.

Another way of phrasing the minimality condition is:

Lemma 65. *σ is minimal if and only if and $\langle \mathcal{A}_R \rangle \subseteq \langle \mathcal{A}_\gamma \rangle$.*

The space of relative stability conditions

For our purposes, the part of the stability condition 'on the disk side' does not matter; we realize this by using an equivalence relation. Let $\sigma \in \mathrm{Stab}(\mathcal{F}(\tilde{S} = S \cup \Delta_m))$ and $\sigma' \in \mathrm{Stab}(\mathcal{F}(\tilde{S}' = S \cup \Delta_n))$ be two relative stability conditions on (S, γ). As above, one can (non-uniquely) pick corresponding arc systems $\mathcal{A}_L \sqcup \mathcal{A}_\gamma \sqcup \mathcal{A}_R$ and $\mathcal{A}'_L \sqcup \mathcal{A}'_\gamma \sqcup \mathcal{A}'_R$ on $\tilde{S}$ and $\tilde{S}'$, and restrict stability conditions to each side.

We will see that we need to be careful about genericity when defining the correct equivalence relation. For motivation let us first define a naive notion of equivalence:

Definition 32. (Naive equivalence) $\sigma \sim_{\text{naive}} \sigma'$ if there is an equivalence of categories

$$\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle \cong \langle \mathcal{A}'_L \sqcup \mathcal{A}'_\gamma \rangle$$

(compatible with the embedding of $\mathcal{F}(S)$ on both sides) such that the restricted stability conditions $\sigma|_L$ and $\sigma'|_L$ agree.

It is clear from the definition above that $\sim_{\text{naive}}$ defines an equivalence relation on the set of relative stability conditions on (S, γ). We would like to define the space of relative stability conditions as the quotient of the space

$$\mathbb{S} = \bigsqcup_{n \geq 2} \mathrm{Stab}(\mathcal{F}(S \cup_\gamma \Delta_n))$$

by the relation $\sim_{\text{naive}}$, but it turns out that this space is ill-behaved. For instance, it is not Hausdorff, because the graph $\Gamma_{\sim_{\text{naive}}} \subset \mathbb{S} \times \mathbb{S}$ of the naive relation is not a closed subset.

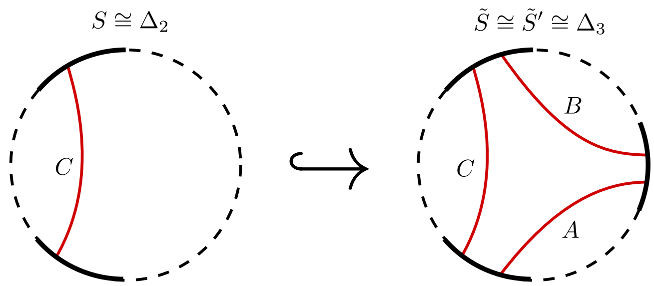

Figure 5.3: The surfaces $S \cong \Delta_2$ and $\tilde{S} \cong \tilde{S}' \cong \Delta_3$. The category $\mathcal{F}(S)$ is equivalent to $\mathrm{Mod}(A_1)$ and $\mathcal{F}(\tilde{S})$ is equivalent to $\mathrm{Mod}(A_2)$.

Example. Take the simple example where $S \cong \Delta_2$ with unique (up to shift) indecomposable object C and $\tilde{S} \cong \tilde{S}' \cong \Delta_3$, with objects A, B, C as below.

We have a distinguished triangle $A \to C \to B$. Consider two infinite families of stability conditions on $\mathcal{F}(\Delta_3)$, $\{\sigma_m = (Z_m, \mathcal{P}_m)\}$ and $\{\sigma'_m = (Z'_m, \mathcal{P}'_m)\}$ with $m \in \mathbb{Z}_+$, on $\mathcal{F}(\Delta_3)$ given by the central charges

$$Z_m(A) = \frac{1}{3} + i\frac{1}{m}, \qquad Z_m(B) = \frac{2}{3} - i\frac{1}{m}$$

$$Z'_m(A) = \frac{2}{3} + i\frac{1}{m}, \qquad Z'_m(B) = \frac{1}{3} - i\frac{1}{m}$$

with A, B and C stable in all of them, picking phases for all these objects between $-1/2$ and $1/2$. Each one of these sequences converges in $\mathrm{Stab}(\mathcal{F}(\Delta_3))$ respectively, to the stability conditions $\sigma_\infty, \sigma'_\infty$ with central charges

$$Z_\infty(A) = \frac{1}{3}, \qquad Z_\infty(B) = \frac{2}{3}$$

$$Z'_\infty(A) = \frac{2}{3}, \qquad Z'_\infty(B) = \frac{1}{3}$$

where A, B are stable but C is only semistable, with Jordan-Hölder factors A, B.

Seen as relative stability conditions on (Δ_2, γ), all the σ_m, σ'_m for any m are equivalent under $\sim_{\text{naive}}$; the subcategory $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ is $\mathcal{F}(\Delta_2) = \langle C \rangle$ and the central charge of C is 1 for all finite m. On the other hand, σ_∞ and σ'_∞ are not equivalent under $\sim_{\text{naive}}$, since for those two $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ is the whole category. Thus $(\sigma_\infty, \sigma'_\infty) \in \overline{\Gamma_{\sim\text{naive}}} \setminus \Gamma_{\sim\text{naive}}$.

As in the example above, the problem always arises when we have relative stability conditions which are non-reduced. Consider a relative condition σ on (S, γ) given by a stability condition on $\mathcal{F}(\tilde{S})$ for some $\tilde{S} = S \cup_\gamma \Delta_n$, where the object C supported on γ has a COSI decomposition $C_1, \ldots, C_N$. Assume that σ is non-reduced; this means that there is a nonempty set of indices $R \subset \{1, \ldots, N\}$ such that the extension map is 'on the right' (ie. $\in \mathrm{Ext}^1(C_{i+1}, C_i)$) and C_i and C_{i+1} have the same phase. Let us suppose that the set R is of the form $j, j+1, \ldots, j+m$ for some $1 \leq j \leq j+m \leq N-2$ with all objects $C_j, \ldots, C_{j+m+1}$ having the same phase ϕ; the general case (where R is the disjoint union of a number of those subsets) will not be any more difficult.

Consider now the reduced arc system given by

$$\mathcal{A}_\gamma^{\mathrm{red}} = \{\gamma_1, \ldots, \gamma_{j-1}, \tilde{\gamma}, \gamma_{j+m+2}, \ldots, \gamma_N\},$$

where $\tilde{\gamma}$ is obtained by concatenating the intervals $\gamma_j, \ldots, \gamma_{j+m+1}$ at the m marked boundaries M_i with index $i \in R$. Let us now define a reduced restriction σ^{red} given by restricting the data of σ to the subcategory $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma^{\mathrm{red}} \rangle$, and then adding to the category $\mathcal{P}_\phi$ the objects supported on $\tilde{\gamma}$.

Lemma 66. *σ^{red} is a stability condition.*

Proof. It suffices to prove that every object in the subcategory $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma^{\mathrm{red}} \rangle$ has an HN decomposition into stable objects also in that same subcategory. Because of Lemma 50, the only way this could fail is if there is some indecomposable object X of $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma^{\mathrm{red}} \rangle$ in whose decomposition some but not all of the stable interval objects $C_j, \ldots, C_{j+m+1}$ appear (if all of them appear we just replace that semistable object with the stable object $\tilde{C}$ supported on $\tilde{\gamma}$). But this cannot happen for phase reasons, following a similar argument as the proof of Lemma 50. □

For completeness let us define $\sigma^{\mathrm{red}} = \sigma|_L$ if σ is reduced. With this definition we can now state the correct notion of equivalence.

Definition 33. (Equivalence) $\sigma \sim \sigma'$ if there is an equivalence of categories

$$\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma^{\mathrm{red}} \rangle \cong \langle \mathcal{A}'_L \sqcup \mathcal{A}'^{\mathrm{red}}_\gamma \rangle$$

(compatible with the embedding of $\mathcal{F}(S)$ on both sides) such that the reduced restricted stability conditions σ^{red} and σ'^{red} agree.

It is clear from the definition that $\sim$ is an equivalence relation on the set $\mathbb{S} = \bigsqcup_{n \geq 2} \mathrm{Stab}(\mathcal{F}(S \cup_\gamma \Delta_n))$.

Lemma 67. *There is a unique minimal and reduced relative stability condition in each equivalence class of the equivalence relation $\sim$.*

Proof. Consider some relative stability condition σ; as above it defines a stability condition σ^{red} on the subcategory $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma^{\mathrm{red}} \rangle$. Note that this subcategory is also of the form $\mathcal{F}(S \cup_\gamma \Delta_n)$, with $n = |\mathcal{A}_\gamma^{\mathrm{red}}| + 1$, and also by construction σ is equivalent to the reduced σ^{red} when both are viewed as relative stability conditions on (S, γ).

Suppose now that we have two stability conditions $\sigma \sim \sigma'$ which are minimal and thus reduced; then the arcs in $\mathcal{A}_R, \mathcal{A}'_R$ can be generated by the other arcs so by compatibility we have

$$\mathcal{F}(\tilde{S}) \cong \langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle \cong \langle \mathcal{A}'_L \sqcup \mathcal{A}'_\gamma \rangle \cong \mathcal{F}(\tilde{S}'),$$

but it is easy to see that no two categories $\mathcal{F}(S \cup_\gamma \Delta_n)$ are equivalent for different n (for example by taking K_0) so $\tilde{S} \cong \tilde{S}'$ (compatibly with the embedding of S) with equivalent stability conditions. □

Definition 34. (Space of relative stability conditions) Let us define $\mathrm{RelStab}(S, \gamma)$ as the set of minimal and reduced stability conditions; this set is given the quotient topology by the identification $\mathrm{RelStab}(S, \gamma) = \mathbb{S}/\sim$,

Proposition 68. *The space* $\mathrm{RelStab}(S, \gamma)$ *is Hausdorff.*

Proof. This is equivalent to showing that the graph $\Gamma_\sim$ of the equivalence relation is closed in $\mathbb{S} \times \mathbb{S}$. Since $\mathbb{S}$ is an disjoint union this is equivalent to showing $\Gamma_\sim$ is closed in each component $\mathrm{Stab}(\mathcal{F}(\tilde{S})) \times \mathrm{Stab}(\mathcal{F}(\tilde{S}'))$.

The spaces $\mathrm{Stab}(\mathcal{F}(\tilde{S}))$ have a wall-and-chamber structure where the walls are the locus of non-reduced stability conditions. By standard arguments, the union of all walls is a locally finite union of real codimension one subsets. The complement is composed of open chambers, and by Lemma 64 the target subcategory $\mathcal{T} = \langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ is constant on each chamber.

In the interior of each chamber

$$\mathcal{C} = \mathcal{C}_\rho \times \mathcal{C}_\sigma \subset \mathrm{Stab}(\mathcal{F}(\tilde{S})) \times \mathrm{Stab}(\mathcal{F}(\tilde{S}')),$$

the locus $\Gamma_\sim$ is the preimage of the diagonal $\Delta \subset \mathrm{Stab}(\mathcal{T}) \times \mathrm{Stab}(\mathcal{T})$, so it is closed by continuity.

Let us look at the walls surrounding the chamber $\mathcal{C}$, and start with a simple codimension one wall W, ie. the locus at the boundary of $\mathcal{C}$ where the phases ϕ_i, ϕ_{i+1} of two adjacent interval objects C_i, C_{i+1} (with an extension to the right) agree. There are two possibilities: $\phi_i < \phi_{i+1}$ or $\phi_i > \phi_{i+1}$ inside of $\mathcal{C}$. In the former case, comparing the target categories we see that the reduced target category $\mathcal{T}_W^{\mathrm{red}}$ on the wall is equal to the usual target category $\mathcal{T}_\mathcal{C}$ in the interior of the chamber, so we can apply the same argument as inside the chamber and conclude that $\Gamma_\sim \cap W$ is closed.

In the latter case $\mathcal{T}_W^{\mathrm{red}}$ is smaller than $\mathcal{T}_\mathcal{C}$, as it doesn't contain the objects C_i, C_{i+1}, only their extension. However, the closure $\overline{\Gamma_\sim \cap \mathcal{C}}$ meets W along a closed locus contained within $\Gamma_\sim \cap W$, as the reduced equivalence condition is strictly weaker than the naive equivalence condition on W. The general case for walls of higher codimension is essentially the same and can be obtained iteratively.

Now, over the entire space $\mathrm{Stab}(\mathcal{F}(\tilde{S})) \times \mathrm{Stab}(\mathcal{F}(\tilde{S}'))$, since each point is surrounded by finitely many reduced chambers and $\Gamma_\sim$ is closed within the closure of each one of them, $\Gamma_\sim$ is the locally finite union of closed subsets. □

Remark. Unlike the space of stability conditions $\mathrm{Stab}(\mathcal{F}(S))$, the space $\mathrm{RelStab}(S, \gamma)$ is not a complex manifold; in fact it is a stratified space, with cells of unbounded dimension.

