



university of
 groningen

faculty of science
 and engineering

Homological Mirror Symmetry and the Thomas-Yau Conjecture

A tale of string theory, algebraic geometry and symplectic geometry

Presented by
 Bram Brongers



**university of
 groningen**

**faculty of science
 and engineering**

University of Groningen

Master's Thesis

To fulfill the requirements for the degree of
 Master of Science in Mathematics
 at University of Groningen under the supervision of
 dr. Nikolay Martynchuk (University of Groningen)
 and
 Prof. dr. Martijn Kool (Utrecht University)

Bram Brongers (s3208249)

July 10, 2023

Contents

Introduction	5
I Homological Mirror Symmetry	11
1 Symplectic Geometry	12
1.1 Symplectic Manifolds	12
1.2 Lagrangian Submanifolds	13
1.3 Momentum Maps and Toric Manifolds	15
2 Calabi-Yau Manifolds and Varieties	18
2.1 Complex Manifolds	18
2.2 Holomorphic Vector Bundles	22
2.3 Kähler Manifolds	25
2.4 Calabi-Yau Manifolds	27
3 The A_∞-Categories	34
3.1 The Derived Category of Coherent Sheaves	35
3.2 The Fukaya Category	40
4 Homological Mirror Symmetry	44
4.1 Mirror Symmetry and String Theory	44
4.2 Example: Mirror Symmetry for $K3$ surfaces	50
4.3 The Homological Mirror Symmetry Conjecture	54
4.4 Example: Homological Mirror Symmetry for Complex Tori	60
4.5 The B -field in Homological Mirror Symmetry	66
II The Thomas-Yau Conjecture	68
5 Donaldson-Thomas Invariants	69
5.1 Stability of Sheaves	69
5.2 The Moduli Space	71
5.3 Donaldson-Thomas Invariants	73
5.4 Gauge Theory and Holomorphic Bundles	76
6 The Kobayashi-Hitchin Correspondence	79
6.1 μ -Stability	79
6.2 The Hermitian-Yang-Mills Equation	80
6.3 Hermitian-Yang-Mills as Symplectic Quotient	81

CONTENTS

7	Hitchin Systems and $P = W$ Phenomena	85
7.1	Mirror Symmetry from Hitchin Systems	85
7.2	The $P = W$ Phenomenon	92
7.3	Enumerative Geometry and $P = W$	97
8	The Thomas-Yau Conjecture	100
8.1	BPS Branes	100
8.2	The Mirror Image	102
8.3	Wall-Crossing	108
8.4	The Thomas-Yau Conjecture	114
8.5	Example: The Elliptic Curve	116
8.6	Example: Cohomogeneity one Lagrangian submanifolds in T^{2n}	118
8.7	Example: Lagrangian fibrations of $\mathbb{C}\mathbb{P}^n$	123
9	The Thomas-Yau-Joyce Conjecture	130
9.1	Bridgeland Stability	130
9.2	Refined Donaldson-Thomas Invariants	134
9.3	The A - and B -model Donaldson-Thomas Invariants	137
10	Conclusion	140
III	Appendices	143
A	Appendix: Differential Geometry	144
A.1	Vector Bundles and Connections	144
A.2	Chern Classes	148
A.3	Holonomy	153
A.4	Spin Structures	154
B	Appendix: Sheaf Theory	155
B.1	Coherent Sheaves	155
B.2	Filtrations of Sheaves	160
C	Appendix: Floer Theory	161
C.1	Almost Complex Structures and Holomorphic Curves	161
C.2	The Maslov Index	162
C.3	Lagrangian-Floer Cohomology	165
D	Homological Algebra Preliminaries	168
D.1	Abelian categories and cohomology	168
D.2	Injective Objects and Sheaf Cohomology	171
D.3	Triangulated Categories	174
D.4	Derived Categories	178
D.5	A_∞ -Categories	182

Introduction

The topic of this thesis is the result of a personal quest to understand mirror symmetry. As such, this project really started in the summer of 2021 when I first started exploring mirror symmetry. It is my hope that readers of this text come to appreciate the beauty of the subject as I perceive it, and that they will be motivated to delve into the subject further. We will assume that the reader is familiar with differential geometry and algebraic geometry, as it is taught in a typical graduate course. We will also use homological algebra at various points. The relevant theory is recalled in the appendices. Having said that, let us begin by attempting to describe what mirror symmetry is in non-technical terms, before giving an overview of the topics covered in this text.

Mirror symmetry is a principle in string theory, so we begin there. The basic premise in string theory is that the fundamental particles should not be modelled by 0-dimensional points, but by 1-dimensional objects, namely strings. The different modes of vibration of these strings correspond to the particles that we observe. For this to work out, space-time as we know it, i.e. \mathbb{R}^4 with the Minkowski metric, has to be replaced by $\mathbb{R}^4 \times \mathcal{X}$, where \mathcal{X} is some tiny geometric shape called a Calabi-Yau threefold. String theory on $\mathbb{R}^4 \times \mathcal{X}$ is also called string theory compactified on \mathcal{X} , and \mathcal{X} is called the compactification. It provides additional dimensions for the strings to vibrate in, but it is supposedly so small that we do not currently have access to the energy scales that are required to detect it. Calabi-Yau manifolds are a very special class of smooth manifolds. The "three" in threefold refers to the complex dimension of \mathcal{X} , which implies it has real dimension 6. This leads to the famous $4 + 6 = 10$ dimensions of string theory.

So string theorists assert that, for every Calabi-Yau threefold \mathcal{X} , we obtain some notion of a string theory. Different choices of \mathcal{X} lead to different theories, just as two instruments with different shapes have different acoustics. Given that we live in a universe, which may or may not be described by string theory, the natural question to ask is: which choice of \mathcal{X} corresponds to the universe that we observe? Mirror symmetry says that this question cannot be answered definitely: for every choice of \mathcal{X} , there is another choice, called the mirror, denoted by \mathcal{X}^\vee , which results in an identical string theory. This observation has startling implications: the string theories obtained from these spaces encode some of their geometry, and as a result, mirror symmetry implies a relation between the geometry of \mathcal{X} , and that of \mathcal{X}^\vee . In general, \mathcal{X} and \mathcal{X}^\vee will not be isomorphic as complex manifolds, or even diffeomorphic. Hence, from the perspective of classical geometry, it seems absurd to suggest that any meaningful properties of \mathcal{X} could be derived from \mathcal{X}^\vee , and vice versa. But mirror symmetry gives us the tools to do so. This is the beauty of mirror symmetry: a completely unexpected duality in the geometry of Calabi-Yau manifolds,¹ which is revealed to us through ideas from physics.

¹Actually, mirror symmetry goes beyond Calabi-Yau manifolds, but we will not discuss that further.

In full generality, string theory is not mathematically precise. So how do we go from string theory to something which is mathematically precise, and allows us to make mathematical claims about the relation between the geometry of \mathcal{X} and its mirror? This is done through a procedure known as topological twisting, which yields a topological quantum field theory (TQFT). What such a TQFT is exactly does not matter right now - the point is that these topological string theories are mathematically precise, and much simpler than the original string theory. There are two ways to obtain a TQFT from a string theory on a given Calabi-Yau manifold \mathcal{X} , and each retains some of the properties of the original string theory. One uses only the symplectic geometry of the manifold (this is called the A -model), and the other uses only the complex geometry of the manifold (this is called the B -model). It can be shown that mirror symmetry exchanges the A -model on \mathcal{X} with the B -model on \mathcal{X}^\vee , and vice versa. In this way, string theory relates two completely different types of geometry, namely algebraic geometry, and symplectic geometry.

In algebraic geometry, structures are very rigid in some precise sense. In symplectic geometry, there is much less rigidity, and this is why the symplectic side of mirror symmetry requires many more technicalities to be addressed. In spite of this, mirror symmetry allows one to make claims about symplectic geometry, based on known ideas from algebraic geometry (and vice versa, in principle). This is the beauty of mirror symmetry, in my eyes, and it is part of the underlying philosophy of the Thomas-Yau conjecture [1, 2]. More famously, it is how mirror symmetry became a mathematical discipline, when Candelas et al. used period integrals (related to Hodge theory, i.e. algebraic geometry) on the mirror quintic, to compute the Gromov-Witten invariants (which are symplectic invariants) of the quintic threefold in $\mathbb{C}\mathbb{P}^4$ [3].

Overview of the text

Our first objective is to make the claim of mirror symmetry mathematically precise. To this end, it should be noted that there are many variants of mirror symmetry, some more precise than others. Indeed, if one were to ask a physicist what mirror symmetry is, the answer might be that, given a Calabi-Yau manifold \mathcal{X} , there is a second Calabi-Yau manifold \mathcal{X}^\vee , such that type IIA string theory compactified on \mathcal{X} is isomorphic to type IIB string theory compactified on \mathcal{X}^\vee . But this is not a mathematically rigorous statement. Therefore, mathematicians have constructed their own mirror symmetry conjectures, primarily within the framework of topological string theory, which is mathematically rigorous. One of these is the conjecture formulated by Kontsevich, which was named homological mirror symmetry. It states the following.

Convention 1. Throughout the text, we will be working with complex manifolds, which will typically be denoted by a calligraphic letter \mathcal{X} or \mathcal{Y} . Sometimes, relevant constructions do not depend on the holomorphic structure on the manifold. This is why we will denote the underlying smooth manifold, which contains strictly less data, by X , Y , etc. When something does not depend on the holomorphic structure, we use X (e.g. $b_k(X)$)

for Betti numbers), but when something does depend on the holomorphic structure, we use \mathcal{X} (e.g. $h^{p,q}(\mathcal{X})$ for Hodge numbers).

Conjecture 0.0.1 (Kontsevich [4]). *Let (\mathcal{X}, ω) and $(\mathcal{X}^\vee, \omega^\vee)$ be a mirror pair of Calabi-Yau threefolds. Then there exists a quasi-equivalence of A_∞ -categories*

$$D^b(\mathcal{X}) \cong D^b \text{Fuk}(X^\vee, \omega^\vee) \quad D^b(\mathcal{X}^\vee) \cong D^b \text{Fuk}(X, \omega)$$

Here, the respective categories are the bounded derived category of coherent sheaves, and the bounded derived Fukaya category. This is a weaker statement than the physicist's mirror symmetry conjecture, but is derived from it by first performing a "topological twisting" and obtaining a topological string theory. In the first part of this thesis, we will explain the mathematical ingredients that go into this conjecture, as well as the notion of a mirror pair of Calabi-Yau manifolds.

To do this, after reviewing the relevant differential geometry, we will briefly touch on the string theoretic notions which motivate homological mirror symmetry, mainly following [5, 6]: the A -model and the B -model which are obtained from a non-linear sigma model with Calabi-Yau target space, after performing some topological twisting, and the notion of a D -brane. The basic way to think about a D -brane is as a boundary condition for the worldsheet of an open string. As an open string propagates through spacetime, it traces out its worldsheet. The paths traced out by the endpoints of the string live on D -branes. For now, say that such data is given by a submanifold $Y \subset X$, where X is the smooth manifold underlying the Calabi-Yau manifold, together with a vector bundle $E \rightarrow Y$ and a connection ∇ on E , i.e. a triple (Y, E, ∇) . In general, this may not be true; one has to work in the large volume limit to obtain a geometric interpretation of D -branes, which is what we will do, without further mention of this.

A D -brane for the A -model, called an A -brane, is given by the data of a graded Lagrangian submanifold $L \subset X$, a vector bundle $E \rightarrow L$ and a flat unitary connection ∇ on E . It can be shown that a D -brane for the B -model, called a B -brane, consists of a triple (\mathcal{Y}, E, ∇) where $\mathcal{Y} \subset \mathcal{X}$ is a complex submanifold and ∇ is a connection whose curvature is of type $(1, 1)$. Equivalently, this means that a B -brane is given by a complex submanifold \mathcal{Y} together with a holomorphic vector bundle $\mathcal{E} \rightarrow \mathcal{Y}$.

As mentioned, the A - and B -model are topological twistings of some string theory. This untwisted theory also has its D -branes. Some of these may arise as A -branes or B -branes from the twisted theories, and the ones that do are called A -type BPS branes and B -type BPS branes, respectively.² The A -type BPS branes are A -branes with an additional criterion imposed: the connection ∇ needs to satisfy a partial differential equation called the Hermitian-Yang-Mills equation. The B -type BPS branes are B -branes with an

²The A -type BPS branes are BPS states for type IIB string theory, whereas B -type BPS branes are BPS states for type IIA string theory. Because this is rather confusing, we will omit any mention of type IIA/IIB string theory.

CONTENTS

additional criterion imposed: they satisfy a partial differential equation called the special Lagrangian condition. All of these terms will be explained in the text.

This brings us to the Thomas-Yau conjecture. It starts with a celebrated theorem known as the Kobayashi-Hitchin correspondence, or as the Donaldson-Uhlenbeck-Yau theorem [7]. The B -type BPS branes carry connections which satisfy the Hermitian-Yang-Mills equation. The Kobayashi-Hitchin correspondence asserts that the existence of a solution to this equation is equivalent to μ -stability - this is an algebro-geometric condition which comes from geometric invariant theory. More specifically, it says that the gauge orbit of the Chern connection contains a unique solution to the Hermitian-Yang-Mills equation if and only if the holomorphic vector bundle is μ -(poly)stable.

The story behind the Kobayashi-Hitchin correspondence can be viewed through the lens of infinite dimensional symplectic reduction, as seen in [8]. This shows how there is an infinite dimensional Lie group acting on the space of B -branes (\mathcal{V}, E, ∇) with fixed Chern character $\text{ch}(E) \in H^{\text{ev}}(X, \mathbb{Q})$, and the orbit of the Chern connection under this group has a unique Hermitian-Yang-Mills representative (i.e. has a B -type BPS brane) if and only if the vector bundle is μ -stable. By analogy, there is a group acting in a Hamiltonian fashion on the space of A -branes (L, E, ∇) with a fixed homology class $[L] \in H_{\frac{1}{2} \dim X}(X, \mathbb{Z})$, and we will work out some of the details that are left implicit in the literature on the subject. The Thomas-Yau conjecture asserts that the orbit of a given A -brane has a unique special Lagrangian representative (i.e. has an A -type BPS brane in its orbit) if and only if the A -brane is stable in some appropriate sense. Furthermore, a stable A -brane should converge to its special Lagrangian representative under the mean curvature flow [2].

After explaining this conjecture, we discuss the Thomas-Yau conjecture for the simplest possible case, namely that of the elliptic curve, which was worked out in [1], but we will provide some further details. Using this, we perform some calculations for cohomogeneity one Lagrangian submanifolds in higher dimensional tori, proving a T^{n-1} -invariant version of the Thomas-Yau conjecture in this case. In fact, we also show the following: given that the Thomas-Yau conjecture holds on a Calabi-Yau manifold \mathcal{X} , it holds in some appropriate T^n -invariant sense on $\mathcal{X} \times T^{2n}$ for T^n -invariant Lagrangian submanifolds. To the authors best knowledge, this result is not stated or proved anywhere in the literature, but we will use the methods employed in [9] to arrive at this result. We then recover higher dimensional tori by taking \mathcal{X} to be the elliptic curve, which is presently the only compact Calabi-Yau manifold for which the Thomas-Yau conjecture is known to hold. We also look at the behaviour of fibres of the standard Lagrangian torus fibration of $\mathbb{C}\mathbb{P}^n$. Whilst not Calabi-Yau, $\mathbb{C}\mathbb{P}^n$ is a Fano variety for which a version of mirror symmetry is known to hold [10]. The mean curvature flow will turn out to preserve the fibration, but exhibits very different behaviour from the Calabi-Yau case. Of course, this is to be expected, given that $\mathbb{C}\mathbb{P}^n$ is not Ricci flat. In fact, it turned out that relatively recently, a much more general result was established in [11], where the mean curvature flow of the fibration of an arbitrary toric Kähler manifold is determined.

Our method of proof is quite different, however, as we will be lifting Lagrangians in $\mathbb{C}\mathbb{P}^n$, to Lagrangians in \mathbb{C}^{n+1} using symplectic reduction, and calculate the mean curvature vector in \mathbb{C}^{n+1} by using the Calabi-Yau property of \mathbb{C}^{n+1} . In [11], a method is employed which only uses the data of the Delzant polytope of the toric manifold, and we will use this to see what happens to the toric fibrations of Hirzebruch surfaces as well.

After performing these calculations, we turn our eye to the Thomas-Yau-Joyce conjecture, which is an updated version of the Thomas-Yau conjecture, taking into account the formulation of a categorical notion of stability, called a Bridgeland stability condition [12]. The Thomas-Yau-Joyce conjecture asserts that the map

$$D^b\text{Fuk}(X, \omega) \rightarrow \mathbb{C} \quad (L, E, \nabla) \mapsto \int_L \Omega$$

defines a Bridgeland stability condition (or, more accurately, the central charge corresponding to such a stability condition). We explain that the machinery of Kontsevich-Soibelman [13] may be used to extract enumerative invariants, called refined Donaldson-Thomas invariants, from a CY3 (Calabi-Yau threefold) category together with a Bridgeland stability condition on it. In his original paper, Thomas noted that there was no *A*-side analogue of Donaldson-Thomas invariants, introduced in [14], on Calabi-Yau threefolds. These invariants "count" the number of *B*-type BPS branes in the large volume limit, in some appropriate sense. If correct, the Thomas-Yau-Joyce conjecture allows one to define this *A*-side analogue using the Kontsevich-Soibelman machinery, and not just in the large volume limit. Using homological mirror symmetry, the Thomas-Yau-Joyce conjecture, and the construction of Kontsevich-Soibelman, one can define the appropriate count of *B*-type BPS branes, outside of the large volume limit as well. This may be interesting because there are currently no known Bridgeland stability conditions on $D^b(\mathcal{X})$ for a general Calabi-Yau threefold. If the Thomas-Yau-Joyce conjecture is true, one can invoke homological mirror symmetry to establish the existence of a Bridgeland stability condition on $D^b(\mathcal{X}^\vee)$.

Along the way, we will also touch on various other topics: closed string mirror symmetry for *K3* surfaces, the SYZ conjecture, moduli spaces of Higgs bundles and the Hitchin system, as well as the $P = W$ phenomenon for abelian varieties as treated in [15]. This allows us to look at a class of hyper-Kähler manifolds which exhibit mirror symmetry in a rather interesting way, related to the geometric Langlands correspondence and Langlands dual groups. We will then be able to see how mirror symmetry can lead to rather surprising mathematical predictions, namely the $P = W$ conjecture.³ The $P = W$ conjecture yields a way to calculate some of the enumerative invariants that are introduced throughout the text. In general, these are difficult to compute, and the $P = W$ conjecture, which has been shown to hold in many cases, simplifies this problem greatly for Calabi-Yau threefolds which are local curves (i.e. the total space of an appropriate rank 2 holomorphic bundle over a Riemann surface).

³The authors of the original paper [16] which contains the $P = W$ conjecture did not seem to be explicitly motivated by mirror symmetry.

Notational Conventions

Throughout the text, we employ certain conventions with regards to notation, and abuse thereof.

X, Y, \dots	Smooth manifolds
$\mathcal{X}, \mathcal{Y}, \dots$	Complex manifolds and complex projective varieties
v, w, \dots	Tangent vectors, vector fields
E, F, \dots	Smooth vector bundles
$\mathcal{E}, \mathcal{F}, \dots$	Holomorphic vector bundles, coherent sheaves
$\mathcal{E}_\bullet, \mathcal{F}_\bullet, \dots$	Complexes of holomorphic vector bundles, coherent sheaves
L	Lagrangian submanifolds
\mathcal{L}	Holomorphic line bundles
$\mathcal{O}_{\mathcal{X}}, K_{\mathcal{X}}, T\mathcal{X}$	Structure sheaf, canonical bundle, holomorphic tangent bundle
∇, A	Connection on a vector bundle
$\Gamma(U, E), H^0(U, E)$	Sections of a sheaf E over an open subset U
$\Omega_{\mathcal{X}}^k$	Sheaf of holomorphic k -forms on \mathcal{X}
$\Omega^k(X, E)$	Differential k -forms with values in a vector bundle E
R^\times	Multiplicative group of units of a ring R

Also, note the following:

1. If we write A for a connection on a vector bundle, this means $A \in \Omega^1(X, \text{End}(E))$ and we are really considering $\nabla = \nabla_0 + A$ with respect to some fixed connection ∇_0 .
2. We write \mathcal{E} for a holomorphic vector bundle, as well as for its sheaf of holomorphic sections.
3. If E is a smooth vector bundle, we write E_x for the fibre at $x \in X$. If \mathcal{E} is a holomorphic vector bundle (or coherent sheaf), we instead write \mathcal{E}_x for its stalk at x .
4. We often regard a complex projective variety as a complex manifold and vice versa, without mentioning this.
5. We will use physics terminology occasionally even though IANAP,⁴ so some of it may be used incorrectly.

⁴I am not a physicist.

Part I

Homological Mirror Symmetry

Chapter 1

Symplectic Geometry

Symplectic manifolds arise in physics, as the phase space of some state space. They encode Hamilton's equations, which are as the foundation of Hamiltonian mechanics. Furthermore, symplectic geometry is closely related to complex geometry, via almost complex structures, discussed in the appendix C.1.1. Kähler manifolds, and in particular Calabi-Yau manifolds, lie at the intersection of these two overlapping domains. In this section, we recap the basics of symplectic geometry, which are required to define the Fukaya category of a symplectic manifold as well as our subsequent discussion of the Thomas-Yau conjecture. This chapter is based on [17, 18].

1.1 Symplectic Manifolds

Let us start off by recalling the basic definitions.

Definition 1.1.1. Let X be a smooth manifold, and $\omega \in \Omega^2(X)$. Then the pair (X, ω) is called a symplectic manifold if ω is closed and non-degenerate. A symplectomorphism between symplectic manifolds (X_1, ω_1) , (X_2, ω_2) is a diffeomorphism $f : X_1 \rightarrow X_2$ such that $f^* \omega_2 = \omega_1$.

Non-degeneracy of ω can be phrased in several different ways. In local coordinates, it amounts to invertibility of a matrix representation of ω . Because the matrix of ω is skew-symmetric, its invertibility implies that X must be even dimensional. We will generally take the dimension of X to be $2n$. More invariantly, non-degeneracy can be stated as saying that ω provides an isomorphism $\varphi : TX \rightarrow T^*X$ through $v \mapsto \iota_v \omega$. The canonical example of a symplectic manifold is given by the cotangent bundle.

Example 1. Let X be any smooth manifold, and consider the tautological 1-form τ on T^*X , which is defined as follows. For $v \in T_{(x,\eta)}(T^*X)$, we define $\tau_{(x,\eta)}(v) := \eta(d\pi_{(x,\eta)} v)$, where $\pi : T^*X \rightarrow X$ is the canonical projection. Then the 2-form $-d\tau$ is a symplectic form on T^*X , called the canonical symplectic form. In standard local coordinates $(x^1, \dots, x^n, p^1, \dots, p^n)$ on T^*X , the 1-form τ is given by $\sum p^i dx^i$, and the symplectic form is given by $\sum dx^i \wedge dp^i$.

This example is typically interpreted physically as saying that X is the state space, and T^*X is the phase space. Hence, the coordinates x^i parameterise the position of some

1.2. LAGRANGIAN SUBMANIFOLDS

particle, while the coordinates p^i parameterise its momentum. The following classical theorem tells us that this is the model space for any symplectic manifold.

Theorem 1.1.2 (The Darboux Theorem). *Let (X, ω) be a symplectic manifold, and $x \in X$. Then there exist local coordinates $(x^1, \dots, x^n, p^1, \dots, p^n)$ centered at x such that $\omega = \sum dx^i \wedge dp^i$.*

Like the cotangent bundle, every symplectic manifold is orientable, since ω^n provides a volume form on X . To describe a physical system, we also need a Hamiltonian. This is a smooth function $H \in C^\infty(X)$.