Compatibility

Consider now two surfaces S and S' with embedded intervals γ, γ' and relative stability conditions $\sigma \in \mathrm{RelStab}(S,\gamma)$ and $\sigma \in \mathrm{RelStab}(S',\gamma')$. Given any two such surfaces, we can glue them by identifying $\gamma = \gamma'$ and obtain a surface $S \cup_\gamma S'$. Since there is a full arc system on this surface containing the arc γ, one can take the ribbon graph dual to this arc system and get a pushout presentation

$$\mathcal{F}(S \cup_\gamma S') = \mathcal{F}(S) \cup_{\mathcal{F}(\gamma)} \mathcal{F}(S').$$

The relative stability conditions σ, σ' have unique minimal and reduced representatives by Lemma 67. However they also have many minimal but non-reduced representatives.

Definition 35. A *compatibility structure* between σ and σ' is the following data:

- Minimal representatives $\tilde{\sigma} \in \mathrm{Stab}(\mathcal{F}(\tilde{S}))$ and $\tilde{\sigma}' \in \mathrm{Stab}(\mathcal{F}(\tilde{S}'))$ of σ and σ'.
- Inclusions of surfaces

$$S \hookrightarrow \tilde{S} \hookrightarrow S \cup_\gamma S', \qquad S' \hookrightarrow \tilde{S}' \hookrightarrow S \cup_\gamma S',$$

such that the images of the embedded intervals in the COSI decompositions of γ and γ' agree as an arc system $\mathcal{A}_\gamma$ inside of $S \cup_\gamma S'$, and the restrictions $\tilde{\sigma}|_{\langle \mathcal{A}_\gamma \rangle}$ and $\tilde{\sigma}|_{\langle \mathcal{A}_\gamma \rangle}$ are the same stability condition in $\mathrm{Stab}(\langle \mathcal{A}_\gamma \rangle)$.

5.2 Cutting and gluing relative stability conditions

In this section, we will explain how to cut (ordinary) stability conditions into relative stability conditions and glue relative stability conditions into (ordinary) stability conditions. This will allow us to reduce the calculations of stability conditions on general surfaces Σ to the calculation of stability conditions on simpler surfaces. Before we present these procedures, we will need to use the following generalization of a slicing.

Definition 36. A *pre-slicing* $\mathcal{P}^{pre}$ on a category $\mathcal{C}$ is a choice of full triangulated subcategories $\mathcal{P}^{pre}_\phi$ for every $\phi \in \mathbb{R}$, such that $\mathrm{Hom}(X,Y) = 0$ if $X \in \mathcal{P}^{pre}_\phi$ and $Y \in \mathcal{P}^{pre}_\psi$, $\phi > \psi$.

Remark. This is the same data as a slicing, except that we don't require the existence of Harder-Narasimhan decompositions for objects.

Definition 37. A *pre-stability condition* on $\mathcal{C}$ is the data of a central charge function $Z : K_0(\mathcal{C}) \to \mathbb{C}$ and a pre-slicing $\mathcal{P}^{pre}$ satisfying the usual compatibility condition $Z(X)/|Z(X)| = e^{i\pi\phi}$ if $X \in \mathcal{P}^{pre}_\phi$.

Let us denote by $\mathrm{PreStab}(\mathcal{C})$ the set of all pre-stability conditions on $\mathcal{C}$. It is obvious that we have an inclusion of sets

$$\mathrm{Stab}(\mathcal{C}) \hookrightarrow \mathrm{PreStab}(\mathcal{C}).$$

Cutting stability conditions

We return to the setting of a surface Σ that is cut into Σ_L, Σ_R by an embedded interval γ supporting a rank one object C.

Consider a stability condition $\sigma \in \mathrm{Stab}(\mathcal{F}(\Sigma))$. By Corollary 56, the object C has a simple COSI decomposition into objects $C_1, \ldots, C_N$ supported on arcs $\gamma_1, \ldots, \gamma_N$, which connect the marked boundary intervals $M_0, \ldots, M_N$. As in subsection 5.1, there is then a full system of arcs

$$\mathcal{A} = \mathcal{A}_L \sqcup \mathcal{A}_\gamma \sqcup \mathcal{A}_R$$

such that every arc in $\mathcal{A}_L$ has a representative contained in Σ_L, every arc in $\mathcal{A}_R$ has a representative contained in Σ_R, and $\mathcal{A}_\gamma = \{\gamma_1, \ldots, \gamma_N\}$.

Each extension between C_i and C_{i+1} happens either on the left (ie. by an extension map $C_i \xrightarrow{+1} C_{i+1}$) or on the right (ie. by an extension map $C_{i+1} \xrightarrow{+1} C_i$). Let m, n be the numbers of indices with extension on the left and right, respectively, plus 2; we have by definition $m - 2 + n - 2 = N + 1 =$ number of marked boundary intervals along the chain.

Then we have surfaces $\tilde{\Sigma}_L = \Sigma_L \cup_\gamma \Delta_m$ and $\tilde{\Sigma}_R = \Sigma_R \cup_\gamma \Delta_n$ such that

$$\mathcal{F}(\tilde{\Sigma}_L) = \langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle, \qquad \mathcal{F}(\tilde{\Sigma}_R) = \langle \mathcal{A}_R \sqcup \mathcal{A}_\gamma \rangle.$$

Consider the restrictions

$$\sigma_L = \sigma|_{\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle}, \qquad \sigma_R = \sigma|_{\langle \mathcal{A}_\gamma \sqcup \mathcal{A}_R \rangle}$$

that is, as in the previous section we take the data given by restricting the central charges and intersecting the slicings with each full subcategory.

Lemma 69. *σ_L, σ_R are stability conditions on $\mathcal{F}(\tilde{\Sigma}_L), \mathcal{F}(\tilde{\Sigma}_R)$.*

Proof. The condition $Z(X) = m(X) \exp(i\pi\phi_X)$ on every semistable object X is satisfied by construction, so we just need to check that every object $X \in \mathcal{F}_L$ has a HN filtration, ie. that $\mathcal{P}_L$ indeed defines a slicing.

It is enough to check this on indecomposable objects. By geometricity, every such object X is represented by an immersed curve in $\tilde{\Sigma}_L$ with indecomposable local system. Consider its image in $\mathcal{F}(\Sigma)$ which is also an immersed curve, and its chain-of-interval decomposition under σ.

If X is an interval object, then both of its ends are on marked boundary components belonging to $\tilde{\Sigma}_L$, and since the associated chain of intervals is isotopic to the support of X, if any of those intervals in in Σ_R, then the chain must cross back to Σ_L, creating a polygon of the sort prohibited by Lemma 50. And if X is a circle object then it is by definition supported on a non-nullhomotopic immersed circle, so by the same argument its chain of intervals cannot cross over to Σ_R without also creating a prohibited polygon. Thus every stable component of the HN decomposition is in $\mathcal{F}_L$. □

We then use the inclusions of marked surfaces $\Sigma_L \hookrightarrow \tilde{\Sigma}_L$ and $\Sigma_R \hookrightarrow \tilde{\Sigma}_R$ to interpret these stability conditions as relative stability conditions:

Definition 38. The cutting map

$$cut_\gamma : \mathrm{Stab}(\mathcal{F}(\Sigma)) \to \mathrm{RelStab}(\Sigma_L, \gamma) \times \mathrm{RelStab}(\Sigma_R, \gamma)$$

sends a stability conditions σ as above to the image of the stability conditions (σ_L, σ_R).

By Lemma 67 every element of RelStab has a unique minimal and reduced representative, so we can alternatively define the cutting map by using the 'reduced restriction' of Lemma 66

$$cut_\gamma(\sigma) = (\sigma_L^{\mathrm{red}}, \sigma_R^{\mathrm{red}}).$$

Lemma 70. *The map* $\mathrm{Stab}(\mathcal{F}(\Sigma)) \xrightarrow{cut_\gamma} \mathrm{RelStab}(\Sigma_L, \gamma) \times \mathrm{RelStab}(\Sigma_R, \gamma)$ *is continuous.*

Proof. We must look separately at the maps to each side; let us prove continuity of the map $\mathrm{Stab}(\mathcal{F}(\Sigma)) \xrightarrow{cut_L} \mathrm{RelStab}(\Sigma_L, \gamma)$. Recall that in subsection 5.1 we define the topology on the RelStab spaces as the quotient topology inherited from $\mathbb{S} = \bigsqcup_n \mathrm{Stab}(S \cup_\gamma \Delta_n)$.

Note that the construction for the map cut_L does not give a manifestly continuous map since the target $\mathcal{T} = \langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ changes across walls in $\mathrm{Stab}(\mathcal{F}(\Sigma))$. We remediate this by locally defining other maps that are continuous, and which agree with cut_L after identifying by the equivalence relation $\sim$.

Let σ be a stability condition on $\mathcal{F}(\Sigma)$ such that $\sigma_L = \sigma|_{\langle \mathcal{A}_L \cup \mathcal{A}_\gamma \rangle}$ is a non-reduced stability condition, and let us say that under σ the object C supported on γ has a decomposition into $C_1, \ldots, C_N$ supported on embedded intervals $\gamma_1, \ldots, \gamma_N$ with respective phases $\phi_1, \ldots, \phi_N$. Non-reducedness means that there is some collection of indices i such that C_i, C_{i+1} have the same phase, and are extended on the right. For simplicity, suppose first that we have a single such index; the general case can be deduced by iterating this argument. Let us denote C_{bot} to be the object obtained by concatenating $C_1, \ldots, C_i$, and C_{top} to be the object obtained by concatenating $C_{i+1}, \ldots, C_N$.

By standard arguments, the locus on which the objects $C_1, \ldots, C_N$ are simple is open, so there is a neighborhood $U \ni \sigma$ on which all these objects are simple, and with a complex isomorphism $U \cong (K_0(\mathcal{F}(\Sigma)))^\vee$. If necessary we further restrict U such that on this open set the $\phi_{i-1} \neq \phi_i$ and $\phi_{i+1} \neq \phi_{i+2}$. This implies that on U the chains $C_1, \ldots, C_i$ and $C_{i+1}, \ldots, C_N$ gives cosi decompositions of C_{bot} and C_{top}, respectively.

Consider now a fixed target category $\mathcal{T}_{\mathrm{fix}}$ given by the target $\mathcal{T}_\sigma = \langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ at σ. We argue that for every stability condition $\sigma' \in U$, $\sigma'|_{\mathcal{T}_{\mathrm{fix}}}$ is a stability condition. Note that this doesn't follow immediately from Lemma 50 since along some chambers in U, the pair C_i, C_{i+1} is not the cosi decomposition of any object so we cannot directly use the non-crossing argument.

Nevertheless, we can use a small modification of that argument. Consider some indecomposable object X in the subcategory $\mathcal{T}_{\mathrm{fix}}$; by geometricity this can be represented by an immersed curve ξ to the left of the chain of intervals, and by the results of Chapter 4, X has a cosi decomposition into intervals $\xi_1, \ldots, \xi_M$ whose concatenation is isotopic to ξ.

Now, since both ends of ξ are to the left of the γ chain, and this chain is divided into two stable chains, extended on the left, the only way that the ξ chain can cross the γ chain is it if crosses the chain for C_{bot} or C_{top} (or both). But again this is prohibited by the noncrossing argument of Lemma 50.

Thus this defines a map $\widetilde{cut}_\gamma : U \to \mathrm{Stab}(\mathcal{T}_{\mathrm{fix}})$ which by construction is continuous and agrees with cut_γ on U; doing this for every wall gives continuity of cut_γ. □

Note that by construction we have representatives $\sigma_L \in \mathrm{Stab}(\mathcal{F}(\tilde{\Sigma}_L))$ and $\sigma_R \in \mathrm{Stab}(\mathcal{F}(\tilde{\Sigma}_R))$ of the relative stability conditions $\sigma_L^{\mathrm{red}}, \sigma_R^{\mathrm{red}}$, and also inclusions of surfaces $\tilde{\Sigma}_L \hookrightarrow \Sigma$ and $\tilde{\Sigma}_R \hookrightarrow \Sigma$. It follows directly from the construction above that:

Lemma 71. *This is a compatibility structure between σ_L^{red} and σ_R^{red}.*

Gluing stability conditions

As in the previous section consider a surface cut into two parts by an embedded interval $\Sigma = \Sigma_L \cup_\gamma \Sigma_R$. Suppose we have relative stability conditions $\sigma_L \in \mathrm{RelStab}(\Sigma_L, \gamma)$ and $\sigma_R \in \mathrm{RelStab}(\Sigma_R, \gamma)$ with some compatibility structure between them (as in Definition 35).

Unpacking this data, we have non-negative integers m and n and stability conditions $\sigma_L = (Z_L, \mathcal{P}_L)$ on

$$\mathcal{F}_L = \mathcal{F}(\tilde{\Sigma}_L) = \mathcal{F}(\Sigma_L \cup_\gamma \Delta_m)$$

and $\sigma_R = (Z_R, \mathcal{P}_R)$ on

$$\mathcal{F}_R = \mathcal{F}(\tilde{\Sigma}_R) = \mathcal{F}(\Sigma_R \cup_\gamma \Delta_n)$$

representing σ_L, σ_R, together with inclusions of marked surfaces $\Sigma_L \hookrightarrow \tilde{\Sigma}_L \hookrightarrow \Sigma$ and $\Sigma_R \hookrightarrow \tilde{\Sigma}_R \hookrightarrow \Sigma$.

The compatibility condition implies that the chain-of-intervals decomposition $C_1^L, \ldots, C_N^L$ of the indecomposable object $C^L \in \mathcal{F}_L$ supported on $\gamma \subset \tilde{\Sigma}_L$ and the chain-of-intervals decomposition $C_1^R, \ldots, C_N^R$ of the indecomposable object $C^R \in \mathcal{F}_R$ supported on $\gamma \subset \tilde{\Sigma}_R$ are of the same length N on both sides, and that the central charges agree, ie.

$$Z_L(C_i^L) = Z_R(C_i^R)$$

for all i. Also compatibility also requires that the extension maps η_i^L and η_i^R go the same direction, ie. either both go forward

$$\eta_i^L \in \mathrm{Ext}^1(C_i^L, C_{i+1}^L) \text{ and } \eta_i^R \in \mathrm{Ext}^1(C_i^R, C_{i+1}^R)$$

or both go backward

$$\eta_i^L \in \mathrm{Ext}^1(C_{i+1}^L, C_i^L) \text{ and } \eta_i^R \in \mathrm{Ext}^1(C_{i+1}^R, C_i^R)$$

so we have the relation $(m-2)+(n-2)=N-1$ due to minimality of σ_L and σ_R.

The compatibility structure gives an identification between the images of $C_1^L, \dots, C_N^L$ and $C_1^L, \dots, C_N^L$ inside of $\mathcal{F}(\Sigma)$; we denote this full subcategory spanned by these arcs $\langle \mathcal{A}_\gamma \rangle$ as in previous sections. This gives a pushout presentation

$$\begin{array}{ccc} \langle \mathcal{A}_\gamma \rangle & \longrightarrow & \mathcal{F}_R \\ \downarrow & & \downarrow j_R \\ \mathcal{F}_L & \xrightarrow{j_L} & \mathcal{F}(\Sigma) \end{array}$$

From this data we will produce a central charge function $K_0(\mathcal{F}(\Sigma)) \to \mathbb{C}$ and a pre-slicing $\mathcal{P}$ on $\mathcal{F}(\Sigma)$.

The central charge

Applying the functor K_0 to the pushout above gives us a diagram of $\mathbb{Z}$-modules

$$\begin{array}{ccc} K_0(\langle \mathcal{A}_\gamma \rangle) & \longrightarrow & K_0(\mathcal{F}_R) \\ \downarrow & & \downarrow \\ K_0(\mathcal{F}_L) & \longrightarrow & K_0(\mathcal{F}(\Sigma)) \end{array}$$

Lemma 72. *This is a pushout of $\mathbb{Z}$-modules.*

Proof. A priori this need not be a pushout, since K_0 does not necessarily commute with colimits. However note that in this case we have an explicit description of the K_0 groups in terms of H^1 groups because of Theorem 44, and the result follows from the fact that we are gluing along a single chain.

More explicitly, note that $K_0(\mathcal{F}(S))$ for some marked surface S is generated by the arcs in an arc system modulo relations coming from polygons. Completing $\mathcal{A}_\gamma$ to a full arc system $\mathcal{A}_L \sqcup \mathcal{A}_\gamma \sqcup \mathcal{A}_R$ we see that since there are no polygons crossing between the two sides of the chain, so the set of relations on $K_0(\mathcal{F}(\Sigma))$ is the union of the sets of relations defining $K_0(\mathcal{F}_L)$ and $K_0(\mathcal{F}_R)$; this implies the statement above. □

By compatibility of the relative stability conditions σ_L and σ_R, the central charges on both sides agree when restricted to $K_0(\langle \mathcal{A}_\gamma \rangle)$, so we get a map $Z : K_0(\mathcal{F}(\Sigma)) \to \mathbb{C}$; this will be our central charge.

The pre-slicing

We will define full subcategories $\mathcal{P}_\phi$ of semistable objects in two steps. Let us first define initial subcategories $\mathcal{P}'_\phi$ by

$$\mathcal{P}'_\phi = j_L((\mathcal{P}_L)_\phi) \cup j_R((\mathcal{P}_R)_\phi)$$

ie. we take the images of the semistable objects under σ_L and σ_R to be stable in $\mathcal{F}(\Sigma)$.

Now let us algorithmically add some objects to the slicing by the following prescription. We first define a particular kind of arrangement of stable objects. Remember that $M_0, \ldots, M_N$ are boundary components of $\mathcal{F}(\Sigma)$ appearing in a chain of intervals that compose to γ. Let us partition $\mathcal{M} = \mathcal{M}_L \sqcup \mathcal{M}_\gamma \sqcup \mathcal{M}_R$ where $\mathcal{M}_\gamma = \{M_0, \ldots, M_N\}$, $\mathcal{M}_L$ are the other boundary components coming from Σ_L and $\mathcal{M}_R$ are the other boundary components coming from Σ_R.

Definition 39. A *lozenge* of stable intervals is the following arrangement of intervals:

- Four distinct marked boundary components $M_\ell, M_r, M_{up}, M_{down}$, where

$$M_\ell \in \mathcal{M}_L, \quad M_r \in \mathcal{M}_R, \quad M_{up}, M_{down} \in \mathcal{M}_\gamma$$

- A chain of intervals $\alpha_1, \ldots, \alpha_a$ linking M_ℓ to M_{up}, such that α_i supports a stable object $A_i \in \mathcal{P}'_{\text{phase}(A_i)}$, and
$$\text{phase}(A_1) \leq \cdots \leq \text{phase}(A_a)$$

- A chain of intervals $\beta_1, \ldots, \beta_b$ linking M_{up} to M_r, such that β_i supports a stable object $B_i \in \mathcal{P}'_{\text{phase}(B_i)}$, and
$$\text{phase}(B_1) \leq \cdots \leq \text{phase}(B_b)$$

- A chain of intervals $\delta_1, \ldots, \delta_d$ linking M_ℓ to M_{down}, such that δ_i supports a stable object $D_i \in \mathcal{P}'_{\text{phase}(D_i)}$, and
$$\text{phase}(D_1) \geq \cdots \geq \text{phase}(D_d)$$

- A chain of intervals $\eta_1, \ldots, \eta_d$ linking M_{down} to M_r, such that η_i supports a stable object $E_i \in \mathcal{P}'_{\text{phase}(E_i)}$, and
$$\text{phase}(E_1) \geq \cdots \geq \text{phase}(E_e)$$

such that the phases of these stable objects satisfy

$$\text{phase}(D_1) \leq \text{phase}(A_1) \leq \text{phase}(D_1) + 1, \quad \text{phase}(B_1) \leq \text{phase}(A_a) \leq \text{phase}(B_1) + 1,$$

$$\text{phase}(B_b) \leq \text{phase}(E_e) \leq \text{phase}(B_b) + 1, \quad \text{phase}(D_d) \leq \text{phase}(E_1) \leq \text{phase}(D_d) + 1.$$

and such that these four chain of stable intervals bound a disk. This is pictured below in Figure 5.4 for ease of presentation.

Consider now the complex number

$$Z(X) := \sum_i Z(A_i) + \sum_i Z(B_i) = \sum_i Z(D_i) + \sum Z(E_i),$$

which is the central charge of the object X supported on the interval from M_ℓ to M_r one gets by successive extensions of the A_i, B_i or D_i, E_i. The equality follows from well-definedness of Z.

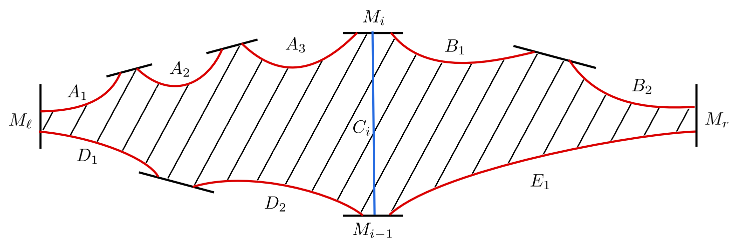

Figure 5.4: A lozenge of stable objects with $a = 3, b = 2, d = 2, e = 1$.

Definition 40. We call such a lozenge *unobstructed* if there is a choice of branch of the argument function $\arg : \mathbb{C}^\times \to \mathbb{R}$ such that the following inequalities between the phases are satisfied:

$$\text{phase}(D_1) \leq \arg(Z(X)) \leq \text{phase}(A_1), \quad \text{phase}(B_b) \leq \arg(Z(X)) \leq \text{phase}(E_e).$$

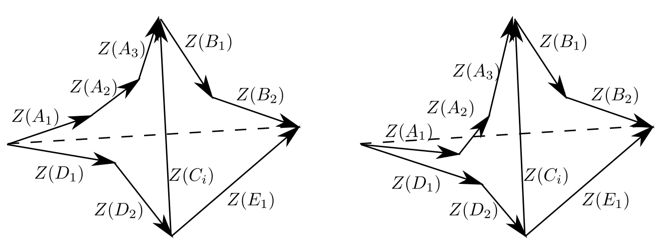

Figure 5.5: The central charges of the objects in an unobstructed lozenge (left) and in an obstructed lozenge (right).

It follows from the inequalities above that if a lozenge is unobstructed then there is only a single choice of $\arg(Z(X))$ satisfying the condition; let's call it $\phi_X \in \mathbb{R}$. For every unobstructed lozenge we find, let us declare that the corresponding X is semistable of phase ϕ_X. So we define $\mathcal{P}_\phi$ to be spanned by all objects in $\mathcal{P}'_\phi$ plus all objects of phase ϕ that we obtained from unobstructed lozenges.

Lemma 73. *The data Z and $\mathcal{P}$ as above define a prestability condition on $\mathcal{F}(\Sigma)$.*

Proof. The compatibility between the argument of Z and the phase of the subcategories $\mathcal{P}$ is automatic from the definition, since every stable object either comes directly from one side or has central charge and phase defined by the formula above. So we only have to prove that $\mathcal{P}$ is in fact a preslicing: we must show that $\text{Hom}(X, Y) = 0$ if $X \in \mathcal{P}_{\phi_X}$ and $Y \in \mathcal{P}_{\phi_Y}$ with $\phi_X > \phi_Y$.

By definition, each full subcategory $\mathcal{P}_\phi$ can be spanned by three full subcategories

$$\mathcal{P}^L_\phi = j_L((\mathcal{P}_L)_\phi), \quad \mathcal{P}^R_\phi = j_R((\mathcal{P}_R)_\phi), \quad \mathcal{P}^\lozenge_\phi,$$

where $\mathcal{P}^{\Diamond}_{\phi}$ has all the objects of phase ϕ obtained from unobstructed lozenges. Note that $\mathcal{P}^{\Diamond}_{\phi}$ is disjoint from the other two, but $\mathcal{P}^{L}_{\phi}$ and $\mathcal{P}^{R}_{\phi}$ are not disjoint; in fact their intersection is spanned by the objects supported on the chain of intervals $\{\gamma_i\}$.

Let us check vanishing of the appropriate homs. It is enough to check on stable objects. If $X, Y \in \mathcal{P}^L$ then

$$\mathrm{Hom}(X, Y) \neq 0 \implies \phi_X \leq \phi_Y$$

automatically since they're both semistable in $\mathcal{F}_L$ and $\mathcal{F}_L \to \mathcal{F}(\Sigma)$ is fully faithful; same for the case $X, Y \in \mathcal{P}^R$. So there are four remaining cases:

1. $X \in \mathcal{P}^{L}_{\phi_X}$ and $Y \in \mathcal{P}^{R}_{\phi_Y}$
2. $X \in \mathcal{P}^{\Diamond}_{\phi_X}$ and $Y \in \mathcal{P}^{L}_{\phi_Y}$
3. $X \in \mathcal{P}^{L}_{\phi_X}$ and $Y \in \mathcal{P}^{\Diamond}_{\phi_Y}$
4. $X \in \mathcal{P}^{\Diamond}_{\phi_X}$ and $Y \in \mathcal{P}^{\Diamond}_{\phi_Y}$

All the other cases can be obtained symmetrically by switching left and right. Let us treat each case separately:

1. We can find representatives of X, Y contained in the images of $\tilde{\Sigma}_L, \tilde{\Sigma}_R$ respectively, such that neither intersects the chain $\{\gamma_i\}$; so there are no intersections between them. The only way we can have $\mathrm{Hom}(X, Y) \neq 0$ is if X and Y are intervals sharing a common boundary component at one of the M_i along the chain, with a boundary path from X to Y.