Definition 1.1.3. Let $H \in C^\infty(X)$ be a smooth function. The Hamiltonian vector field v_H of H is defined by

$$dH = \iota_{v_H} \omega$$

Observe that H is preserved under the flow of v_H , since $v_H H = dH(v_H) = \omega(v_H, v_H) = 0$. If we take v_H to describe the dynamics of a physical system, then this means we are constrained to the level sets of H , which is also known as conservation of energy.

We could also consider a time-dependent Hamiltonian, which would be a smooth function $H_t \in C^\infty(X \times I)$. These will be of some importance for the Thomas-Yau conjecture. In particular, we will need the notion of a Hamiltonian isotopy. Let us start with a symplectic isotopy. We have a group of symplectic diffeomorphisms $\varphi : X \rightarrow X$, which we denote $\text{Symp}(X, \omega)$. It is a subgroup of $\text{Diff}(X)$, and inherits a topology. A symplectic isotopy is a continuous family of symplectomorphisms $\varphi_t : X \rightarrow X$, which can also be viewed as a smooth map $\varphi_t : X \times I \rightarrow X$ such that $\varphi_{t_0} : X \rightarrow X$ is a symplectomorphism for all $t_0 \in I$. A symplectic isotopy is generated by a unique time-dependent vector field $v_t : X \times I \rightarrow TX$, defined by

$$\frac{d}{dt} \varphi_t = v_t \circ \varphi_t$$

Definition 1.1.4. Let $\varphi_t : X \times I \rightarrow X$ be a symplectic isotopy. Then φ_t is said to be a Hamiltonian isotopy if there exists $H \in C^\infty(X \times I)$ such that for all $t \in I$,

$$\iota_{v_t} \omega = dH_t$$

1.2 Lagrangian Submanifolds

Objects which are of central importance in symplectic geometry are so-called Lagrangian submanifolds. To define these, we first define some preliminary notions. We consider ω_x as a bilinear form on the vector space $T_x X$. Let $W \subseteq T_x X$ be a linear subspace. Then we define $W^\perp := \{v \in T_x X \mid \omega(v, w) = 0 \quad \forall w \in W\}$ to be its symplectic complement. A linear subspace W is called isotropic if $W \subseteq W^\perp$, co-isotropic if $W^\perp \subseteq W$, and Lagrangian if $W = W^\perp$. Every Lagrangian subspace must have dimension $\frac{1}{2} \dim V$, so Lagrangian subspaces can also be characterised as maximally isotropic subspaces.

1.2. LAGRANGIAN SUBMANIFOLDS

Definition 1.2.1. Let $Y \subseteq X$ be an embedded submanifold of a symplectic manifold. Then Y is called Lagrangian if $T_x Y \subseteq T_x X$ is a Lagrangian subspace for all $x \in Y$.

Of course, (co-)isotropic submanifolds are defined in analogous fashion. By definition, a submanifold Y is Lagrangian if and only if $\dim Y = \frac{1}{2} \dim X$ and $\omega|_Y \equiv 0$.

Example 2. The most important example of a Lagrangian submanifold is the zero section of the cotangent bundle T^*X with its canonical symplectic form. Weinstein's Lagrangian neighbourhood theorem states that any closed Lagrangian submanifold looks like this. More precisely, every Lagrangian submanifold $L \subseteq X$ admits an open neighbourhood U and a symplectomorphism $\varphi : U \rightarrow V \subseteq T^*L$, such that $\varphi(L)$ is the zero section of T^*L .

Lagrangian submanifolds generalise the idea of conservation of energy, as we will now see through the lens of integrable systems.

Definition 1.2.2. Let (X, ω) be a symplectic manifold. Then there exists a Poisson bracket on $C^\infty(X)$, defined by

$$\{f, g\} := \omega(v_f, v_g)$$

We recall that a Poisson bracket is a Lie bracket on $C^\infty(X)$ which is a derivation with respect to the multiplicative structure, as well as the Lie algebra structure. That is, it also satisfies $\{f \cdot g, h\} = f \cdot \{g, h\} + g \cdot \{f, h\}$. The requirement that $d\omega = 0$ in the definition of a symplectic form, amounts precisely to the requirement that this bracket satisfies the Jacobi identity. Two functions are said to Poisson commute when their Poisson bracket vanishes. Note that $\omega(v_f, v_g) = df(v_g) = v_g f$, so this amounts to saying that the functions are preserved under each other's flows. We can locally write a submanifold as $Y = f^{-1}(0)$, where $f = (f_1, \dots, f_n) : X \rightarrow \mathbb{R}^n$ is smooth with 0 as a regular value.

Proposition 1.2.3. [18] *The submanifold Y is Lagrangian if and only if the functions f_i Poisson commute pairwise. That is, $\{f_i, f_j\}|_Y = 0$ for all $i, j \geq 1$.*

Recall that $\ker df_x = T_x Y$. Since $d(f_i) = \iota_{v_{f_i}} \omega$, this implies that a Hamiltonian $H \in C^\infty(X)$ satisfies $(v_H)_x \in T_x Y \iff \{f_i, H\} = 0$ for all i .

Definition 1.2.4. Let (X, ω, H) be a Hamiltonian system. A function $f \in C^\infty(X)$ is called an integral of motion if $\{f, H\} = 0$, i.e. if $v_H f = 0$.

Thus, if a Lagrangian submanifold is defined by n -many integrals of motion (which are necessarily in involution), then the trajectory of the flow of v_H is completely contained in any Lagrangian submanifold that it intersects.

Definition 1.2.5. Let (X, ω, H) be a Hamiltonian system. If the system admits n independent Poisson commuting integrals of motion, then the system is said to be completely integrable.

1.3. MOMENTUM MAPS AND TORIC MANIFOLDS

Theorem 1.2.6 (Arnold-Liouville [18]). *Let (X, ω, H) be a completely integrable Hamiltonian system with integrals of motion f_i . Suppose that 0 is a regular value of $f = (f_1, \dots, f_n)$. Then there exists an open interval $(-\varepsilon, \varepsilon) = I$ such that $f^{-1}(E) := L_E$ is a Lagrangian submanifold for all $E \in I$. If L_0 is compact and connected, then $L_E \cong T^n$.*

The manifold L_0 is always compact if f is a proper map. In the case where $L_E \cong T^n$, the trajectories are constrained to these tori. Arnold also proved that this motion is linear on the torus. In particular, completely integrable systems give us examples of Lagrangian torus fibrations. Such fibrations are also intimately related to so-called SYZ (Ströminger-Yau-Zaslow) approach to mirror symmetry, originating from their famous paper [19]. Completely integrable systems are the "holy grail" of Hamiltonian mechanics, because this gives us a (relatively) simple way to describe the time evolution of the system.

Every integrable system admits a Lagrangian torus fibration by taking as the fibres the common level sets of the integrals of motion. A fibration is distinctly different from a fibre bundle, because a Lagrangian torus fibration may admit singular fibres.

Definition 1.2.7. Let (X, ω) be a symplectic manifold. A Lagrangian torus fibration of X is a proper continuous map $\pi : X \rightarrow B$ such that there is a dense open subset $B_{\text{reg}} \subset B$ for which $\pi^{-1}(B_{\text{reg}}) \rightarrow B_{\text{reg}}$ is a smooth submersion whose fibre is a Lagrangian torus. If $B_{\text{sing}} = B \setminus B_{\text{reg}}$, then the fibre of π over B_{sing} is a connected stratified space whose strata are isotropic in X .

Example 3. Suppose that B is an integral affine manifold, which means that it admits an atlas whose associated transition functions take values in $\text{GL}(n, \mathbb{Z}) \times \mathbb{R}^n \subset \text{Diff}(\mathbb{R}^n)$. We can choose a local \mathbb{Z}^n -system $T_{\mathbb{Z}}B \subset TB$, which exists globally because of the assumption on the transition functions. This yields a dual lattice $T_{\mathbb{Z}}^*B \subset T^*B$ after choosing a metric. Define $X = T^*B/T_{\mathbb{Z}}^*B$, which carries a natural projection map $\pi : X \rightarrow B$. The manifold X is a symplectic manifold, since the natural symplectic structure on T^*B descends to X . Since each T_b^*B is a Lagrangian submanifold of T^*B , each fibre of π is a Lagrangian submanifold which is diffeomorphic to $T_b^*B/\mathbb{Z}^n \cong T^n$, which means that $\pi : X \rightarrow B$ is a Lagrangian torus fibration.

1.3 Momentum Maps and Toric Manifolds

Later on, we will be considering a very special kind of integrable system, namely a toric manifold. To define these, we first need the notion of a momentum map. Suppose that a Lie group G acts on a symplectic manifold (X, ω) by symplectomorphisms. Then each $\xi \in \mathfrak{g}$ defines a vector field

$$(v_{\xi})_x := \left. \frac{d}{dt} \right|_{t=0} x \cdot \exp(t\xi)$$

1.3. MOMENTUM MAPS AND TORIC MANIFOLDS

The group action is said to be Hamiltonian if there exists a smooth function $\mu : X \rightarrow \mathfrak{g}^*$ which is G -equivariant (w.r.t. the coadjoint representation of G on \mathfrak{g}^*) and satisfies

$$\iota_{v_\xi} \omega = d\langle \mu, \xi \rangle$$

for all $\xi \in \mathfrak{g}$. Note that $\mu^{-1}(0)$ is invariant under the G -action (in fact, this is true for any central element of \mathfrak{g}^*).

Definition 1.3.1. Suppose that G acts on (X, ω) with momentum map μ . Suppose G acts freely and properly on $\mu^{-1}(0)$. Then

$$X // G := \mu^{-1}(0)/G$$

is called the symplectic reduction of X (at 0).

Example 4. Consider the group action

$$\mathbb{C}^n \times S^1 \rightarrow \mathbb{C}^n \quad ((z_1, \dots, z_n), e^{i\theta}) \mapsto (z_1 e^{i\theta}, \dots, z_n e^{i\theta})$$

It has a momentum map given by $\mu(z_1, \dots, z_n) = |z_1|^2 + \dots + |z_n|^2$. Symplectic reduction at 1 yields

$$\mathbb{C}^n // S^1 = \mu^{-1}(1)/S^1 = S^{2n-1}/S^1 \cong \mathbb{C}\mathbb{P}^{n-1}$$

Definition 1.3.2. A toric manifold consists of a compact symplectic manifold (X, ω) together with a Hamiltonian T^n -action.

Theorem 1.3.3 (Delzant). *The image of the momentum map of a toric manifold is a convex polytope in \mathbb{R}^n , called the Delzant polytope. The fibres are the T^n -orbits.*

Example 5. Consider the T^n -action on $\mathbb{C}\mathbb{P}^n$ given by

$$([z_0 : \dots : z_n], e^{i\theta_1}, \dots, e^{i\theta_n}) \mapsto [z_0 : z_1 e^{i\theta_1} : \dots : z_n e^{i\theta_n}]$$

This action has a momentum map given by

$$\mu : [z_0 : \dots : z_n] \mapsto \frac{1}{|z|^2} (|z_1|^2, \dots, |z_n|^2)$$

The Delzant polytope is the regular n -simplex in \mathbb{R}^n , which we denote Δ_n . The fibres of μ over $\text{int}(\Delta_n)$ are Lagrangian tori. The fibres of μ over the k -dimensional faces of Δ_n are k -dimensional isotropic tori. For instance, take $\mathbb{C}\mathbb{P}^1 \cong S^2$. Then the Delzant polytope is $I = [0, 1]$. The fibre over $\{0\}$ and $\{1\}$ is a single point, and the fibre over $c \in (0, 1)$ is a circle. This is the fibration associated to the height function $h : S^2 \supset \mathbb{R}^3 \rightarrow I \subset \mathbb{R}$ given by $h(x, y, z) = z$, which can easily be visualised.