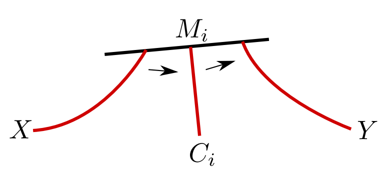

   Consider then C_i and shift its grading so that the morphism $X \to C_i$ is in degree zero; then by index arguments the morphism $C_i \to Y$ is also in degree zero. But since these three objects are stable we have

$$\phi_X \leq \phi_{C_i} \leq \phi(Y).$$

2. Let X be obtained from an unobstructed lozenge with notation as in Definition 39, and $Y \in \mathcal{F}_L$. Consider the distinguished triangle $B \to X \to A$ and let us apply the functor $\mathrm{Hom}(-, Y)$ to get a distinguished triangle

$$\mathrm{Hom}(A, Y) \to \mathrm{Hom}(X, Y) \to \mathrm{Hom}(B, Y).$$

Since Y comes from $\mathcal{F}_L$, it has a representative that stays to the left of the chain and therefore of B so by assumption we have $\mathrm{Hom}(B,Y) = 0$. Thus if $\mathrm{Hom}(X,Y) \neq 0$ then $\mathrm{Hom}(A,Y) \neq 0$. Since A is given by the iterated extension of the A_i, there must be some A_i with $\mathrm{Hom}(A_i,Y) \neq 0$; but A_i and Y are both in the image of $\mathcal{F}_L$ we must have $\phi_{A_i} \leq \phi_Y$, and also by construction $\phi_X \leq \phi_{A_1}$ so we have

$$\phi_X \leq \phi_{A_1} \leq \phi_{A_i} \leq \phi_Y.$$

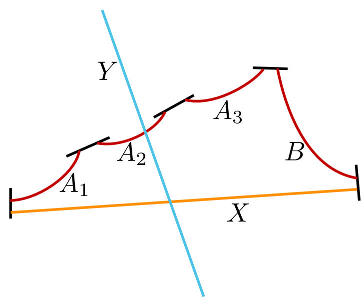

3. Suppose we have an unobstructed lozenge with sides A, B, D, E and diagonal Y. A similar argument as in case (2) shows that if $\mathrm{Hom}(X,Y) \neq 0$, then $\mathrm{Hom}(X,D) \neq 0$, and then for some i we have $\mathrm{Hom}(X,D_i) \neq 0$

$$\phi_X \leq \phi_{D_i} \leq \phi_{D_d} \leq \phi_Y.$$

4. This case can be obtained by an iterated version of the argument in case (2). Let us denote the two lozenges by A_X, B_X, D_X, E_X with diagonal X and A_Y, B_Y, D_Y, E_Y with diagonal Y. Suppose that $\mathrm{Hom}(X,Y) \neq 0$, and consider the triangle $D_X \to X \to E_X$. Consider first the case $\mathrm{Hom}(D_X,Y) = 0$ then $\mathrm{Hom}(E_X,Y) \neq 0$. Now consider the triangle $B_Y \to Y \to A_Y$. Since E_X and A_Y have representatives contained in the right and the left side, respectively, and don't share a boundary component we have $\mathrm{Hom}(E_X, A_Y) = 0$ so we must have $\mathrm{Hom}(E_X, B_Y) \neq 0$. But then there must be indices i, j such that $\mathrm{Hom}((E_X)_i, (B_Y)_j) \neq 0$ so then

$$\phi_X \leq \phi_{(E_X)_i} \leq \phi_{(B_Y)_j} \leq \phi_Y.$$

The other case is $\mathrm{Hom}(D_X,Y) \neq 0$. Consider the triangle $B_Y \to Y \to A_Y$. By an analogous argument we can find indices i, j such that

$$\phi_X \leq \phi_{(A_X)_i} \leq \phi_{(D_Y)_j} \leq \phi_Y.$$

□

Uniqueness of compatibility structure

In the same setting as the previous subsection, let $\Gamma \subset \mathrm{RelStab}(\Sigma_L, \gamma) \times \mathrm{RelStab}(\Sigma_R, \gamma)$ be the locus of pairs of relative stability conditions (σ_L, σ_R) such that there exists a compatibility condition between σ_L and σ_R.

Lemma 74. *For each $(\sigma_L, \sigma_R) \in \Gamma$, there is a unique compatibility structure between σ_L and σ_R up to equivalence.*

Proof. Let us first prove that the numbers m, n defining $\tilde{\Sigma}_L, \tilde{\Sigma}_R$ are unique. Consider the subset

$$\mathbb{M}_\sigma \subset \mathbb{S} = \bigsqcup_{n \geq 2} \mathrm{Stab}(\mathcal{F}(\Sigma_L \cup_\gamma \Delta_n))$$

of its minimal (but possibly not reduced) representatives. Given $\tilde{\sigma} \in \mathbb{M}_\sigma$ we consider the COSI decomposition of the rank one object C supported on γ as before, and define the numbers $int(\tilde{\sigma}), ext(\tilde{\sigma})$ to be respectively the number of internal/external extensions in the γ chain, ie. the number of indices i such that the corresponding extension happens on the left/right, ie. by an extension map $\in \mathrm{Ext}^1(C_{i+1}, C_i)/\in \mathrm{Ext}^1(C_i, C_{i+1})$. This defines constructible functions $int, ext : \mathbb{M}_\sigma \to \mathbb{Z}_{\geq 0}$ such that $int(\tilde{\sigma}) + ext(\tilde{\sigma}) = N - 1$, where $N - 1$ is the total length of the object C under $\tilde{\sigma}$.

We argue that the function int is constant; by Lemma 67 there is a unique minimal and reduced representative σ^{red} of every relative stability condition. However, reduced restriction does not change the int of a stability condition, so $int(\tilde{\sigma}) = int(\tilde{\sigma}^{\mathrm{red}}) = int(\sigma^{\mathrm{red}})$ on all of $\mathbb{M}_\sigma$. We define the same functions on the right side for the relative stability condition $\sigma' \in \mathrm{RelStab}(\Sigma_R, \gamma)$. Compatibility implies that $int(\tilde{\sigma}) = ext(\tilde{\sigma}'), ext(\tilde{\sigma}) = int(\tilde{\sigma}')$, but since int is constant there is only one possibility for the value of ext. Comparing with the gluing map we have $m = ext(\tilde{\sigma}), n = ext(\tilde{\sigma}')$.

This determines the isomorphism type of the surfaces $\tilde{\Sigma}_L$ and $\tilde{\Sigma}_R$. Consider now the inclusion of marked surfaces $j_L : \tilde{\Sigma}_L \hookrightarrow \Sigma_L \cup_\gamma \Sigma_R$. By definition of compatibility structure, $j_L|_{\Sigma_L}$ agrees with the inclusion $\Sigma_L \hookrightarrow \Sigma_L \cup_\gamma \Sigma_R$, so the 'left part' of j_L is fixed; j_L is determined up to equivalence by the images of the extra $m - 2$ marked boundary intervals in the disk Δ_m attached along γ (two of the marked boundary intervals are fixed to the ends of γ).

Analogously, j_R is determined up to equivalence by the image of the extra $n - 2$ marked boundary intervals of Δ_n. But the images of the extra $m - 2$ marked intervals under j_L is contained in the image of the marked intervals coming from Σ_R under j_R, so they are fixed; the same is true for the image of the extra $n - 2$ marked intervals under j_R. Minimality implies that the subcategory $\langle \mathcal{A}_L \sqcup \mathcal{A}_\gamma \rangle$ is the whole category $\mathcal{F}(\tilde{\Sigma}_R)$ so once we fix σ, the representative $\tilde{\sigma}$ is completely determined by its restriction to $\langle \mathcal{A}_\gamma \rangle \cong \mathcal{F}(\Delta_{N+1})$.

By the classification of stability conditions on the Fukaya category of a disk presented in [56, Section 6.2], stability conditions on $\mathcal{F}(\Delta_{N+1})$ are entirely determined by the central charges and phases of the $N + 1$ intervals in the chain. Let us label the marked boundary intervals $M_0, \ldots, M_N$ in sequence. We argue that the central charges and phases of the objects $C_1, \dots, C_N$ are unique using the following 'zip-up' procedure. Consider first the object C_1; since M_0 is in the common image of Σ_L and Σ_R, and M_1 is 'internal' (in the subset counted by the *int* function) to either of those surfaces, the interval supporting C_1 is contained in the image of either Σ_L or Σ_R, so its central charge $Z(C_1)$ and phase ϕ_1 are fixed by either σ_L or σ_R.

Suppose without loss of generality that the interval supporting C_1 is in the image of Σ_L, and consider now C_2. There are two possibilities for M_2; either it is internal to Σ_L or to Σ_R. In the former case since M_1 and M_2 are in the image of the same side Σ_L, $Z(C_2)$ and ϕ_2 are fixed by σ_L. In the latter case, C_2 is not in the image of either Σ_L or Σ_R, but we consider the concatenation C_{1+2} given by extending at M_1; both ends of this object are in the image of Σ_R so the central charge $Z(C_{1+2})$ of this (non-stable) object is fixed by σ_R. So $Z(C_2) = Z(C_{1+2}) - Z(C_1)$ is also fixed. Moreover, among the shifts of C_2, there is a unique one with the extension map at M_1 in the correct degree, so ϕ_2 is also fixed. Proceeding by induction we find that all $Z(C_i), \phi_i$ are fixed by the initial data σ_L, σ_R. □

Relation between cutting and gluing maps

Because of the uniqueness of compatibility structure proven above and Lemma 73, we can define a gluing map

$$\mathrm{RelStab}(\Sigma_L, \gamma) \times \mathrm{RelStab}(\Sigma_R, \gamma) \supset \Gamma \xrightarrow{glue_\gamma} \mathrm{PreStab}(\mathcal{F}(\Sigma_L \cup_\gamma \Sigma_R))$$

which produces a prestability condition.

A priori it is not obvious whether these are actual stability conditions, however this can be shown to be the case when we start with an actual stability condition $\sigma \in \mathcal{F}(\Sigma)$.

Theorem 75. *The composition*

$$\mathrm{Stab}(\mathcal{F}(\Sigma)) \xrightarrow{cut_\gamma} \Gamma \xrightarrow{glue_\gamma} \mathrm{PreStab}(\mathcal{F}(\Sigma))$$

is equal to the canonical inclusion $\mathrm{Stab}(\mathcal{F}(\Sigma)) \hookrightarrow \mathrm{PreStab}(\mathcal{F}(\Sigma))$.

Note that the theorem can be also stated as saying that the gluing map lands in $\mathrm{Stab}(\mathcal{F}(\Sigma))$ and gives an right-inverse to the cutting map. It is then immediate from the definitions that this is also a left-inverse; the cutting map forgets all the stable objects coming from the lozenges so the composition

$$\Gamma \xrightarrow{glue_\gamma} \mathrm{Stab}(\mathcal{F}(\Sigma)) \xrightarrow{cut_\gamma} \Gamma$$

is the identity on pairs of compatible relative stability conditions.

We will need the following lemma in the proof of 75:

Lemma 76. *Let X be a stable interval object (under σ), with a representative that crosses the interval γ once. Then there is an unobstructed lozenge (under σ_L, σ_R) with diagonal X. Conversely, the diagonal of every unobstructed lozenge is stable under σ.*

Proof. Let $C_1, \ldots, C_N$ be the COSI decomposition of the object C supported on γ. Note that X cannot cross this chain multiple times, since this would create a polygon of the sort prohibited by Lemma 50. Let us say then that X intersects one C_j transversely. Then we have $\mathrm{Ext}^1(C_j, X) \cong \mathrm{Hom}(X, C_j) \cong k$; consider the corresponding extension and cone

$$C_j \to A \oplus E \to X, \qquad B \oplus D \to X \to C_j.$$

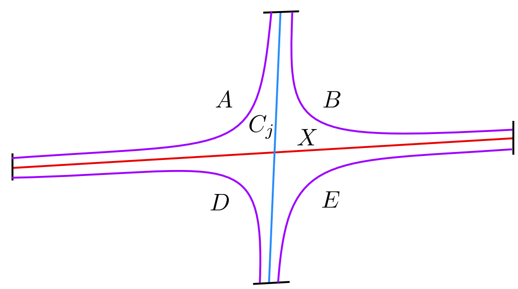

Each one of the objects A, B, D, E is an embedded interval object and by Proposition 53 has a COSI decomposition; we denote the objects in these chains by $\{A_i\}, \{B_i\}, \{D_i\}, \{E_i\}$, respectively.

We argue that $\{A_i\}$ and $\{B_i\}$ only have extensions on the right, and $\{D_i\}, \{E_i\}$ only have extensions on the left. Note first that the chains of intervals $\{A_i\}$, $\{D_i\}$ and the interval γ don't intersect mutually, since this would contradict Lemma 50. Consider the chain made up of supports of the A_i and $D_i[-1]$. This chain together with γ bounds a disk, therefore every extension is on the right; this translates to extensions on the right $\in \mathrm{Ext}^1(A_i, A_{i+1})$ and extensions on the left $\mathrm{Ext}^1(D_{i+1}, D_i)$. An analogous argument applies to B and E; note that since none of these chains crosses γ, and γ separates Σ, they do not intersect one another.

Thus we have a lozenge whose diagonal is X; it remains to prove it is unobstructed. Suppose that the lozenge A, B, D, E is obstructed; therefore we must have at least one of the following inequalities

$$\phi_{A_1} \leq \phi_X, \quad \phi_{D_1} \geq \phi_X, \quad \phi_{B_b} \geq \phi_X, \quad \phi_{E_e} \leq \phi_X.$$

Suppose first that $\phi_{A_1} < \phi_X$. Consider then the object X' given by the iterated extension of $A_2, \ldots A_a, B_1, \ldots B_b$, we then have a distinguished triangle

$$X' \to X \to A_1$$

and the map $X \to A_1$ cannot be zero since X' is indecomposable (by Theorem 43), which cannot happen since $\phi_X > \phi_{A_1}$. The other cases are similar; moreover, the case of coinciding phases poses no further problems since we can always take σ to be appropriately generic (since we need to be off of finitely many walls).

This proves one of the directions. For the converse, suppose that we have an unobstructed lozenge A, B, D, E as above, with diagonal object X which is not stable. By construction X is an embedded interval, so it has a chain-of-interval decomposition $\{X_i\}$ under σ. There are two mutually exclusive cases:

1. There are representatives for all the X_i contained in the lozenge, ie. contained in the disk bounded by the lozenge or running along its sides.

2. At least one of the representatives necessarily crosses out of the lozenge.

The concatenation of the chain $\{X_i\}$ is isotopic to the object X. Therefore in case (2), if the chain crosses out of the lozenge along one of the sides it must cross back in, and along the *same* side, since each of the objects A, B, D, E cuts the surface into two. Therefore we have a configuration prohibited by Lemma 50.

As for case (1), every extension between X_i and X_{i+1} must happen at one of the marked components along the boundary of the lozenge. Note that even though the chain $\{X_i\}$ may not be simple (intervals could in principle double back), it must not cross itself by the same lemma, and therefore there are only two options: either X_i and X_{i+1} share a boundary component along the top of the lozenge (ie. along A or B sides) and the extension happens on the right, or it is along the bottom (ie. along D or E sides) and the extension happens on the left. Suppose that at least one of the intervals X_i ends on the A side; let i be maximal among such indices. Then X_{i+1} stretches between the A side and another side of the lozenge, however its phase is smaller than X_i so this contradicts the existence of a nontrivial extension on the right $\in \mathrm{Ext}^1(X_i, X_{i+1})$. The same argument can be applied along any of the other sides, in the case where no interval ends on the A side. Therefore there cannot be more than one stable interval, and X itself is stable. □

The lemma above should be interpreted as stating that the unobstructed lozenges "see" all the stable interval objects that were eliminated by cutting along γ.

Proof. (of Theorem 75) For clarity let us denote $\sigma = (Z, \mathcal{P}) \in \mathrm{Stab}(\mathcal{F}(\Sigma))$, $(\sigma_L, \sigma_R) = ((Z_L, \mathcal{P}_L), (Z_R, \mathcal{P}_R))$ for its image under the cutting map, and $\sigma_g = (Z_g, \mathcal{P}_g)$ for the pre-stability condition glued out of σ_L and σ_R. It is clear that the central charges Z and Z_g are the same; it is enough to check on a set of generators and we can pick the arc system $\mathcal{A}_L \sqcup \mathcal{A}_\gamma \sqcup \mathcal{A}_R$ where the central charges agree by construction.

As for the (pre)slicings, the inclusion $\mathcal{P}_g \subseteq \mathcal{P}$ is a direct consequence of Lemma 76, since every diagonal of an unobstructed lozenge under σ_L, σ_R is stable under σ. As for the inclusion $\mathcal{P} \subseteq \mathcal{P}_g$, by Theorem 47 every stable object is either a stable embedded interval or a stable embedded circle; again by Lemma 76 the stable embedded intervals correspond to unobstructed lozenges and appear in $\mathcal{P}_g$, and as for the stable circles, they must not cross the chain $\{\gamma_i\}$ by Lemma 50 so they are either entirely contained in $\mathcal{F}_L$ or $\mathcal{F}_R$ and therefore also appear in $\mathcal{P}_g$. So $\mathcal{P}_g$ is in fact a slicing and equal to $\mathcal{P}$. □

5.3 Calculations

In the previous section, we outlined a procedure for cutting stability conditions on $\mathcal{F}(\Sigma)$ along some embedded interval γ into relative stability conditions. This procedure only works when the object supported on γ has a simple COSI decomposition, and from Lemma 52 we know that embedded intervals cutting the surface into two necessarily have this property.

Consider some general surface Σ with genus g and punctures $p_0, p_1, \ldots, p_n$ with $m_0, m_1, \ldots, m_n$ marked boundaries, respectively. Suppose that $m_0 \geq 2$. We can then decompose the surface into a disk with some number of marked boundary intervals, possibly some annuli with two marked boundary intervals on the outer boundary circle, and possibly some punctured tori with two marked boundary intervals on the boundary circle.

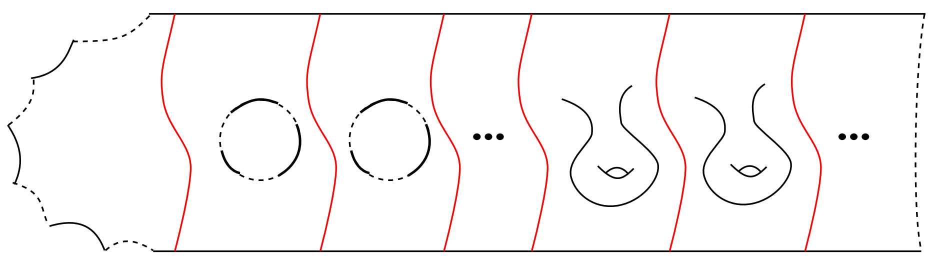

Figure 5.6: A decomposition of the surface Σ into a disk, possibly several annuli and possibly several punctured tori.

Note that for each one of these pieces, when modified by gluing some disk Δ_n along a boundary, give rise to the following kinds of surfaces:

1. The disk Δ_n with $n \geq 2$ marked boundary intervals
2. The annulus $\Delta^*_{p,q}$ with p, q marked boundary intervals on the outer and inner boundary circle, respectively
3. The punctured torus T^*_n with n marked boundary intervals

on which we need to calculate the space of (ordinary) stability conditions.

By the main theorem of [56] (Theorem 5.3) the locus of HKK stability conditions in $\mathrm{Stab}(\mathcal{F}(\Sigma))$ is a union of connected components. Thus, if every stability condition can be continuously deformed into an HKK stability condition, then all stability conditions are HKK stability conditions.

We will use this strategy for the three base cases; in fact we will prove that every stability condition can be continuously deformed to a stability condition with finite heart. This argument already appears for the case of the disk and the annulus in [56]; we will reproduce it in greater detail so that its use in the context of the punctured torus is clearer.

Finite-heart stability conditions

The definitions and lemmas here seem to be standard in the literature to some extent and may appear with different formulations; for clarity we will assemble them here.

Definition 41. A stability condition $\sigma \in \mathrm{Stab}(\mathcal{D})$ is *finite-heart* if the corresponding heart $\mathcal{H}$ is a finite abelian category, ie. a finite length abelian category that furthermore only has finitely many isomorphism classes of simple objects.

Note that finite-length only means that every object is finite-length but those lengths could be unbounded; this doesn't happen in the cases we care about because of the following standard fact.

Lemma 77. *If $\mathcal{H}$ is finite-length and $\mathrm{rk}(K_0(\mathcal{H})) = \mathrm{rk}(K_0(\mathcal{D})) < \infty$ then $\mathcal{H}$ is finite, and in particular the number of isomorphism classes of simple objects is equal to $\mathrm{rk}(K_0(\mathcal{D}))$.*

We have the following criterion to determine when some stability condition is finite-heart, based on the set of stable phases $\Phi \in S^1$, ie. the set of phases of stable objects.

Lemma 78. *If Φ has a gap around zero (ie. $S^1 \setminus \Phi$ contains an open interval $I \ni 0$) and $K_0(\mathcal{D}) < \infty$ then σ is finite-heart.*

Remark. This fact is used in [56] but left unstated. The clear statement and proof of this lemma were informed to me by F. Haiden.

Proof. Note that ϕ is symmetric under a $\mathbb{Z}_2$ rotation so having a gap around zero means that Φ is contained in a strict cone in the upper half-plane. Thus there is $K > 0$ such that $|\Im(Z(X))| > K \cdot |\Re(Z(X))|$ for any semistable object X. We will argue that the set of semistable imaginary parts

$$\{\Im(Z(E))|0 \neq E \in \mathcal{P}_\phi, \phi \in \mathbb{R}\}$$

is discrete. Suppose that there is an accumulation point, which without loss of generality we assume to be $a > 0$; we can then pick a sequence of pairwise non-isomorphic semistable objects $\{E_n\}$ such that $\lim_{n\to\infty} |\Im(Z(E_n)) - a| = 0$; in particular for $\delta > 0$ we can pick the sequence such that $|\Im(Z(E_n)) - a| < \delta$ for every n, so picking $0 < \delta < a$ gives $|\Re(Z(E_n))| < K(a + \delta)$

But since Λ is finite rank and the E_n are all distinct, we have $\lim_{n\to\infty} \|E_n\| = \infty$. We then have

$$|Z(E_n)| < |\Im(Z(E_n))| + |\Re(Z(E_n))| \leq (K + 1)|\Re(Z(E_n))| \leq (K + 1)K(a + \delta) = \text{const.}$$

So we have $\lim_{n\to\infty} \frac{|Z(E_n)|}{\|E_n\|} = 0$ contradicting the support condition.

So since the set of imaginary parts of objects in the heart $\mathcal{H}$ is discrete and bounded below by zero, any strictly descending chain of objects is finite, and therefore $\mathcal{H}$ is finite-length, and thus σ is finite-heart by the assumption $\mathrm{rk}(K_0(\mathcal{D})) < \infty$. ☐

Using the formalism of S-graphs presented in Section 6 of [56], one can prove the following lemma (which is implicitly used in the proofs of Theorems 6.1 and 6.2 of that same paper)

Lemma 79. *If σ is a finite-heart stability condition on $\mathcal{F}(\Sigma)$ then it is an HKK stability condition.*

For each of the three base cases, we will see that every stability condition can be deformed to a finite-heart stability condition.

The disk

(Section 6.2 of [56]) We have $\mathcal{F}(\Delta_n) \cong \mathrm{Mod}(\mathbb{A}_{n-1})$, which up to shift has finitely many indecomposable objects. Thus any heart is a finite abelian category, and every stability condition is finite-heart and therefore HKK.

The annulus

There are two different kinds of annulus; one where the nontrivial circle is gradable, ie. has index zero, and one where it has index nonzero. Consider first the annulus $\Delta^*_{p,q,(m)}$ with p, q marked boundary components and grading $m \neq 0$ around the circle.

We argue that the set of stable phases is finite. Let us fix some embedded interval object I_0 to have winding number zero, and measure the winding number of every other interval or circle with reference to it. By the classification of objects, there are only finitely many primitive (ie. non multiple) classes in $K_0(\mathcal{F}(\Delta^*_{p,q,(m)}))$ whose winding number is less than some fixed N in absolute value, so if there are infinitely many non-isomorphic stable objects there must be a sequence of stable objects X_i with winding number $\to \infty$.

Consider some object X_i with winding number N_i which intersects I_0 transversely N_i many times. Since the circle has index $m \neq 0$, this contributes classes to both $\mathrm{Ext}^*(I_0, X_i)$ and $\mathrm{Ext}^*(X_i, I_0)$ in a range spanning $(m-1)N_i$ degrees. But this is impossible as $N_i \to \infty$ since the stable components of I_0 have a minimum and maximum phase.

Consider now the annulus with zero grading. We have $\mathcal{F}(\Delta^*_{p,q,(0)}) \cong \mathrm{Mod}(\tilde{\mathbb{A}}_{p+q-1})$. So we have $\Gamma = K_0(\mathcal{F}(\Delta^*_{p,q,(0)})) = \mathbb{Z}^{p+q}$, and denote by $S \subset \Gamma$ the subgroup generated by the circle around the annulus. Let $E \subset \Gamma$ be the set of classes of indecomposable objects. By the classification of objects E/S is finite so the only possible accumulation point in the set of stable phases Φ is $\arg(Z(S))$. After a rotation (which can be arbitrarily small) we can guarantee that Φ has a gap around zero and apply Lemma 78.

The punctured torus

The calculation of this case is new. From the cutting procedure we know that only need to consider the punctured torus T^*_n with $n \geq 2$ marked boundary components. In fact there are many inequivalent such punctured tori, with different gradings. Let us pick simple closed curves L and M as longitude and meridian, and denote by i_L, i_M the index of the grading along them. By picking different curves we get indices differing by an action of $\mathrm{SL}(2,\mathbb{Z})$ so the set of distinct graded punctured tori is $\mathbb{Z}^2/\mathrm{SL}(2,\mathbb{Z})$. The orbits of $\mathrm{SL}(2,\mathbb{Z})$ on $\mathbb{Z}^2$ are labelled by gcd, so each orbit contains a unique pair of the form $(0, m)$.

Let us fix a grading such that $(i_L, i_M) = (0, m)$. It will be important for us to know what are the circle objects. The classes in $\pi_1(T^*)$ which are representable by simple closed curves are the curves winding (p, q) times around the longitude and meridian, with $\gcd(p, q) = 1$, plus the curve $MLM^{-1}L^{-1}$, ie. the circle around the puncture.

For any of these tori, the index of the circle around the puncture is always 2 for topological reasons (it bounds a punctured torus) so this curve is never gradable. On the torus with $(i_L, i_M) = (0, m \neq 0)$ torus the index of the (p, q) curve is $mq \neq 0$ if $q \neq 0$, so all of the embedded circle objects are supported on the longitude L. On the torus with $(i_L, i_M) = (0, 0)$, every simple closed curve is gradable and supports embedded circle objects.

Remark. This is the fundamental reason why the calculation for the $(0, 0)$ will be more involved than the case of the annulus; in that case the lattice spanned by the circle objects inside of $K_0(\mathcal{D})$ is rank one, so there can be at most one direction of phase accumulation. In the punctured torus, the central charges of stable objects could in principle occupy every direction of the lattice, making Φ dense; we will prove that this doesn't happen generically.