1.3. MOMENTUM MAPS AND TORIC MANIFOLDS

Example 6. There is another class of examples that we will return to, namely the Hirzebruch surfaces which are indexed by an integer. We denote the n -th Hirzebruch surface by \mathcal{H}_n , and its Delzant polytope is given by the vertices $\{(0,0), (n+1,0), (0,1), (1,1)\}$. One can use this data to reconstruct the T^2 -action on \mathbb{C}^4 , and its momentum map, from which the n -th Hirzebruch surface is obtained by symplectic reduction. We refer to [17] for the details of this construction, which we outline now. Since we are given the vertices, we can find the primitive inward pointing normal vectors of the polytope, and they are $\{v_i\} = \{(0,1), (-1,-n), (0,-1), (1,0)\}$. Next we wish to determine the maps π and ι in an exact sequence

$$0 \xrightarrow{\iota} \mathbb{R}^2 \rightarrow \mathbb{R}^4 \xrightarrow{\pi} \mathbb{R}^2 \rightarrow 0$$

which are the Lie algebras of tori. In particular, \mathbb{R}^4 corresponds to the Lie algebra of T^4 acting on \mathbb{C}^4 in the standard manner, which we will use in a moment. The map π is determined by $e_i \mapsto v_i$, so we have

$$\begin{aligned} (1,0,0,0) &\mapsto (0,1) & (0,1,0,0) &\mapsto (-1,-n) \\ (0,0,1,0) &\mapsto (0,-1) & (0,0,0,1) &\mapsto (1,0) \end{aligned}$$

Hence $\text{im } \iota = \ker \pi = \text{span}_{\mathbb{R}}\{e_1 + e_3, ne_1 + e_2 + e_4\}$, and these vectors are the columns of the matrix of the map ι . Then the momentum map we need for the construction of the Hirzebruch surfaces is given by $\iota^* \circ \mu_0$, where $\mu_0 : \mathbb{C}^4 \rightarrow \mathbb{R}^4$ is the momentum map of the standard T^4 -action on \mathbb{C}^4 , i.e. $(z_1, z_2, z_3, z_4) \mapsto (e^{i\theta_1} z_1, e^{i\theta_2} z_2, e^{i\theta_3} z_3, e^{i\theta_4} z_4)$ with momentum map $(z_1, z_2, z_3, z_4) \mapsto (|z_1|^2, |z_2|^2, |z_3|^2 - n - 1, |z_4|^2 - 1)$, with the additional constants coming from the data of the polytope. So our momentum map is found to be

$$\mu(z_1, z_2, z_3, z_4) = (|z_1|^2 + |z_3|^2 - n - 1, n|z_1|^2 + |z_2|^2 + |z_4|^2 - 1)$$

with the torus action being given by

$$(z_1, z_2, z_3, z_4) \mapsto (e^{i\theta_1} e^{in\theta_2} z_1, e^{i\theta_2} z_2, e^{i\theta_1} z_3, e^{i\theta_2} z_4)$$

for $(e^{i\theta_1}, e^{i\theta_2}) \in T^2$. The symplectic reduction $\mu^{-1}(0)/T^2$ is defined to be \mathcal{H}_n . These manifolds are also complex manifolds, inheriting their complex structure from \mathbb{C}^4 . One can show that $\mathcal{H}_n \cong \mathbb{P}(\mathcal{O} \oplus \mathcal{O}(-n))$, where $\mathcal{O} \oplus \mathcal{O}(-n) \rightarrow \mathbb{C}\mathbb{P}^1$ is a rank 2 holomorphic vector bundle, defined in the next chapter. In particular, $\mathcal{H}_0 \cong \mathbb{C}\mathbb{P}^1 \times \mathbb{C}\mathbb{P}^1$.

The behaviour of these fibres is typical of toric manifolds: they are themselves smooth tori, although their dimensions vary. If we have an integrable system which is not toric, the fibres may acquire singularities, like the pinched tori that appear in fibrations of a $K3$ surface, as we will see later.

Chapter 2

Calabi-Yau Manifolds and Varieties

Homological mirror symmetry is a statement about a pair of Calabi-Yau manifolds (or varieties). These are special kinds of Kähler manifolds, which in turn are special kinds of complex manifolds. As such, we start by discussing complex manifolds, and work our way up to Calabi-Yau manifolds. The content in this chapter is based on [20, 21, 22].

2.1 Complex Manifolds

Definition 2.1.1. A complex manifold \mathcal{X} is a topological manifold X of dimension¹ $2n$ which admits an atlas $\{U_\alpha, \varphi_\alpha\}$ such that

$$\varphi_\alpha \circ \varphi_\beta^{-1} : \mathbb{C}^n \supseteq \varphi_\beta(U_\beta \cap U_\alpha) \rightarrow \mathbb{C}^n$$

is holomorphic for all non-empty intersections.

The notion of holomorphicity of functions and maps between complex manifolds is defined in the same way that it is defined for smooth manifolds, namely by verifying the definitions in local charts.

Example 7. The first non-trivial example of a complex manifold is $\mathbb{C}\mathbb{P}^n$, and it plays a very important role in complex geometry. We recall that $\mathbb{C}\mathbb{P}^n = (\mathbb{C}^{n+1} \setminus \{0\})/\mathbb{C}^\times$ by definition, and we have homogeneous coordinates $[z_0 : \cdots : z_n]$ on $\mathbb{C}\mathbb{P}^n$, which are defined up to \mathbb{C}^\times -scaling. We have the standard open cover $\{U_i\}_{i=0, \dots, n}$ where the sets $U_i = D(z_i) := \{[z_0 : \cdots : z_n] \in \mathbb{C}\mathbb{P}^n \mid z_i \neq 0\}$ are called the standard affine opens of $\mathbb{C}\mathbb{P}^n$. We consider a homeomorphism $U_i \cong \mathbb{C}^n$ given by

$$\varphi_i : U_i \rightarrow \mathbb{C}^n \quad [z_0 : \cdots : z_n] \mapsto \left(\frac{z_0}{z_i}, \dots, \widehat{z_i}, \dots, \frac{z_n}{z_i} \right)$$

The transition maps are then found to be

$$\varphi_j \circ \varphi_i^{-1}(w_1, \dots, w_n) = \varphi_j([w_1 : \cdots : 1 : \cdots : w_n]) = \left(\frac{w_1}{w_j}, \dots, \frac{1}{w_j}, \dots, \widehat{w_j}, \dots, \frac{w_n}{w_j} \right)$$

where we have assumed that $j < i$. The case $j > i$ is quite similar, and in both cases, we see that we obtain holomorphic maps, so that $\mathbb{C}\mathbb{P}^n$ is indeed a complex manifold.

¹Whenever we talk about the dimension of \mathcal{X} , we mean $\dim \mathcal{X} = \dim_{\mathbb{C}} X$. Whenever we talk about the dimension of X , we mean $\dim X = \dim_{\mathbb{R}} X$.

2.1. COMPLEX MANIFOLDS

Example 8. Preimages of regular values of holomorphic maps between complex manifolds are again complex manifolds, by the holomorphic implicit function theorem. In particular, homogeneous polynomials whose Jacobian have maximal rank everywhere define complex manifolds $\mathcal{X} \subseteq \mathbb{C}\mathbb{P}^n$, and Chow's theorem states that every closed submanifold of $\mathbb{C}\mathbb{P}^n$, which we will call a projective complex manifold, arises in this way, as the intersection of vanishing loci of homogeneous polynomials. That is, every projective complex manifold defines a smooth projective algebraic variety over \mathbb{C} . In particular,

$$\{[z_0 : \cdots : z_n] \in \mathbb{C}\mathbb{P}^n \mid \sum z_i^{n+1} = 0\}$$

is a complex manifold of dimension $n - 1$, which is a Calabi-Yau manifold, as we will show later.

Serre's GAGA theorems state that we have an equivalence of categories between projective complex manifolds and projective smooth varieties. Furthermore, if we denote by \mathcal{X}^{an} the projective manifold with its Euclidean topology, and by \mathcal{X} the projective smooth variety which it corresponds to, there is an equivalence of categories between their coherent sheaves: $\text{Coh}(\mathcal{X}^{\text{an}}) \cong \text{Coh}(\mathcal{X})$. All of this is simply to say that we can translate back and forth between algebraic geometry, and complex differential geometry on projective manifolds.

On a complex manifold, we can take local holomorphic coordinates $z^k : U_\alpha \rightarrow \mathbb{C}^n$ and anti-holomorphic coordinates $\bar{z}^k : U \rightarrow \mathbb{C}^n$. These also contain smooth coordinates via $z^k = x^k + iy^k$. We can define the 1-forms

$$dz^k := dx^k + idy^k \quad d\bar{z}^k := dx^k - idy^k$$

Dual to these are the vector fields

$$\frac{\partial}{\partial z^k} = \frac{1}{2} \left(\frac{\partial}{\partial x^k} + i \frac{\partial}{\partial y^k} \right) \quad \frac{\partial}{\partial \bar{z}^k} = \frac{1}{2} \left(\frac{\partial}{\partial x^k} - i \frac{\partial}{\partial y^k} \right)$$

These define an almost complex structure (see C.1.1). Namely, take $U \subseteq X$ sufficiently small so that $TX|_U$ is trivial, which makes $\Gamma(U, TX|_U)$ a free $C^\infty(U)$ -module. Then we define J on a basis for this free module via

$$J\left(\frac{\partial}{\partial x^k}\right) = \frac{\partial}{\partial y^k} \quad J\left(\frac{\partial}{\partial y^k}\right) = -\frac{\partial}{\partial x^k}$$

This local definition for J is globally well-defined because the transition functions are holomorphic, since we are on a complex manifold. The endomorphism J induces a decomposition $TX \otimes \mathbb{C} = TX^{1,0} \oplus TX^{0,1}$ into its $\pm i$ eigenspaces. When the distribution $TX^{1,0}$ is integrable, the almost complex structure J comes from a complex structure. This is the Newlander-Nirenberg theorem, which we rephrase slightly in a moment.

2.1. COMPLEX MANIFOLDS

The decomposition of $TX \otimes \mathbb{C}$ induces one on $T^*X \otimes \mathbb{C}$ via pullback. This also yields $\wedge^k T^*X \otimes \mathbb{C} = \oplus_{p+q=k} \wedge^p T^*X^{1,0} \otimes \wedge^q T^*X^{0,1}$. Consequently, any almost complex manifold has $\Omega^k(X, \mathbb{C}) = \oplus_{p+q=k} \Omega^{p,q}(X)$. In local coordinates, a basis for $\Omega^{p,q}(U)$ is given by

$$\{dz^{j_1} \wedge \cdots \wedge dz^{j_p} \wedge d\bar{z}^{k_1} \wedge \cdots \wedge d\bar{z}^{k_q} \mid j_1 < \cdots < j_p, \quad k_1 < \cdots < k_q\}$$

One can define $\partial = \pi^{p+1,q} \circ d$ and $\bar{\partial} = \pi^{p,q+1} \circ d$. By definition, then, we have

$$\partial : \Omega^{p,q}(X) \rightarrow \Omega^{p+1,q}(X) \quad \bar{\partial} : \Omega^{p,q}(X) \rightarrow \Omega^{p,q+1}(X)$$

Theorem 2.1.2 (Newlander-Nirenberg [23]). *An almost complex structure is integrable if and only if $d = \partial + \bar{\partial}$. Equivalently, if and only if $\bar{\partial}^2 = 0$.*

On a complex manifold, one defines the operators ∂ and $\bar{\partial}$ on functions by locally setting

$$\partial f = \sum_k \frac{\partial f}{\partial z^k} dz^k \quad \bar{\partial} f = \sum_k \frac{\partial f}{\partial \bar{z}^k} d\bar{z}^k$$

It should be clear how to extend this definition to (p, q) -forms.

Example 9. Suppose that Σ is a real oriented surface, and equip Σ with a Riemannian metric g . We can define the Hodge star operator $\star : \wedge^k T\Sigma \rightarrow \wedge^{2-k} T\Sigma$. In particular, we get $\star : T\Sigma \rightarrow T\Sigma$, and for dimensional reasons, $\star^2 = -\text{id}$. Thus, every oriented surface with a Riemannian metric admits an almost complex structure. Since $\Omega^{2,0}(\Sigma) = 0$, again for dimensional reasons, we see that $\bar{\partial}^2 = 0$, so the complex structure is necessarily integrable.

Let us also mention the example of complex tori, as they allow for an explicit description of the so-called complex moduli space.

Example 10. Let Λ be a lattice of maximal rank in \mathbb{C}^n . Then \mathbb{C}^n/Λ is a complex manifold, and it will be called a complex torus. We will use these as a testing ground for various ideas throughout this text.

It is instructive to consider the case $n = 1$. Since \mathbb{C} is the universal covering space of \mathbb{C}/Λ , any holomorphic map $f : \mathbb{C}/\Lambda_1 \rightarrow \mathbb{C}/\Lambda_2$ between two complex tori lifts to a holomorphic map $F : \mathbb{C} \rightarrow \mathbb{C}$. Suppose F_1, F_2 are both lifts of f . Define $G = F_1 - F_2$ and let $\pi_i : \mathbb{C} \rightarrow \mathbb{C}/\Lambda_i$ be the projection maps. Then for all $z \in \mathbb{C}$ we get

$$f \circ \pi_1(z) = \pi_2 \circ F_1(z) = \pi_2 \circ F_2(z) \implies \pi_2 \circ F_1(z) - \pi_2 \circ F_2(0) = 0 \implies G(z) \in \Lambda_2$$

By continuity of G and the fact that a lattice is discrete, we must have $G = \text{const.} \in \Lambda_2$. Now, we claim that $F(z) = \alpha z + \beta$ for some $\alpha \in \mathbb{C}$ such that $\alpha\Lambda_1 \subseteq \Lambda_2$. For $\lambda \in \Lambda_1$,

2.1. COMPLEX MANIFOLDS

define $G_\lambda(z) = F(z+\lambda) - F(z)$. Since $F(z)$ and $F(z+\lambda)$ are both lifts of f , we must have $G_\lambda(z) \in \Lambda_2$ constant. Furthermore, $\frac{d}{dz}F(z+\lambda) = \frac{d}{dz}F(z)$ so that $\frac{d}{dz}F$ is Λ_1 -periodic. Thus, it is a bounded, holomorphic function on \mathbb{C} , which must be constant. It follows that $F(z) = \alpha z + \beta$ for some complex numbers α, β . Finally, if $\lambda \in \Lambda_1$, then $G_\lambda(0) = \alpha\lambda + \beta - \beta = \alpha\lambda \in \Lambda_2$, so indeed $\alpha\Lambda_1 \subseteq \Lambda_2$.

Theorem 2.1.3. *Two complex tori \mathbb{C}/Λ_1 and \mathbb{C}/Λ_2 are isomorphic if and only if there exists an $\alpha \in \mathbb{C}^\times$ such that $\alpha\Lambda_1 = \Lambda_2$.*

Proof. Suppose $\alpha\Lambda_1 = \Lambda_2$. Let $m_\alpha(z) = \alpha z$. Since $m_\alpha(\Lambda_1) = \Lambda_2$, this descends to a holomorphic map $\widetilde{m}_\alpha : \mathbb{C}/\Lambda_1 \rightarrow \mathbb{C}/\Lambda_2$. This map has a holomorphic inverse given by $\widetilde{m}_{\alpha^{-1}} : \mathbb{C}/\Lambda_2 \rightarrow \mathbb{C}/\Lambda_1$ so that $\mathbb{C}/\Lambda_1 \cong \mathbb{C}/\Lambda_2$.

Conversely, suppose $\mathbb{C}/\Lambda_1 \cong \mathbb{C}/\Lambda_2$. Let $f : \mathbb{C}/\Lambda_1 \rightarrow \mathbb{C}/\Lambda_2$ be a biholomorphism, with lift $F : \mathbb{C} \rightarrow \mathbb{C}$, and inverse lift $G : \mathbb{C} \rightarrow \mathbb{C}$. Then $F(z) = \alpha z + \beta$ and $G(z) = \gamma z + \delta$. Looking at $F \circ G(z) = \alpha\gamma z + \alpha\delta + \beta$, we see that $\alpha\gamma = 1$. We know that

$$\begin{aligned} \alpha\Lambda_1 \subseteq \Lambda_2 \quad \gamma\Lambda_2 \subseteq \Lambda_1 &\implies \\ \alpha\gamma\Lambda_1 = \Lambda_1 \subseteq \gamma\Lambda_2 \subseteq \Lambda_1 \quad \alpha\gamma\Lambda_2 = \Lambda_2 \subseteq \alpha\Lambda_1 \subseteq \Lambda_2 & \end{aligned}$$

Therefore, $\alpha\Lambda_1 = \Lambda_2$. □

Now, if $\{w_1, w_2\} \subset \mathbb{C}$ defines a lattice Λ in \mathbb{C} , then we can take $\alpha = w_1^{-1}$ so that $\alpha w_1 = 1$. Set $\alpha w_2 := \tau$. We have $\alpha\Lambda = \mathbb{Z} + \tau\mathbb{Z}$. Therefore, every complex torus is isomorphic to $\mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$, where $\text{im } \tau \neq 0$. If $\text{im } \tau < 0$ then $\mathbb{Z} - \tau\mathbb{Z}$ defines the same torus. Thus, every complex torus can be represented by a complex parameter $\tau \in \mathbb{H}$, where \mathbb{H} is the complex upper half plane. This complex torus is denoted \mathcal{C}_τ .

Theorem 2.1.4. *Let $\tau, \sigma \in \mathbb{H}$. Then $\mathcal{C}_\tau \cong \mathcal{C}_\sigma$ if and only if τ and σ lie in the same $SL(2, \mathbb{Z})$ -orbit, where the group acts by Möbius transformations.*

Proof. We need to quotient out by the choice of basis of the lattice Λ , which means taking $\mathbb{H}/GL(2, \mathbb{Z})$ for some group action of $GL(2, \mathbb{Z})$. Any element of this group has determinant ± 1 , and we can restrict to $SL(2, \mathbb{Z})$ because we have already made the identification $\mathbb{Z} \oplus \tau\mathbb{Z} = \mathbb{Z} \oplus -\tau\mathbb{Z}$. The action of $SL(2, \mathbb{Z})$ on \mathbb{H} is given by $\mathbb{Z}^2 \mapsto \mathbb{C}$, determined by $(1, 0) \mapsto 1$ and $(0, 1) \mapsto \tau$. Precomposing this with an element of $SL(2, \mathbb{Z})$ acting on \mathbb{Z}^2 , one finds

$$\tau \mapsto \frac{a\tau + b}{c\tau + d}$$

where a, b, c, d are the matrix entries. In other words, $SL(2, \mathbb{Z})$ acts by Möbius transformations. □

Thus, the coordinate τ on $\mathbb{H}/SL(2, \mathbb{Z})$ parameterises isomorphism classes of complex structures on the topological torus $S^1 \times S^1$.

2.2. HOLOMORPHIC VECTOR BUNDLES

More generally, for a given smooth manifold X , one can always define the complex moduli space of X .

Definition 2.1.5. The complex moduli space of a smooth manifold X is defined as

$$\mathcal{M}_c(X) = \{\text{Integrable complex structures } J \text{ on } X\} / \sim$$

where $J_0 \sim J_1 \iff J_0 = \varphi^* J_1$ for some $\varphi \in \text{Diff}(X)$.

We would like this to be a projective complex manifold itself, but this is typically not the case. So instead of studying the global object $\mathcal{M}_c(X)$, one typically studies local deformations of a given complex structure. This is done by establishing the existence of a first order neighbourhood $\text{Def}(\mathcal{X})$ of $\mathcal{X} = (X, J)$ in $\mathcal{M}_c(X)$, together with a morphism π called the universal deformation $\pi : \mathcal{U} \rightarrow \text{Def}(\mathcal{X})$. It satisfies the property that, given any proper holomorphic submersion $\mathcal{S} \rightarrow \mathcal{B}$ with $\pi^{-1}(b_0) \cong \mathcal{X}$ for some $b_0 \in \mathcal{B}$, there exists a uniquely defined morphism $\Phi : \mathcal{B} \supset V \rightarrow \text{Def}(\mathcal{X})$ such that $\mathcal{S}|_V \cong \Phi^* \mathcal{U}$ for some sufficiently small neighbourhood V of b_0 . In general, $\text{Def}(\mathcal{X})$ will be a germ of a complex analytic space. For Calabi-Yau manifolds, one can show the following.

Theorem 2.1.6 (Bogomolov-Tian-Todorov [21]). *Let \mathcal{X} be a Calabi-Yau manifold with no global holomorphic vector fields. Then $\text{Def}(\mathcal{X})$ exists, and it is a germ of a complex manifold, with tangent space $H^1(\mathcal{X}, T\mathcal{X})$.*

Here, $H^1(\mathcal{X}, T\mathcal{X})$ may be interpreted as sheaf cohomology of the sheaf of holomorphic vector fields, or as the Dolbeault cohomology group to be discussed later. In either case, the point is that the complex manifolds which will be of interest to us have a smooth complex moduli space of dimension $\dim_{\mathbb{C}} \mathcal{M}_c(X) = \dim_{\mathbb{C}} H^1(\mathcal{X}, T\mathcal{X})$. However, $\mathcal{M}_c(X)$ may not be compact, let alone projective. Moduli spaces are very intricate objects and we shall not discuss these intricacies presently. We refer the reader to [21, 20, 24] and references therein.

2.2 Holomorphic Vector Bundles

Once we have a smooth manifold, most of differential geometry is done by considering appropriate vector bundles and sections thereof. The natural analogue is the notion of a holomorphic vector bundle, which will be very important for us.

Definition 2.2.1. Let \mathcal{X} be a complex manifold. A holomorphic vector bundle \mathcal{E} over \mathcal{X} consists of an open cover $\{U_\alpha\}$ together with holomorphic maps $\varphi_{\alpha\beta} : U_{\alpha\beta} \rightarrow \text{GL}(n, \mathbb{C})$ for all intersections $U_{\alpha\beta}$, such that the cocycle condition $\varphi_{\alpha\beta} \circ \varphi_{\beta\gamma} \circ \varphi_{\gamma\alpha} = \text{id}$ is satisfied on triple intersections, and $\varphi_{\alpha\alpha} = \text{id}$.

²Observe that $\text{GL}(n, \mathbb{C}) \subseteq \mathbb{C}^{n^2}$ is a complex manifold, since it is an open subset of a complex vector space.

2.2. HOLOMORPHIC VECTOR BUNDLES

Evidently, forgetting about the complex structure, every holomorphic vector bundle is a smooth vector bundle. The converse is not true. We denote the smooth complex vector bundle which underlies a holomorphic vector bundle \mathcal{E} by E . As for smooth vector bundles, all the natural linear algebra operations apply to holomorphic vector bundles to produce holomorphic vector bundles. We are referring to Hom-bundles, dual bundles, wedge powers, tensor products, and so forth. With this in mind, let us give some examples of holomorphic vector bundles.

Example 11. Consider the complex manifold $\mathbb{C}\mathbb{P}^n$, and take the product $\mathbb{C}\mathbb{P}^n \times \mathbb{C}^{n+1}$. View elements of $\mathbb{C}\mathbb{P}^n$ as lines $\ell \subset \mathbb{C}^{n+1}$ and define

$$\mathcal{O}(-1) := \{(\ell, w) \in \mathbb{C}\mathbb{P}^n \times \mathbb{C}^{n+1} \mid w \in \ell\}$$

This associates a complex line to each point $\ell \in \mathbb{C}\mathbb{P}^n$, and it is easily verified in local coordinates that this defines a sub-bundle of the trivial bundle $\mathbb{C}\mathbb{P}^n \times \mathbb{C}^{n+1}$. We leave it as an exercise to the reader to verify that the transition functions are in fact given by $\varphi_{ij} = z_j/z_i$. Its dual bundle is denoted by $\mathcal{O}(1)$, whose transition functions are then $\varphi_{ij} = z_i/z_j$. We denote $\mathcal{O}(k) = \mathcal{O}(1)^{\otimes k}$, with negative powers obviously being tensor powers of the dual bundle instead.

One can show that every line bundle $\mathbb{C}\mathbb{P}^n$ is isomorphic to $\mathcal{O}(k)$ for some k . By using the definition in terms of cocycles, it is easy to see that $\mathcal{O}(k) \otimes \mathcal{O}(l) = \mathcal{O}(k+l)$, and also $\mathcal{O}(0) = \mathcal{O}$, the trivial bundle. As such, the line bundles on $\mathbb{C}\mathbb{P}^n$ form a group isomorphic to \mathbb{Z} . More generally, the group of line bundles on a complex manifold (or variety) is called the Picard group, denoted by $\text{Pic}(\mathcal{X})$. The group operation is the tensor product, for which taking the dual bundle defines an inverse.