The $(0, m \neq 0)$ torus

Let us denote $\mathcal{D} = \mathcal{F}(T^*_{n,(0,m)})$ where n is the number of marked boundaries. This case will be very similar to the index zero annulus. There is only one type of embedded circle object L, since no other circles are gradable. Let $\Gamma = K_0(\mathcal{D})$ and $E \subset \Gamma$ be the set of classes of stable objects.

We argue that the set $E/\langle L\rangle$ is finite. Suppose otherwise, and note that by the classification of embedded curves, the number of embedded curves with winding numbers (p, q) with $|q \leq N|$ is infinite, but they form finitely many orbits in $K_0(\mathcal{D})$ under the action of the subgroup $\langle L\rangle$. Thus, if we have an infinite sequence of stable objects $\{E_i\}$ with winding numbers (p_i, q_i) and pairwise distinct classes $[E_i] \in K_0(\mathcal{D})/\langle L\rangle$, there is a subsequence with $\lim_{i\to\infty} |q_i| = \infty$.

This is impossible in any stability condition. Note that an object with winding q_i along the meridian intersects L transversely $|q_i|$ times; but since $m \neq 0$ the difference in degree between each two consecutive intersections is $|m|$, so the amplitude of nonzero degrees in both $\mathrm{Hom}(E_i, L)$ and $\mathrm{Hom}(L, E_i)$ is $m(q_i - 1)$. Since $|q_i| \to \infty$ we can find stable objects E_i with arbitrarily large amplitude morphisms in both directions which is impossible since L has some HN decomposition with finitely many semistable factors, having a minimum and a maximum phase.

From the fact that $E/\langle L\rangle$ is finite we can proceed as in the annulus case, and after an infinitesimal rotation we can guarantee that any stability condition has an gap in Φ.

The $(0,0)$ torus

Let us denote $\mathcal{D} = \mathcal{F}(T^*_{n,(0,0)})$, where n is the number of marked boundaries. We will first need some facts about $K_0(\mathcal{D})$. By Theorem 5.1 of [56] there is an isomorphism

$$K_0(\mathcal{F}(\Sigma, M)) = H_1(\Sigma, M; \mathbb{Z}_\tau),$$

where $\mathbb{Z}_\tau$ is the $\mathbb{Z}$-local system associated to the orientation double cover of the foliation. In our case, since we are looking at the foliation with $(0,0)$ winding, $\mathbb{Z}_\tau$ is trivial.

Let us pick an explicit set of generators of $K_0(\mathcal{D})$ as below: first choose a basis of $H_1(T, \mathbb{Z})$ and a labeling $M_1, \ldots, M_N$ of the marked boundary components. The classes $[L]$ and $[M]$ are represented by circles around the longitude and meridian, and $[E_i], i = 1, \ldots, N$ is represented by intervals that connect adjacent M_i and M_{i+1} along the boundary. Consider the object X winding around the longitude with ends at M_1, M_N. Extending it by $E_1, \ldots, E_{n-1}$ and by E_n both give L, so in K_0 we have $\sum_{i=1}^n [E_i] = 0$

So the classes $[L], [M], [E_1], \ldots, [E_{n-1}]$ give a basis of $K_0(\mathcal{D})$. Since every immersed curve has well-defined winding numbers, we have a projection map

$$w : K_0(\mathcal{D}) \longrightarrow \mathbb{Z}^2 = Span([L], [M])$$

taking a curve of (p, q) winding numbers to $p[L] + q[M]$. The following lemma tells us that the distribution of stable phases is not essentially changed by w.

Lemma 80. *For any sequence of stable objects $\{X_k\}$ (with all X_k pairwise distinct) if $\lim_{k\to\infty} \arg(Z(X_k))$ exists then*

$$\lim_{k\to\infty} \arg(Z(w([X_k]))) = \lim_{k\to\infty} \arg(Z(X_k))$$

Proof. By the classification of indecomposables, X is represented by some circle or interval with winding (p, q). If X is a circle we already have $w([X]) = [X]$. Given embedded interval with boundaries on M_1, M_i, one can express it as the concatenation of p copies of the interval winding along the longitude with both ends at M_1 (whose class is $[L]$), q copies of the interval winding along the meridian with both ends at M_1 (whose class is $[M]$) and a chain of intervals $E_1, \ldots, E_{i-1}$ connecting M_1 to M_i. This chain can wind around the circle any number of times, but since $\sum_{j=1}^n [E_j] = 0$, its class is always $[E_1] + \cdots + [E_i]$. Applying $|Z(\cdot)|$, since this sum is bounded above we have

$$|Z(X) - Z(w(X))| \leq C$$

for some fixed constant C.

Consider now the stable objects X_k. Without loss of generality suppose that $\lim_{k\to\infty} \arg(Z(X_k)) = 0$ (ie. the positive real direction). These objects can be represented by embedded intervals; note that there are finitely many embedded intervals with fixed winding numbers. Thus in the infinite sequence of distinct objects $\{X_k\}$ we must have $p_k^2 + q_k^2 \to \infty$ so therefore $|Z(X_k)| \to \infty$ and $\Re(Z(X_k)) \to +\infty$.

The triangle inequality,

$$|Z(X_k)| - C \leq |Z(w(X_k))| \leq |Z(X_k)| + C$$

also implies similar inequalities for the real and imaginary parts. Since $|\Re(Z(X_k))| \to \infty$ we have

$$\lim_{k\to\infty} \frac{|\Im(Z(w(X_k)))|}{|\Re(Z(w(X_k)))|} \leq \lim_{k\to\infty} \frac{|\Im(Z(X_k))| + C}{|\Re(Z(X_k))| - C} = \lim_{k\to\infty} \frac{|\Im(Z(X_k))|}{|\Re(Z(X_k))|} = 0.$$

so $\lim_{k\to\infty} \arg(Z(w(X_k))) = \lim_{k\to\infty} \arg(Z(X_k)) = 0$. □

Corollary 81. *If the set of stable phases Φ is dense in S^1 then the set*

$$\Phi_w = \{\arg(Z(w(X))) \mid X \text{stable}\}$$

is also dense in S^1

We will also need to know a bit more about which objects necessarily intersect transversely.

Lemma 82. *Consider two embedded objects X and Y with winding numbers (p_X, q_X) and (p_Y, q_Y), respectively. If $|p_X q_Y - q_X p_Y| \geq 2$ then X and Y intersect transversely.*

Proof. If X is a circle with $(p_X, q_X) = (1, 0)$, then any Y with $|q_Y| \geq 1$ intersects X transversely; if X is an embedded interval with $(p_X, q_X) = (1, 0)$, then circles with $|q_Y| \geq 1$ intersect X transversely but intervals with $|q_Y| = 1$ may not. On the other hand, winding more times around the meridian by requiring $|q_Y| \geq 2$ necessarily causes a transverse intersection. Applying the right element of $\mathrm{SL}(2,\mathbb{Z})$ that sends $(1, 0) \mapsto (p_X, q_X)$ gives the statement of the lemma. □

The following lemma gives an existence result for a certain kind of stable object.

Lemma 83. *Let $\sigma \in \mathrm{Stab}(\mathcal{D})$ be a stability condition on $\mathcal{D} = \mathcal{F}(T^*_N)$. Then there is some stable object represented by an embedded interval with nonzero winding and ends at different marked boundaries.*

Proof. Suppose otherwise; by the classification of embedded curves, there are three remaining possibilities for a stable object:

1. A semistable circle with winding $\neq (0, 0)$,
2. A semistable interval with winding $\neq (0, 0)$ both ends on the same marked boundary,
3. A semistable interval with $(0, 0)$ winding and ends possibly on different marked boundaries.

Two objects of type (2) ending on the same marked boundary M will have extension morphisms between them, but we argue that if they have different classes in $K_0(\mathcal{D})$ these morphisms cannot appear in the HN decomposition of any object. By keeping track of the grading with respect to the $(0,0)$ grading on the torus, we note that if we grade the intervals such that $\deg(f) = 1$, then $\deg(g) = 0$. Thus $\phi_B \leq \phi_A$ and by genericity $\phi_B \neq \phi_A$ since $[A] \neq [B]$, so $f \in \mathrm{Ext}^1(A, B)$ cannot appear in the HN decomposition.

Thus every interval with winding (p, q), $\gcd(p, q) = 1$ and ends on the same marked boundary must be semistable, since there is no way to express it as a valid extension of the objects above. We argue that this is impossible in a generic stability condition. Take for example the semistable interval J with winding $(1, 0)$ and both ends on some marked boundary M, and consider another embedded interval J' with winding $(0, 1)$, with ends on M and $M' \neq M$. By assumption, J' is not semistable so it must have a chain-of-intervals

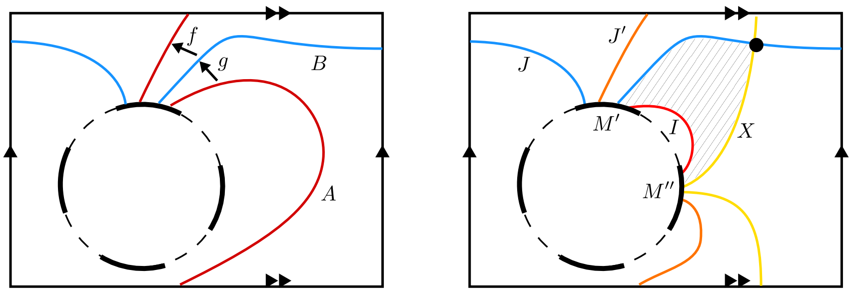

Figure 5.7: Left: two stable objects A, B of type (2). Right: one stable object J of type (2) and one (not semistable) embedded interval J', in whose decomposition some object I of type (3) must appear, causing a prohibited polygon (shaded) to appear.

decomposition with at least two distinct phases; consider the interval objects in this chain that end at M; since the other end of the chain is at another marked boundary, among these objects there must be at least one semistable interval J'_0 of type (3) above (ie. with zero winding). We see immediately that such an interval has an essential transversal intersection with J; therefore the rest of the chain (after J'_0) must cross J as well. But this configuration is prohibited by Lemma 50.

So there must be some semistable interval object I' with nonzero winding and ends on different marked boundary intervals. If I is not stable, consider its Jordan-Hölder filtration into stable objects; among these there must be one stable interval object I connecting two distinct marked boundaries. □

Using the lemmas above, in the following calculation we show that an adequately generic stability condition does not have dense phases in S^1.

Lemma 84. *Let $\sigma \in \mathrm{Stab}(\mathcal{D})$ be a stability condition on $\mathcal{D} = \mathcal{F}(T^*_N)$. Then possibly after a infinitesimal deformation the set of stable phases Φ has a gap, ie. $S^1 \setminus \Phi$ contains an open interval.*

Proof. By the previous lemma, there must be some stable interval I with nontrivial winding and ends on distinct marked boundary components. Applying an appropriate $\mathrm{SL}(2,\mathbb{Z})$ automorphism, we can assume this stable interval I has winding numbers $(1,0)$, ie. winds around the longitude once. Let L be the rank one trivial circle object also with winding number $(1,0)$.

The subset of $\mathrm{Stab}(\mathcal{D})$ where I is stable is open by standard results [24] so there is a neighborhood U of σ where I is stable. From the description of $K_0(\mathcal{D})$ we know that $[I] \neq [L]$, so $Z(I), Z(L)$ are not parallel in the complement of a codimension one wall. Thus, possibly after an infinitesimal deformation inside of U, we can guarantee that I is stable and $Z(I), Z(L)$ have different arguments.

Consider the trivial rank one objects L and M (which may or may not be stable) supported along the longitude and meridian, with gradings so that

$$\deg(M, I) = \deg(M, L) = 0$$

and for simplicity let us rotate and scale the stability condition so that $Z(L) = 1$. Since $[L] \neq [M]$ and we fixed $Z(L) \in \mathbb{R}$, for a generic stability condition we must have $Z(M) \notin \mathbb{R}$. Let us treat the case $\Im(Z(M)) > 0$ first; the other case follows from an analogous argument.

Suppose now that Φ is dense in S^1; by Lemma 81 ,Φ_w is dense too. For a choice of winding numbers (p,q), let us denote by

$$\mathcal{X}_{p,q} = \{(p', q') \mid q > 0, |pq' - qp'| \geq 2\} \subset \mathbb{Z}^2$$

the set of winding numbers whose objects necessarily intersect transversely with objects of winding number (p,q), with positive winding around the meridian.

The set $\mathcal{X}_{1,0}$ corresponding to I is given by $q \geq 2$; so at infinity $\mathcal{X}_{1,0}$ approaches a sector (with angle π). Remember that for any N there are only finitely many indecomposable objects with winding satisfying $p^2 + q^2 \leq N$. By density of Φ_w we can find some stable object X_0 with winding numbers $(p_0, q_0) \in \mathcal{X}_{1,0}$.

Consider now the set $\mathcal{X}_{1,0} \cap \mathcal{X}_{p_0,q_0}$; this set is composed of lattice points inside of two components of a subset of $\mathbb{R} \times \mathbb{R}_+$. At infinity, the right component approaches a sector with angle spanning $(0, \arctan(q_0/p_0))$ and the left component approaches a sector at $(\arctan(q_0/p_0), \pi)$. Note that here we are choosing arctan to be valued between 0 and π. Using density, let us pick some object X_1 with (p_1, q_1) in the right component, and X_{-1} with (p_{-1}, q_{-1}) in the left component. Note that since the sectors span positive angles we can pick these objects with q_1, q_{-1} arbitrarily large; since

$$|\Im(Z(X)) - \Im(Z(w(X))| = |\Im(Z(X)) - q_X \Im(Z(M))|$$

is bounded for any indecomposable object X we can also guarantee that $\Im(Z(X_1))$ and $\Im(Z(X_{-1}))$ are positive.