Definition 2.2.2. Let \mathcal{X} be a complex manifold with underlying smooth manifold X . We define $T\mathcal{X} := TX^{1,0}$ to be the holomorphic tangent bundle, $K_{\mathcal{X}} := \wedge^n T^*\mathcal{X}$ to be the canonical bundle of \mathcal{X} , and $\mathcal{O}_{\mathcal{X}} := X \times \mathbb{C}$ to be the trivial holomorphic vector bundle.

Every holomorphic vector bundle comes with a canonical differential operator.

Theorem 2.2.3. *Let $\mathcal{E} \rightarrow \mathcal{X}$ be a holomorphic vector bundle with smooth vector bundle E . There exists a natural operator*

$$\bar{\partial} : \Omega^{p,q}(X, E) \rightarrow \Omega^{p,q+1}(X, E)$$

called the Dolbeault operator, which squares to zero.

The Dolbeault operator is defined locally as $\bar{\partial}$ acting on vector valued functions. This is well-defined globally because the transition functions are holomorphic, so the $\bar{\partial}$ -operator kills this term off. We leave the details as an exercise.

Definition 2.2.4. Let $\mathcal{E} \rightarrow \mathcal{X}$ be a holomorphic vector bundle, and let $\eta \in \Omega^{p,0}(X, E)$.

2.2. HOLOMORPHIC VECTOR BUNDLES

Then η is called holomorphic if $\bar{\partial}\eta = 0$. The holomorphic sections of \mathcal{E} yield a sheaf, denoted by the same symbol. Its sections will be denoted $\Gamma(\mathcal{X}, \mathcal{E})$ or $H^0(\mathcal{X}, \mathcal{E})$.³

In particular, if $\mathcal{E} = \mathcal{O}_{\mathcal{X}}$ and $p = 0$, we recover the definition of a holomorphic function, since $\bar{\partial}f = 0$ if and only if the Cauchy-Riemann equations are satisfied. Holomorphic sections of $\mathcal{O}(k)$ over $\mathbb{C}\mathbb{P}^n$ may be identified with polynomials of degree k in n variables. After homogenisation, these correspond to homogeneous degree k polynomials in $n + 1$ variables, and so projective algebraic varieties are cut out by sections of various $\mathcal{O}(k)$ over $\mathbb{C}\mathbb{P}^n$.

A morphism of holomorphic vector bundles is a holomorphic section of $\mathcal{H}om(\mathcal{E}, \mathcal{F})$, and an isomorphism of holomorphic vector bundles is an invertible morphism. Importantly, there are holomorphic vector bundles which are isomorphic as smooth vector bundles, but not as holomorphic vector bundles, as we now illustrate. We will use the notion of a meromorphic section of a line bundle, which is defined in local coordinates by the same conditions as a meromorphic function $\mathbb{C}^n \rightarrow \mathbb{C}$. Well-definedness is non-trivial, but we will not digress into this. See [20].

Example 12. Let \mathcal{X} be a complex torus and $z_0 \in \mathcal{X}$. Let $U_0 = \mathcal{X} \setminus \{z_0\}$ and U_1 a small disk centered at z_0 . Define a function $g : U_0 \cap U_1 \cong \mathbb{D} \setminus \{0\} \rightarrow \mathbb{C}^\times$ by $g(z) = 1/z$. Declare \mathcal{L}_{z_0} to be the line bundle associated to this cocycle. Then \mathcal{L}_{z_0} is a holomorphic line bundle, which admits a holomorphic section with a simple zero at z_0 , and no other zeroes or poles.

Proof. The transition function $g(z)$ is clearly holomorphic on its domain, so \mathcal{L}_{z_0} is a holomorphic line bundle over \mathcal{X} . Define a section $s \in H^0(\mathcal{X}, \mathcal{L}_{z_0})$ by defining it on the cover $\{U_0, U_1\}$. We set $s|_{U_0} = 1$ and $s|_{U_1} = z$. On the overlap, we have $z \mapsto z/z = 1$ and so these locally defined functions glue to give a section $s \in H^0(\mathcal{X}, \mathcal{L}_{z_0})$ with the desired property. \square

Given \mathcal{L}_{z_0} as above, it is clear that $\mathcal{L}_{z_0}^*$ admits a section with a single simple pole at z_0 , and no other zeroes and poles. Namely by taking the section $\sigma \in H^0(\mathcal{X}, \mathcal{L}_{z_0}^*)$ such that $\sigma(s) = 1$. Now we can construct a non-trivial holomorphic line bundle with $c_1(\mathcal{L}) = 0$, as follows. Take $z_0, z_1 \in \mathcal{X}$, $z_0 \neq z_1$ and consider (\mathcal{L}_{z_0}, s_0) and (\mathcal{L}_{z_1}, s_1) , where the s_i are sections as in the proof. Then we have a section $s_0 \otimes \sigma_1 \in \mathcal{L}_{z_0} \otimes \mathcal{L}_{z_1}^*$ which has a simple pole at z_1 , and a simple zero at z_0 . Supposing that $\mathcal{L}_{z_0} \otimes \mathcal{L}_{z_1}^*$ is trivial, we can identify $s_0 \otimes \sigma_1$ with a meromorphic function. As we know from complex analysis, we can identify meromorphic functions with holomorphic maps $f : \mathcal{X} \rightarrow \mathbb{C}\mathbb{P}^1$, mapping singularities to the point at infinity. We get a map of degree one $f : \mathcal{X} \rightarrow \mathbb{C}\mathbb{P}^1$, which has to be an isomorphism (see [25]). However, this

³We make an exception for holomorphic sections of $\wedge^k T^* \mathcal{X}$. Its sheaf of holomorphic sections will be denoted by $\Omega_{\mathcal{X}}^k$.

2.3. KÄHLER MANIFOLDS

is impossible, as can already be seen at the level of topological spaces. Consequently, $\mathcal{L} = \mathcal{L}_{z_0} \otimes \mathcal{L}_{z_1}^*$ is non-trivial as a holomorphic bundle. However, it is clear that $L_{z_0} \cong L_{z_1}$ as smooth vector bundles by looking at their Chern classes. Therefore $c_1(L) = 0$, implying that L is trivial as a smooth line bundle. Thus, we have found two vector bundles, \mathcal{L} and $\mathcal{O}_{\mathcal{X}}$, which are not isomorphic as holomorphic vector bundles, although they are isomorphic as smooth vector bundles.

In conclusion, a given smooth vector bundle $E \rightarrow X$ may admit many different holomorphic structures. The space of all holomorphic structures on E can be expressed as a quotient space, just as the complex moduli space of a smooth manifold. We will elaborate on this in the second part of the text.

An important fact that we will use many times throughout the text, is that holomorphic vector bundles have an intrinsically defined notion of cohomology, namely their Dolbeault cohomology groups.

Definition 2.2.5. Let $\mathcal{E} \rightarrow \mathcal{X}$ be a holomorphic vector bundle. The Dolbeault cohomology groups $H^{p,q}(\mathcal{X}, \mathcal{E})$ are defined as

$$H^{p,q}(\mathcal{X}, \mathcal{E}) := \frac{\ker(\bar{\partial} : \Omega^{p,q}(X, E) \rightarrow \Omega^{p,q+1}(X, E))}{\text{im}(\bar{\partial} : \Omega^{p,q-1}(X, E) \rightarrow \Omega^{p,q}(X, E))}$$

By standard results from sheaf cohomology⁴, we have an isomorphism between the Dolbeault cohomology groups, and the sheaf cohomology groups:

$$H^q(\mathcal{X}, \Omega_{\mathcal{X}}^p \otimes \mathcal{E}) \cong H^{p,q}(\mathcal{X}, \mathcal{E})$$

When $\mathcal{E} = \mathcal{O}_{\mathcal{X}}$, we simply denote these groups by $H^{p,q}(\mathcal{X})$.

2.3 Kähler Manifolds

Next, we introduce Kähler manifolds. A Riemannian metric g on a complex manifold is called Hermitian if $g(Jv, Jw) = g(v, w)$.

Definition 2.3.1. A Kähler manifold (\mathcal{X}, g) is a complex manifold \mathcal{X} together with a Hermitian metric g such that $\omega(v, w) := g(Jv, w)$ defines a symplectic form on \mathcal{X} .

The symplectic form obtained from g in this way is called the Kähler form of the Kähler metric g .

Proposition 2.3.2. *The Kähler form ω is of type $(1, 1)$.*

⁴See D.2.3.

2.3. KÄHLER MANIFOLDS

Proof. It suffices to check this for a vector space with inner product g and complex structure J . We note that $\wedge^2 V^* \otimes \mathbb{C} = \wedge^{2,0} V^* \oplus \wedge^{1,1} V^* \oplus \wedge^{0,2} V^*$. By definition and \mathbb{C} -linearity, $\omega \in \wedge^{2,0} V^* \oplus \wedge^{0,2} V^* \iff J^* \omega = -\omega$, and $\omega \in \wedge^{1,1} V^* \iff J^* \omega = \omega$. Thus, it suffices to verify that $J^* \omega = \omega$. This is straightforward when we use the characterisation $\omega(v, w) = g(Jv, w)$, and the fact that g is Hermitian.

$$(J^* \omega)(v, w) = \omega(Jv, Jw) = g(J^2 v, Jw) = (J^* g)(Jv, w) = g(Jv, w) = \omega(v, w)$$

□

Example 13. Let U_i denote the standard affine opens of $\mathbb{C}\mathbb{P}^n$. Define

$$\omega_i = \frac{i}{2\pi} \partial \bar{\partial} \log(1 + |z|^2)$$

A routine verification confirms that these ω_i agree on overlaps, to give a globally defined 2-form ω , which is closed since $\partial^2 = 0 = \bar{\partial}^2$. Thus, it remains to verify that

$$\begin{aligned} \partial \bar{\partial} \log(1 + |z|^2) &= \frac{\sum dz^k \wedge d\bar{z}^k}{1 + |z|^2} - \frac{(\sum \bar{z}^k dz^k) \wedge (\sum z^k d\bar{z}^k)}{(1 + |z|^2)^2} = \\ &= \frac{1}{(1 + |z|^2)^2} \sum ((1 + |z|^2) \delta_{ij} - \bar{z}_i z_j) dz^i \wedge d\bar{z}^j \end{aligned}$$

defines a Riemannian metric. That is, the matrix $((1 + |z|^2) \delta_{ij} - \bar{z}_i z_j)_{ij} := h$ should be positive definite. This is done using Cauchy-Schwarz. Namely, for $0 \neq v \in \mathbb{C}^n$ we have⁵

$$h(v, v) = v^T h \bar{v} = |v|^2 + |z|^2 |v|^2 - v^T \bar{z} z^T \bar{v} = |v|^2 + |z|^2 |v|^2 - |\langle z, v \rangle|^2 > 0$$

The matrix h clearly defines a Hermitian metric. Thus, $\mathbb{C}\mathbb{P}^n$ is indeed a Kähler manifold.

Example 14. Every complex submanifold $\iota : \mathcal{Y} \hookrightarrow \mathcal{X}$ of a Kähler manifold is again a Kähler manifold. Indeed, if g is a Kähler metric on \mathcal{X} , then $\iota^* g$ is a Hermitian metric on \mathcal{Y} . This holds since ι is an immersion. Furthermore, $d\iota^* \omega = \iota^* d\omega = 0$, so that $\iota^* g$ is Kähler. This means that every projective manifold $\iota : \mathcal{X} \hookrightarrow \mathbb{C}\mathbb{P}^n$ is a Kähler manifold by the previous example.

It turns out that every (closed) Riemann surface is projective, and thus Kähler. However, this is not the easiest way to prove that a Riemann surface is a Kähler manifold.

⁵We let $|\cdot|$ and $\langle \cdot, \cdot \rangle$ denote the standard norm/Hermitian inner product on \mathbb{C}^n .

2.4. CALABI-YAU MANIFOLDS

Indeed, our construction of a complex structure on an orientable surface already tells us how to do this.

Example 15. Recall that we construct a complex structure on an oriented surface using a Riemannian metric, namely the Hodge star isomorphism. By construction, the Riemannian metric is Hermitian with respect to the complex structure, since the Hodge star is an isometry. Since a Riemann surface has real dimension 2, the associated 2-form is trivially closed, and hence symplectic, giving the Riemann surface a Kähler metric.

Example 16. Every complex torus \mathbb{C}^n/Λ is a Kähler manifold when equipped with a flat metric.

A rather remarkable fact about Kähler manifolds is that the (p, q) decomposition of forms descends to cohomology. This is known as the Hodge decomposition theorem. For background, see [22] which covers the topic extensively.

Theorem 2.3.3 (Hodge Decomposition Theorem). *Let \mathcal{X} be a compact Kähler manifold. Then there is a natural decomposition*

$$H^k(X, \mathbb{C}) = \bigoplus_{p+q=k} H^{p,q}(\mathcal{X})$$

By the universal coefficients theorem, we have $H^k(X, \mathbb{R}) \subset H^k(X, \mathbb{C})$. We then set $H^{p,q}(\mathcal{X}, \mathbb{R}) = H^{p,q}(\mathcal{X}) \cap H^{p+q}(X, \mathbb{R})$. The Kähler class refers to the cohomology class $[\omega] \in H^{1,1}(X, \mathbb{R})$, but we simply denote it by ω . The set of all cohomology classes which may arise as a Kähler class is called the Kähler cone (since it is invariant under scaling by $\mathbb{R}_{>0}$ for obvious reasons), and it is an open subset of $H^{1,1}(\mathcal{X}, \mathbb{R})$.

2.4 Calabi-Yau Manifolds

Definition 2.4.1. An almost Calabi-Yau manifold (\mathcal{X}, g) is a Kähler manifold (\mathcal{X}, g) whose canonical bundle is holomorphically trivial: $K_{\mathcal{X}} \cong \mathcal{O}_{\mathcal{X}}$.

The triviality of $K_{\mathcal{X}}$ is equivalent to the existence of a holomorphic $(n, 0)$ -form which is nowhere vanishing. This form $\Omega \in \Omega_{\mathcal{X}}^n(X)$ will be called a holomorphic volume form. By Yau's resolution of the Calabi conjecture, every almost Calabi-Yau manifold admits a unique Ricci flat Kähler metric which represents the Kähler class.

Definition 2.4.2. A Calabi-Yau manifold (\mathcal{X}, g, Ω) is a Ricci flat Kähler manifold together with a holomorphic volume form. The holomorphic volume form is normalised up to S^1 -scaling by

$$\omega^n = n!(-1)^{n(n-1)/2} \Omega \wedge \bar{\Omega}$$

2.4. CALABI-YAU MANIFOLDS

Unless stated otherwise, Calabi-Yau manifolds will be assumed to be compact.

Example 17. Every complex torus \mathbb{C}^n/Λ with a flat metric is Calabi-Yau. A holomorphic volume form is given by $dz_1 \wedge \cdots \wedge dz_n$.

Up to diffeomorphism, compact Calabi-Yau manifolds are quite rare. In complex dimension 1, the only example of a Calabi-Yau manifold is topologically a torus \mathbb{C}/Λ . In complex dimension 2, one has \mathbb{C}^2/Λ and the $K3$ -surface as smooth manifolds which admit a Calabi-Yau structure. In complex dimension 3, there is a much larger number (possibly infinite - this is unknown according to [26]) of Calabi-Yau manifolds that are topologically distinct. We will construct some examples further below (2.4.1). It is important to note that, just as for complex tori, there are non-isomorphic Calabi-Yau structures on the same smooth manifold.

One can also consider Calabi-Yau manifolds in the strict sense. What this means is the following. The metric g gives a holonomy group (see the Appendix A.3) via the Levi-Civita connection. If \mathcal{X} is a Calabi-Yau manifold, then $\text{Hol}(g) \subseteq \text{SU}(n)$. When we have equality, we say that \mathcal{X} is a strict Calabi-Yau manifold. Thus, \mathbb{C}^n/Λ with its flat metric is not a strict Calabi-Yau (unless $n = 1$), as the holonomy is trivial. This can be translated into an algebraic condition (see [21]). The condition for a Calabi-Yau manifold to be strict is equivalent to $H^q(\mathcal{X}, \mathcal{O}_{\mathcal{X}}) = 0$ for $0 < q < n$. Sheaf cohomology is defined on algebraic varieties as well.

Definition 2.4.3. A Calabi-Yau variety is a projective variety \mathcal{X} such that $K_{\mathcal{X}} \cong \mathcal{O}_{\mathcal{X}}$. If the sheaf cohomology groups $H^q(\mathcal{X}, \mathcal{O}_{\mathcal{X}})$ vanish for $0 < q < n$, then \mathcal{X} is called a strict Calabi-Yau variety.

For us, Calabi-Yau varieties will always be smooth, and we will refer to them simply as Calabi-Yau manifolds. This is because strict Calabi-Yau manifolds of dimension ≥ 3 are always algebraic, so we can consider them as projective algebraic varieties, or as projective manifolds.

2.4.1 Constructing Calabi-Yau Threefolds

We will now construct some Calabi-Yau manifolds, although we will not construct their Ricci flat Kähler metrics (also known as Calabi-Yau metrics). Indeed, there are no explicitly known Calabi-Yau metrics on compact Calabi-Yau manifolds (apart from flat metrics on tori), although their existence is guaranteed by Yau's resolution of the Calabi conjecture.

Theorem 2.4.4 (Calabi-Yau [21]). *Let (\mathcal{X}, g') be a compact Kähler manifold with $c_1(X) = 0$. Then there exists a unique Ricci flat Kähler metric g on \mathcal{X} in the Kähler class of g' . We call g the Calabi-Yau metric of \mathcal{X} with Kähler class ω .*

2.4. CALABI-YAU MANIFOLDS

To check whether the manifolds we construct are distinct, we use their Hodge numbers. When \mathcal{X} is compact, the vector spaces $H^{p,q}(\mathcal{X})$ are finite dimensional vector spaces, and we define the integers $h^{p,q}(\mathcal{X}) := \dim_{\mathbb{C}} H^{p,q}(\mathcal{X})$. They are called the Hodge numbers of \mathcal{X} , and they are deformation invariant. That is to say, when we have a holomorphic family $\mathcal{S} \rightarrow \mathcal{B}$ of Kähler manifolds which is a proper submersion onto some base space \mathcal{B} , then the Hodge numbers of the fibres will agree. Recall that $b_k(X) := \dim_{\mathbb{R}} H_k(X, \mathbb{R})$ are called the Betti numbers.

Theorem 2.4.5. [20] *Let \mathcal{X} be a Kähler manifold. Then $h^{p,q}(\mathcal{X})$ satisfy the following relations:*

1. $h^{p,q}(\mathcal{X}) = h^{q,p}(\mathcal{X})$ (conjugate symmetry)
2. $h^{p,q}(\mathcal{X}) = h^{n-p, n-q}(\mathcal{X})$ (Serre duality)
3. $h^{p,q}(\mathcal{X}) = h^{n-q, n-p}(\mathcal{X})$ (the Hodge star isomorphism)

Furthermore, $b_k(X) = \sum_{p+q=k} h^{p,q}(\mathcal{X})$ (the Hodge decomposition).

With this in mind, let us start constructing some Calabi-Yau manifolds.

Example 18. The original example of mirror symmetry was discovered by physicists, and pertained to the Calabi-Yau threefold known as the Fermat quintic. It is defined as the zero locus in $\mathbb{C}\mathbb{P}^4$ of a homogeneous degree 5 polynomial in the same number of variables:

$$\mathcal{Q} := \{(z_0 : z_1 : z_2 : z_3 : z_4) \in \mathbb{C}\mathbb{P}^4 \mid \sum_i z_i^5 = 0\}$$

In other words, \mathcal{Q} is the vanishing locus of a holomorphic section of the bundle $\mathcal{O}(5)$ over $\mathbb{C}\mathbb{P}^4$. We can use this to show that it is a Calabi-Yau manifold. Firstly, we observe that \mathcal{Q} is non-singular because its Jacobian vanishes if and only if $z = 0$ which does not occur in $\mathbb{C}\mathbb{P}^4$. Furthermore, the knowledge that $\mathcal{Q} = Z(s)$ for $s \in H^0(\mathcal{X}, \mathcal{O}(5))$ allows us to calculate the Chern class $c(Q)$ because the normal bundle can be identified with $\mathcal{O}(5)$, i.e. the smooth vector bundle underlying $\mathcal{O}(5)$. Thus, we have an isomorphism $T\mathbb{C}\mathbb{P}^4|_Q \cong TQ \oplus \mathcal{O}(5)|_Q$, which implies that $c(T\mathbb{C}\mathbb{P}^4|_Q) = c(Q)c(\mathcal{O}(5))$. We let x denote the hyperplane class, i.e. the generator of the cohomology ring $H^*(\mathbb{C}\mathbb{P}^4)$. We also denote $x = \iota^*x$ by abuse of notation. By construction, $c(\mathcal{O}(5)) = 1 + 5x$. By elementary arguments, one can check that $c(T\mathbb{C}\mathbb{P}^4) = (1+x)^5$, and in fact $c(T\mathbb{C}\mathbb{P}^n) = (1+x)^{n+1}$. This yields

$$c(T\mathbb{C}\mathbb{P}^4|_Q) = (1+x)^5 = 1 + 5x + 10x^2 + 10x^3 = (1 + c_1(Q) + c_2(Q) + c_3(Q))(1 + 5x) \implies \\ c_1(Q) = 0 \quad c_2(Q) = 10x^2 \quad c_3(Q) = -40x^3$$

So indeed, \mathcal{Q} is a Calabi-Yau manifold, since it is a Kähler manifold with $c_1(Q) = 0$.

2.4. CALABI-YAU MANIFOLDS

From the same computation, we can calculate the Euler characteristic $\chi(Q)$ of Q .

$$\begin{aligned}\chi(Q) &= \langle -40x^3, [Q] \rangle = -40 \cdot 5 = -200 \implies \\ b_0(Q) - b_1(Q) + b_2(Q) - b_3(Q) + b_4(Q) - b_5(Q) + b_6(Q) &= \\ 1 - 0 + 1 - \dim_{\mathbb{R}} H^3(Q, \mathbb{R}) + 1 - 0 + 1 &= -200\end{aligned}$$

Thus, we must have $\dim_{\mathbb{R}} H^3(Q, \mathbb{R}) = 204$. Since $h^{3,0}(Q) = h^{0,3}(Q) = 1$, it follows that $h^{2,1}(Q) = h^{1,2}(Q) = 101$. Here, we have used the Lefschetz hyperplane theorem to conclude that $H_2(Q, \mathbb{R}) \cong \mathbb{R}$, which implies that $b_2(Q) = h^{1,1}(Q) = 1$. The same theorem tells us that Q is a strict Calabi-Yau manifold, i.e. $h^{1,0}(Q) = h^{2,0}(Q) = 0$. In conclusion, we have proved the following.

Theorem 2.4.6. *The quintic threefold Q is a strict Calabi-Yau manifold with Hodge diamond⁶ given by*

$$\begin{array}{cccc} & & & & 1 \\ & & & & 0 & & 0 \\ & & & & 0 & & 1 & & 0 \\ & & & & 1 & & 101 & & 101 & & 1 \\ & & & & 0 & & 1 & & 0 \\ & & & & 0 & & 0 \\ & & & & 1 \end{array}$$

The above computations make it evident that the machinery from algebraic geometry is very useful, and can easily be generalised to e.g. complete intersections in $\mathbb{C}\mathbb{P}^N$. We will construct Calabi-Yau threefolds $\mathcal{X} \subseteq \mathbb{C}\mathbb{P}^5$ as complete intersections explicitly, and mention how to generalise this construction to $\mathbb{C}\mathbb{P}^N$ more generally.