We would like to iterate this process; at the nth step we will have objects $\{X_k\}_{-n \leq k \leq n}$ with winding numbers (p_k, q_k) running clockwise in angle, ie. $0 \leq \arctan(q_k/p_k) \leq \pi$ is decreasing. The set

$$\mathcal{X}_{1,0} \cap \mathcal{X}_{p_{-n},q_{-n}} \cap \cdots \cap \mathcal{X}_{p_0,q_0} \cap \cdots \cap \mathcal{X}_{p_n,p_q}$$

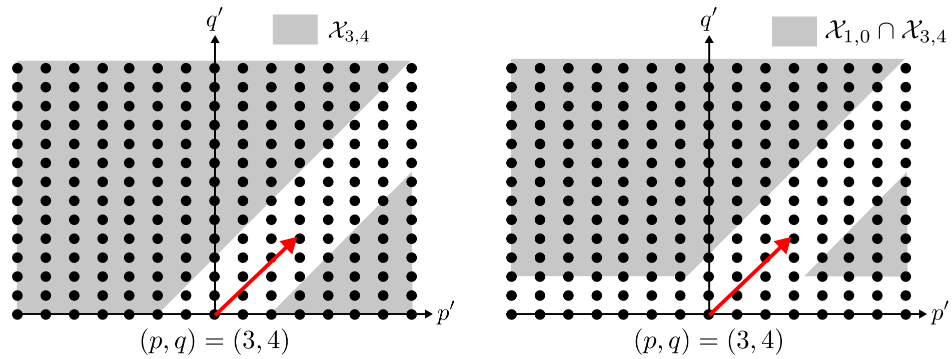

Figure 5.8: Left: the set $\mathcal{X}_{p,q}$ for $(p, q) = (3, 4)$ is composed of the $\mathbb{Z}^2$ dots inside of the shaded area. Note that all these sets have two parts, each of which at infinity approaches a sector with finite angle. Right: after the first iteration we consider $\mathcal{X}_{1,0} \cap \mathcal{X}_{3,4}$. Note that after any number of iterations the each side of this set still approaches a sector with finite angle at infinity.

at infinity approaches two sectors at $(0, \arctan(q_n/p_n))$ and $(\arctan(q_{-n}, p_{-n}), \pi)$; since each of these sectors has nonzero angle we can use density and repeat the process by picking stable objects X_{-n-1}, X_{n+1} in each sector, also both with central charge with positive imaginary part. Also from density of Φ it follows that we can pick objects such that

$$\lim_{k\to+\infty} \arctan(q_k/p_k)) = 0, \quad \lim_{k\to-\infty} \arctan(q_k/p_k)) = \pi.$$

Iterating to infinity we get stable objects $\ldots, X_{-1}, X_0, X_1, \ldots$ all mutually transversely intersecting, that also transversely intersect I as well. Taking appropriate shifts we can guarantee that all these objects have phases $0 \leq \phi_k \leq 1$. We then get that

$$\lim_{k\to+\infty} \phi_k = 0, \quad \lim_{k\to-\infty} \phi_k = 1.$$

Let d_k be the degree of the intersection between X_k and I, and f_k be the degree of the intersection between X_k and X_{k+1}. Let us shift I such that $d_0 = -1$. The triangles with sides X_k, X_{k+1}, I give the relations $d_k = d_{k+1} + f_k$. Since all the objects are stable we have inequalities for the phases

$$\phi_k \leq \phi_I + d_k \leq \phi_k + 1, \qquad \phi_k \leq \phi_{k+1} + f_k \leq \phi_k + 1.$$

But we chose the shifts such that all the ϕ_k are in $(0, 1)$, so we must have $f_k = 0$ for all k, and therefore $d_k = -1$ for all k, so $\phi_k - 1 \leq \phi_I \leq \phi_k$.

Taking the two limits $k \to +\infty$ and $k \to -\infty$ gives us $\phi_I = \phi_L = 0$ which contradicts the genericity of σ. □

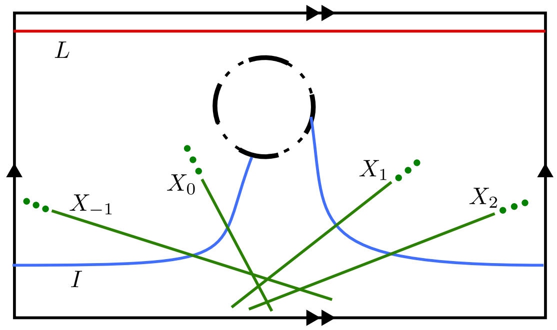

Figure 5.9: Stable circle L and stable interval I with ends on different boundary components, together with transversely intersecting stable objects $X_i, i \in \mathbb{Z}$.

5.4 Conclusions

The calculations for the three base cases above show that the every generic stability condition on those categories is an HKK stability condition; because of [56, Theorem 5.3] the image of the moduli of HKK stability conditions in $\mathrm{Stab}(\mathcal{D})$ is an union of connected components, so for all these cases there are only HKK stability conditions.

The cutting and gluing procedures allow us to reduce the calculation to the three base cases, and because of Theorem 75 this proves Theorem 60: every stability condition on a graded surface Σ is an HKK stability condition, ie. given by a quadratic differential with essential singularities.

Future directions

An obvious direction of future inquiry is the extension of the definition of relative stability conditions to Fukaya categories of higher-dimensional spaces.

With inspiration in the conjectures of Kontsevich [71], the wrapped Fukaya category of a Weinstein manifold has recently been proven [50, 49, 51] to localize to a cosheaf of categories on the Lagrangian skeleton of the Weinstein manifold, the same way that the Fukaya categories that we considered here can be calculated by a cosheaf on the dual graph.

The naive generalization of Definition 29 to this cosheaf in higher dimensions is easy to write down, but it still unclear whether one has the same nice results. We believe that the main difficulty in establishing similar results in more generality is that we lack equivalents of Lemmas 47, 50 and Theorem 53; and more fundamentally we are not aware of geometric representability results such as Theorem 43 in higher dimensions. Note that these were very important to prove our results, since even defining the cutting and gluing maps required:

1. Constraining the isomorphism type of the category $\mathrm{HNEnv}(X)$ for a certain class of object X (Lemma 52).

2. Having a non-crossing Lemma 50, which lets us separate the HN decomposition of some objects into a left and a right side.

We are of the opinion that answering the analogous questions for higher dimensions is the first step towards progress in that direction.

Another area of future research is to explore the relations between relative stability conditions and the work of Dimitrov, Haiden, Katzarkov and Kontsevich [37, 35, 36], which relates stability conditions on Fukaya categories of surfaces to questions about dynamics on the surface. In particular, it is likely that the cutting and gluing procedures of Section 5.2 can be used to say something about the distribution of stable phases for general surfaces; once we cut a surface Σ into disks, annuli and punctured tori, the collection of stable objects in $\mathcal{F}(\Sigma)$ can be produced algorithmically from collections of stable objects on each piece. It appears that one could use this to give a partial answer to Question 4.9 of [37], about the existence of conditions on a triangulated category $\mathcal{T}$ constraining the distribution of accumulation points in the set of stable phases; this will be a topic of future research.

Bibliography

[1] Mohammed Abouzaid. "A cotangent fibre generates the Fukaya category". In: *Advances in Mathematics* 2.228 (2011), pp. 894–939.

[2] Mohammed Abouzaid. "A geometric criterion for generating the Fukaya category". In: *Publications mathématiques de l'IHÉS* 112 (2010), pp. 191–240.

[3] Mohammed Abouzaid. "On the wrapped Fukaya category and based loops". In: *Journal of Symplectic Geometry* 10.1 (2012), pp. 27–79.

[4] Mohammed Abouzaid and Paul Seidel. "An open string analogue of Viterbo functoriality". In: *Geometry & Topology* 14.2 (2010), pp. 627–718.

[5] Mina Aganagic et al. "Topological strings, D-model, and knot contact homology". In: *Advances in Theoretical and Mathematical Physics* 18.4 (2014), pp. 827–956.

[6] Paul S Aspinwall and Michael R Douglas. "D-brane stability and monodromy". In: *Journal of High Energy Physics* 2002.05 (2002), p. 031.

[7] Paul S Aspinwall and Albion Lawrence. "Derived categories and zero-brane stability". In: *Journal of High Energy Physics* 2001.08 (2001), p. 004.

[8] Paul Aspinwall et al. *Dirichlet branes and mirror symmetry.* Vol. 4. American Mathematical Soc., 2009.

[9] Michael F Atiyah. "Topological quantum field theory". In: *Publications Mathématiques de l'IHÉS* 68 (1988), pp. 175–186.

[10] Denis Auroux. "A beginner's introduction to Fukaya categories". In: *Contact and symplectic topology.* Springer, 2014, pp. 85–136.

[11] John C Baez and James Dolan. "Higher-dimensional algebra and topological quantum field theory". In: *Journal of Mathematical Physics* 36.11 (1995), pp. 6073–6105.

[12] Arend Bayer and Emanuele Macrì. "The space of stability conditions on the local projective plane". In: *Duke Mathematical Journal* 160.2 (2011), pp. 263–322.

[13] Arend Bayer, Emanuele Macrì, and Yukinobu Toda. *Bridgeland stability conditions on threefolds I: Bogomolov-Gieseker type inequalities.* 2011. arXiv: `1103.5010`.

[14] I.N. Bernstein, I.M. Gel'Fand, and V.A. Ponomarev. "Coxeter functors and Gabriel's theorem". In: *Representation Theory: Selected Papers* 69 (1982), p. 157.

[15] Roman Bezrukavnikov and Mikhail Kapranov. "Microlocal sheaves and quiver varieties". In: *arXiv:1506.07050* (2015).

[16] Philip Boalch et al. "Quasi-Hamiltonian geometry of meromorphic connections". In: *Duke Mathematical Journal* 139.2 (2007), pp. 369–405.

[17] Philip Paul Boalch. "Geometry and braiding of Stokes data; fission and wild character varieties". In: *Annals of Mathematics* (2014), pp. 301–365.

[18] Philip Boalch and Daisuke Yamakawa. *Twisted wild character varieties*. 2015. arXiv: 1512.08091.

[19] Alexey I Bondal and Mikhail M Kapranov. "Enhanced triangulated categories". In: *Mathematics of the USSR-Sbornik* 70.1 (1991), p. 93.

[20] Christopher Brav and Tobias Dyckerhoff. "Relative Calabi–Yau structures". In: *Compositio Mathematica* 155.2 (2019), pp. 372–412.

[21] Christopher Brav and Tobias Dyckerhoff. *Relative Calabi-Yau structures II: Shifted Lagrangians in the moduli of objects*. 2018. arXiv: 1812.11913.

[22] Tom Bridgeland. *Stability conditions and Kleinian singularities*. 2005. arXiv: math/0508257.

[23] Tom Bridgeland. "Stability conditions on triangulated categories". In: *Annals of Mathematics* (2007), pp. 317–345.

[24] Tom Bridgeland and Ivan Smith. "Quadratic differentials as stability conditions". In: *Publications mathématiques de l'IHÉS* 121.1 (2015), pp. 155–278.

[25] Damien Calaque et al. "Shifted Poisson structures and deformation quantization". In: *Journal of topology* 10.2 (2017), pp. 483–584.

[26] Philip Candelas et al. "A pair of Calabi-Yau manifolds as an exactly soluble superconformal theory". In: *Nuclear Physics B* 359.1 (1991), pp. 21–74.

[27] Claude Cibils. "Hochschild homology of an algebra whose quiver has no oriented cycles". In: *Representation Theory I Finite Dimensional Algebras*. Springer, 1986, pp. 55–59.

[28] Kai Cieliebak and Yakov Eliashberg. *From Stein to Weinstein and back: symplectic geometry of affine complex manifolds*. Vol. 59. American Mathematical Soc., 2012.

[29] John Collins, Alexander Polishchuk, et al. "Gluing stability conditions". In: *Advances in Theoretical and Mathematical Physics* 14.2 (2010), pp. 563–608.

[30] Kevin Costello. "Topological conformal field theories and Calabi–Yau categories". In: *Advances in Mathematics* 210.1 (2007), pp. 165–214.

[31] Kevin J Costello. *The Gromov-Witten potential associated to a TCFT*. 2005. arXiv: arXiv:math/0509264.

[32] William Crawley-Boevey and Peter Shaw. "Multiplicative preprojective algebras, middle convolution and the Deligne–Simpson problem". In: *Advances in Mathematics* 201.1 (2006), pp. 180–208.

[33] Andrea D'Agnolo and Masaki Kashiwara. "Riemann-Hilbert correspondence for holonomic D-modules". In: *Publications mathématiques de l'IHÉS* 123.1 (2016), pp. 69–197.

[34] Pierre Deligne, Bernard Malgrange, and Jean-Pierre Ramis. "Singularités irrégulières, in Documents Mathématiques". In: (2007).

[35] George Dimitrov. "Bridgeland stability conditions and exceptional collections". PhD thesis. uniwien, 2015.

[36] George Dimitrov and Ludmil Katzarkov. "Bridgeland stability conditions on the acyclic triangular quiver". In: *Advances in Mathematics* 288 (2016), pp. 825–886.

[37] George Dimitrov et al. "Dynamical systems and categories". In: *The influence of Solomon Lefschetz in geometry and topology* 50 (2014), pp. 133–170.

[38] Michael R Douglas. "D-branes, categories and N= 1 supersymmetry". In: *Journal of Mathematical Physics* 42.7 (2001), pp. 2818–2843.

[39] Michael R Douglas, Bartomeu Fiol, and Christian Römelsberger. "Stability and BPS branes". In: *Journal of High Energy Physics* 2005.09 (2005), p. 006.

[40] Tobias Ekholm, Lenhard Ng, and Vivek Shende. "A complete knot invariant from contact homology". In: *Inventiones mathematicae* 211.3 (2018), pp. 1149–1200.

[41] Tobias Ekholm et al. "Knot contact homology". In: *Geometry & Topology* 17.2 (2013), pp. 975–1112.

[42] Yakov Eliashberg. *Weinstein manifolds revisited*. 2017. arXiv: `1707.03442`.

[43] Vladimir Fock and Alexander Goncharov. "Moduli spaces of local systems and higher Teichmüller theory". In: *Publications Mathématiques de l'IHÉS* 103 (2006), pp. 1–211.

[44] Sergey Fomin, Michael Shapiro, and Dylan Thurston. "Cluster algebras and triangulated surfaces. Part I: Cluster complexes". In: *Acta Mathematica* 201.1 (2008), pp. 83–146.

[45] Daniel S Freed. "Lectures on topological quantum field theory". In: *Integrable Systems, Quantum Groups, and Quantum Field Theories*. Springer, 1993, pp. 95–156.

[46] K. Fukaya et al. *Lagrangian Intersection Floer Theory: Anomaly and Obstruction, Part I*. AMS/IP studies in advanced mathematics. American Mathematical Society.

[47] Davide Gaiotto, Gregory W Moore, and Andrew Neitzke. "Wall-crossing, Hitchin systems, and the WKB approximation". In: *Advances in Mathematics* 234 (2013), pp. 239–403.

[48] Dennis Gaitsgory and Nick Rozenblyum. *A study in derived algebraic geometry: Volume I: correspondences and duality*. Vol. 1. American Mathematical Soc., 2017.

[49] Sheel Ganatra, John Pardon, and Vivek Shende. *Covariantly functorial wrapped Floer theory on Liouville sectors*. 2017. arXiv: `1706.03152`.

[50] Sheel Ganatra, John Pardon, and Vivek Shende. *Microlocal Morse theory of wrapped Fukaya categories*. 2018. arXiv: `arXiv:1809.08807`.

[51] Sheel Ganatra, John Pardon, and Vivek Shende. *Structural results in wrapped Floer theory*. 2018. arXiv: `1809.03427`.

[52] Michael Gekhtman, Michael Zalmanovich Shapiro, and Alek D Vainshtein. "Cluster algebras and Poisson geometry". In: *Moscow Mathematical Journal* 3.3 (2003), pp. 899–934.

[53] Ezra Getzler. "Batalin-Vilkovisky algebras and two-dimensional topological field theories". In: *Communications in mathematical physics* 159.2 (1994), pp. 265–285.

[54] Stéphane Guillermou, Masaki Kashiwara, Pierre Schapira, et al. "Sheaf quantization of Hamiltonian isotopies and applications to nondisplaceability problems". In: *Duke Mathematical Journal* 161.2 (2012), pp. 201–245.

[55] Sergei Gukov and Edward Witten. *Branes and quantization*. 2008. arXiv: `0809.0305`.

[56] Fabian Haiden, Ludmil Katzarkov, and Maxim Kontsevich. "Flat surfaces and stability structures". In: *Publications mathématiques de l'IHÉS* 126.1 (2017), pp. 247–318.

[57] Kentaro Hori et al. *Mirror symmetry*. Vol. 1. American Mathematical Soc., 2003.

[58] Akishi Ikeda. "Stability conditions on CY_N categories associated to A_n-quivers and period maps". In: *Mathematische Annalen* 367.1-2 (2017), pp. 1–49.

[59] Dominic Joyce. "Configurations in abelian categories - I: Basic properties and moduli stacks". In: *Advances in Mathematics* 203.1 (2006), pp. 194–255.

[60] Dominic Joyce. "Conjectures on Bridgeland stability for Fukaya categories of Calabi-Yau manifolds, special Lagrangians, and Lagrangian mean curvature flow". In: *arXiv preprint arXiv:1401.4949* (2014).

[61] Dominic Joyce. "Special Lagrangian submanifolds with isolated conical singularities. I. Regularity". In: *Annals of Global Analysis and Geometry* 25.3 (2004), pp. 201–251.

[62] Anton Kapustin. "Holomorphic reduction of N= 2 gauge theories, Wilson-'t Hooft operators, and S-duality". In: *arXiv:hep-th/0612119* (2006).

[63] Anton Kapustin and Dmitri Orlov. "Lectures on mirror symmetry, derived categories, and D-branes". In: *Russian Mathematical Surveys* 59.5 (2004), p. 907.

[64] Anton Kapustin and Dmitri Orlov. "Remarks on A-branes, mirror symmetry, and the Fukaya category". In: *Journal of Geometry and Physics* 48.1 (2003), pp. 84–99.

[65] Anton Kapustin and Dmitri Orlov. "Vertex algebras, mirror symmetry, and D-branes: the case of complex tori". In: *Communications in mathematical physics* 233.1 (2003), pp. 79–136.

[66] Masaki Kashiwara and Pierre Schapira. *Microlocal study of sheaves*. Société mathématique de France, 1985.

[67] Masaki Kashiwara and Pierre Schapira. *Sheaves on Manifolds*. Vol. 292. Springer Science & Business Media, 2013.

[68] Ludmil Katzarkov, Maxim Kontsevich, and Tony Pantev. *Hodge theoretic aspects of mirror symmetry*. 2008. arXiv: 0806.0107.

[69] Ludmil Katzarkov and Leonardo Soriani. "Homological Mirror Symmetry, coisotropic branes and P=W". In: *European Journal of Mathematics* 4.3 (2018), pp. 1141–1160.

[70] Bernhard Keller and Idun Reiten. "Cluster-tilted algebras are Gorenstein and stably Calabi–Yau". In: *Advances in Mathematics* 211.1 (2007), pp. 123–151.

[71] Maxim Kontsevich. "Homological algebra of mirror symmetry". In: *Proceedings of the international congress of mathematicians*. Springer. 1995, pp. 120–139.

[72] Maxim Kontsevich. "Symplectic geometry of homological algebra". In: *available at the author's webpage* (2009).

[73] Maxim Kontsevich and Yan Soibelman. "Stability structures, motivic Donaldson-Thomas invariants and cluster transformations". In: *arXiv:0811.2435* (2008).

[74] Chunyi Li. *On stability conditions for the quintic threefold*. 2018. arXiv: 1810.03434.

[75] Yi Li. "Topological sigma models and generalized geometries". PhD thesis. California Institute of Technology, 2005.

[76] Bong Lian, Kefeng Liu, and Shing-Tung Yau. *Mirror principle I*. 1997. arXiv: alg-geom/9712011.

[77] Jean-Louis Loday. *Cyclic homology*. Vol. 301. Springer Science & Business Media, 2013.

[78] Jacob Lurie et al. "On the classification of topological field theories". In: *Current developments in mathematics* 2008 (2009), pp. 129–280.

[79] Emanuele Macrì. "Stability conditions on curves". In: *Mathematical research letters* 14.4 (2007), pp. 657–672.

[80] Bernard Malgrange. "La classification des connexions irrégulieres a une variable". In: *Mathématiques et Physique* (1979), pp. 381–390.

[81] David Nadler. "Arboreal singularities". In: *Geometry & Topology* 21.2 (2017), pp. 1231–1274.

[82] David Nadler. "Fukaya categories as categorical Morse homology". In: *Symmetry, Integrability and Geometry: Methods and Applications* 10.0 (2014), pp. 18–47.

[83] David Nadler. "Microlocal branes are constructible sheaves". In: *Selecta Mathematica* 15.4 (2009), pp. 563–619.

[84] David Nadler. *Non-characteristic expansions of Legendrian singularities*. 2015. arXiv: 1507.01513.

[85] David Nadler. *Wrapped microlocal sheaves on pairs of pants*. 2016. arXiv: 1604.00114.

[86] David Nadler and Eric Zaslow. "Constructible sheaves and the Fukaya category". In: *Journal of the American Mathematical Society* 22.1 (2009), pp. 233–286.

[87] Nikita Nekrasov and Edward Witten. "The omega deformation, branes, integrability and Liouville theory". In: *Journal of High Energy Physics* 2010.9 (2010), p. 92.

[88] *Notes on A-infinity algebras, A-infinity categories and non-commutative geometry. I.*

[89] So Okada. "On stability manifolds of Calabi-Yau surfaces". In: *International Mathematics Research Notices* 2006.9 (2006), pp. 58743–58743.

[90] So Okada. *Stability Manifold of* $\mathbb{P}^1$. 2004. arXiv: math/0411220.

[91] Tony Pantev et al. "Shifted symplectic structures". In: *Publications mathématiques de l'IHÉS* 117.1 (2013), pp. 271–328.

[92] Alexander Postnikov. *Total positivity, Grassmannians, and networks*. 2006. arXiv: math/0609764.

[93] Pavel Safronov. "Quasi-Hamiltonian reduction via classical Chern–Simons theory". In: *Advances in Mathematics* 287 (2016), pp. 733–773.

[94] Pavel Safronov. "Symplectic implosion and the Grothendieck-Springer resolution". In: *Transformation Groups* 22.3 (2017), pp. 767–792.

[95] Graeme B Segal. "The definition of conformal field theory". In: *Differential geometrical methods in theoretical physics*. Springer, 1988, pp. 165–171.

[96] Paul Seidel. *Fukaya categories and Picard-Lefschetz theory*. Vol. 10. European Mathematical Society, 2008.

[97] Ashoke Sen. "Tachyon condensation on the brane antibrane system". In: *Journal of High Energy Physics* 1998.08 (1998), p. 012.

[98] Eric Sharpe. "D-branes, derived categories, and Grothendieck groups". In: *Nuclear Physics B* 561.3 (1999), pp. 433–450.

[99] Vivek Shende. *The conormal torus is a complete knot invariant*. 2016. arXiv: 1604.03520.

[100] Vivek Shende and Alex Takeda. *Calabi-Yau structures on topological Fukaya categories*. 2016. arXiv: 1605.02721.

[101] Vivek Shende, David Treumann, and Harold Williams. *On the combinatorics of exact Lagrangian surfaces*. 2016. arXiv: 1603.07449.

[102] Vivek Shende, David Treumann, and Eric Zaslow. "Legendrian knots and constructible sheaves". In: *Inventiones mathematicae* 207.3 (2017), pp. 1031–1133.

[103] Vivek Shende et al. *Cluster varieties from Legendrian knots*. 2015. arXiv: `1512.08942`.

[104] Carlos T Simpson. "Higgs bundles and local systems". In: *Publications Mathématiques de l'IHÉS* 75 (1992), pp. 5–95.

[105] Laura Starkston. "Arboreal singularities in Weinstein skeleta". In: *Selecta Mathematica* 24.5 (2018), pp. 4105–4140.

[106] George Gabriel Stokes. "On the discontinuity of arbitrary constants which appear in divergent developments". In: *Transactions of the Cambridge Philosophical Society* 10 (1864), p. 105.

[107] Alex A Takeda. *Relative stability conditions on Fukaya categories of surfaces*. 2018. arXiv: `1811.10592`.

[108] Constantin Teleman. "Five lectures on topological field theory". In: *Geometry and Quantization of Moduli Spaces*. Springer, 2016, pp. 109–164.

[109] Richard P Thomas and S-T Yau. *Special Lagrangians, stable bundles and mean curvature flow*. 2001. arXiv: `math/0104197`.

[110] Dylan P Thurston et al. "From dominoes to hexagons". In: *Proceedings of the 2014 Maui and 2015 Qinhuangdao Conferences in Honour of Vaughan FR Jones' 60th Birthday*. Centre for Mathematics and its Applications, Mathematical Sciences Instituted. 2017, pp. 399–414.

[111] Bertrand Toën. "The homotopy theory of dg-categories and derived Morita theory". In: *Inventiones mathematicae* 167.3 (2007), pp. 615–667.

[112] Bertrand Toën and Michel Vaquié. "Moduli of objects in dg-categories". In: *Annales scientifiques de l'Ecole normale supérieure*. Vol. 40. 3. 2007, pp. 387–444.

[113] Edward Witten. "Topological sigma models". In: *Communications in Mathematical Physics* 118.3 (1988), pp. 411–449.

[114] Daisuke Yamakawa. "Geometry of multiplicative preprojective algebra". In: *International Mathematics Research Papers* 2008 (2008).

[115] Wai-kit Yeung. *Relative Calabi-Yau completions*. 2016. arXiv: `1612.06352`.