Example 19. If we want a threefold in $\mathbb{C}\mathbb{P}^5$, we will need $\mathcal{X} = Z(s_1) \cap Z(s_2)$ for some transversal intersection. Thus, we will take $s_i \in \Gamma(\mathbb{C}\mathbb{P}^5, \mathcal{O}(n_i))$ and assume that \mathcal{X} is a complete intersection of these two homogeneous polynomials. Denote $\mathcal{X}_i = Z(s_i)$, so that $\mathcal{X} = \mathcal{X}_1 \cap \mathcal{X}_2$. Then we have a Cartesian diagram:

$$\begin{array}{ccc} \mathcal{X} & \xleftarrow{j} & \mathcal{X}_2 \\ \downarrow i & & \downarrow \iota \\ \mathcal{X}_1 & \xleftarrow{\kappa} & \mathbb{C}\mathbb{P}^5 \end{array}$$

⁶See 4.1 for what we mean by the Hodge diamond.

2.4. CALABI-YAU MANIFOLDS

By the same argument we employed for the quintic, the ideal sheaf of \mathcal{X}_1 is $\mathcal{O}(-n_1)$ which yields an exact sequence of sheaves

$$0 \rightarrow \mathcal{O}(-n_1) \rightarrow \mathcal{O} \rightarrow \kappa_* \mathcal{O}_{\mathcal{X}_1} \rightarrow 0$$

Pulling back by ι is exact (as can be verified locally), which yields an exact sequence

$$0 \rightarrow \mathcal{O}(-n_1)|_{\mathcal{X}_2} \rightarrow \iota^* \mathcal{O}_{\mathcal{X}} \rightarrow \iota^* \kappa_* \mathcal{O}_{\mathcal{X}_1} \rightarrow 0$$

Now, we have $\iota^* \mathcal{O}_{\mathcal{X}} = \mathcal{O}_{\mathcal{X}_2}$, and the Cartesian diagram yields $\iota^* \kappa_* \mathcal{O}_{\mathcal{X}_1} = j_* \iota^* \mathcal{O}_{\mathcal{X}_1} = j_* \mathcal{O}_{\mathcal{X}}$. Therefore, we obtain an exact sequence of sheaves

$$0 \rightarrow \mathcal{O}(-n_1)|_{\mathcal{X}_2} \rightarrow \mathcal{O}_{\mathcal{X}_2} \rightarrow j_* \mathcal{O}_{\mathcal{X}} \rightarrow 0$$

Thus, by definition, the ideal sheaf of $\mathcal{X} \xrightarrow{j} \mathcal{X}_2$ is $\mathcal{O}(-n_1)|_{\mathcal{X}_2}$. Using this, we can once again calculate $c_1(X)$, because $c(X_2) = c(X)c(\mathcal{O}(n_1)|_{\mathcal{X}_2})$ via the splitting $TX_2|_X \cong TX \oplus \mathcal{O}(n_1)|_X$. Clearly, a computation analogous to the one we carried out for the quintic yields

$$c(X_2) = 1 + (6 - n_2)x + (15 - 6n_2 + n_2^2)x^2 + (20 - 15n_2 + 6n_1^2 - n_2^3)x^3 + c_4(X_2)$$

We omit the fourth Chern class because it will not be relevant for us. We can now calculate:

$$\begin{aligned} c(X_2) &= c(X)c(\mathcal{O}(n_1)) \iff \\ 1 + (6 - n_2)x + (15 - 6n_2 + n_2^2)x^2 + (20 - 15n_2 + 6n_1^2 - n_2^3)x^3 + c_4(X_2) &= \\ (1 + c_1(X) + c_2(X) + c_3(X))(1 + n_1x) &\iff \\ c_1(X) &= (6 - n_1 - n_2)x \\ c_2(X) &= (15 - 6n_1 - 6n_2 + n_1^2 + n_2^2 + n_1n_2)x^2 \\ c_3(X) &= (-n_1^3 - n_1^2n_2 + 6n_1^2 - n_1n_2^2 + 6n_1n_2 - 15n_1 - n_2^3 + 6n_2^2 - 15n_2 + 20)x^3 \end{aligned}$$

If we wish for \mathcal{X} to be Calabi-Yau, we must select $n_1, n_2 > 0$ such that $n_1 + n_2 = 6$. More generally, if we want to obtain a complete intersection in $\mathbb{C}\mathbb{P}^N$, we must take $N - 3$ sections $s_i \in \Gamma(\mathbb{C}\mathbb{P}^N, \mathcal{O}(n_i))$ such that $\sum n_i = N$. But how do we know that the Calabi-Yau threefolds we obtain in this way are topologically distinct? For this, we calculate some more of the Hodge numbers. It is not difficult to show by induction on the dimension that $h^{1,0}(\mathcal{X}) = h^{2,0}(\mathcal{X}) = 0$ for complete intersections, using the Lefschetz hyperplane theorem. The same argument then yields $h^{1,1}(\mathcal{X}) = 1$. It remains for us to calculate $h^{1,2}(\mathcal{X})$. We do this via the same argument as for the quintic. That is, $c_3(X) = e(X)$, and furthermore, the cup product is Poincaré dual to intersections. Therefore, the Poincaré dual of $[X] = [X_1 \cap X_2]$ is $c_1(X_1) \wedge c_1(X_2) = n_1 n_2 x^2$. We can then calculate:

$$\begin{aligned} \chi(X) &= \langle c_3(X), [X] \rangle = \langle n_1 n_2 x^2 c_3(X), \mathbb{C}\mathbb{P}^5 \rangle = \\ n_1 n_2 (-n_1^3 - n_1^2 n_2 + 6n_1^2 - n_1 n_2^2 + 6n_1 n_2 - 15n_1 - n_2^3 + 6n_2^2 - 15n_2 + 20) \end{aligned}$$

2.4. CALABI-YAU MANIFOLDS

Here, we used that $\langle x^5, \mathbb{C}\mathbb{P}^5 \rangle = 1$. It follows that

$$\begin{aligned} \chi(X) &= b_0(X) - b_1(X) + b_2(X) - b_3(X) + b_4(X) - b_5(X) + b_6(X) = \\ &= 1 - 0 + 1 - b_3(X) + 1 - 0 + 1 = 4 - b_3(X) \iff \\ b_3(X) &= 4 - n_1 n_2 (-n_1^3 - n_1^2 n_2 + 6n_1^2 - n_1 n_2^2 + 6n_1 n_2 - 15n_1 - n_2^3 + 6n_2^2 - 15n_2 + 20) \iff \\ h^{2,1}(\mathcal{X}) &= 1 - n_1 n_2 (-n_1^3 - n_1^2 n_2 + 6n_1^2 - n_1 n_2^2 + 6n_1 n_2 - 15n_1 - n_2^3 + 6n_2^2 - 15n_2 + 20)/2 \end{aligned}$$

Let us examine if we can choose n_1, n_2 such that we obtain a threefold which is manifestly different from the quintic. Clearly, our options are:

$$\begin{aligned} n_1 = 1 \quad n_2 = 5 &\implies h^{2,1}(\mathcal{X}) = 101 \\ n_1 = 2 \quad n_2 = 4 &\implies h^{2,1}(\mathcal{X}) = 89 \\ n_1 = 3 \quad n_2 = 3 &\implies h^{2,1}(\mathcal{X}) = 73 \end{aligned}$$

As such, we have obtained (at least) two new distinct Calabi-Yau threefolds, as complete intersections in $\mathbb{C}\mathbb{P}^5$. Observe that the case $n_1 = 1$ corresponds to $\mathcal{X}_1 \cong \mathbb{C}\mathbb{P}^4$, which results in $\mathcal{X}_1 \cap \mathcal{X}_2 \cong \mathcal{Q}$. It is clear that the procedure may be applied by intersecting $N - 3$ submanifolds in $\mathbb{C}\mathbb{P}^N$, as we remarked earlier. The computation for the Hodge numbers will become increasingly tedious when doing this, but it gives a whole range of Calabu-Yau manifolds.

2.4.2 Calibrations on Calabi-Yau Manifolds

For the Thomas-Yau conjecture, as well as the SYZ approach to mirror symmetry, we will need the notion of a calibration, and the associated calibrated submanifolds.

Definition 2.4.7. A calibrated manifold is a Riemannian manifold (X, g) together with a closed k -form φ such that $0 \leq \varphi|_V \leq \text{Vol}_g|_V$ for all oriented k -dimensional $V \subseteq T_x M$ and all $x \in X$.

The k -form φ is called the calibration.

Definition 2.4.8. Let (X, g, φ) be a calibrated manifold. Then a calibrated submanifold is a submanifold $Y \subseteq X$ such that $\varphi|_{T_x Y} = \text{Vol}_g|_{T_x Y}$ for all $x \in Y$.

The following is proved in [27].

Theorem 2.4.9. *Let (\mathcal{X}, ω) be a Kähler manifold. Then ω^k for $1 \leq k \leq \dim \mathcal{X}$ is a calibration on \mathcal{X} , and the calibrated submanifolds are precisely complex submanifolds $\mathcal{Y} \subseteq \mathcal{X}$. Suppose that (\mathcal{X}, ω) is additionally Calabi-Yau. Then $\text{Re} \exp(i\theta)\Omega$ is a calibration for Lagrangian submanifolds of (X, ω) , and the calibrated submanifolds are called special Lagrangian submanifolds.*

2.4. CALABI-YAU MANIFOLDS

Thus, special Lagrangian submanifolds are Lagrangian submanifolds $L \subset X$ for which there exists a choice of holomorphic volume form such that $\operatorname{Re} \Omega|_L = \operatorname{Vol}_g$. Equivalently, there exists a choice of holomorphic volume form such that $\operatorname{Im} \Omega|_L = 0$. Special Lagrangians are necessarily oriented submanifolds. These special Lagrangian submanifolds are going to be the central objects of interest for the Thomas-Yau conjecture.

Proposition 2.4.10. *Let $Y \subset X$ be a calibrated submanifold. Then Y is volume minimising within its homology class in X .*

Proof. Let $Y \subseteq X$ be a calibrated submanifold and suppose that $[Y] = [Y']$ for some cycle $Y' \subseteq X$. Then by Stokes's theorem,

$$\operatorname{Vol}(Y) = \int_Y \operatorname{Vol}_g = \int_Y \varphi = \int_{Y'} \varphi \leq \int_{Y'} \operatorname{Vol}_g = \operatorname{Vol}(Y')$$

So indeed, calibrated submanifolds are volume minimising in their homology class. \square

The implication is that calibrated submanifolds are solutions to a difficult problem: finding minimal submanifolds. The second fundamental form $A \in \Gamma(Y, \operatorname{Sym}^2(T^*Y) \otimes NY)$ for a submanifold $Y \subset X$ is defined as $\nabla_v^X w = \nabla_v^Y w + A(v, w)$, where ∇^X and ∇^Y are the Levi-Civita connections of the respective metrics and NY is the normal bundle of Y is $TX|_Y$. The mean curvature vector \vec{H} is defined to be the trace of this tensor, so that $\vec{H}_Y \in \Gamma(Y, NY)$.

Definition 2.4.11. A submanifold $Y \subset X$ of a Riemannian manifold (X, g) is called minimal if $\vec{H}_Y = 0$. Equivalently, if it locally minimises the volume functional.

Solving the equations for the minimisation of the volume functional is, in general, a very difficult problem. Since calibrated submanifolds are volume minimising, they are also minimal submanifolds, and so provide a solution to this problem.

Example 20. If $\dim \mathcal{X} = 1$, then special Lagrangian submanifolds are the same as geodesics.

If $\dim \mathcal{X} = 2$, a Lagrangian submanifold Y is special Lagrangian if and only if there exists a complex structure on X which makes Y a complex submanifold. Because of this, special Lagrangians really only become interesting if $\dim \mathcal{X} \geq 3$.

In $\mathcal{X} = \mathbb{C}^n$ with its standard Calabi-Yau structure, let $\vec{\phi} \in \mathbb{R}^n$ be fixed, and consider

$$\Pi^{\vec{\phi}} = \{(e^{i\phi_1} x_1, \dots, e^{i\phi_n} x_n) \in \mathbb{C}^n \mid (x_1, \dots, x_n) \in \mathbb{R}^n\}$$

Then $\operatorname{Re} \Omega|_{\Pi^{\vec{\phi}}} = \pm \cos(\sum \phi_i)$. So for $\Pi^{\vec{\phi}}$ to be special Lagrangian, we need $\pm \cos(\sum \phi_i) = 1 \iff \sum \phi_i \in \pi\mathbb{Z}$. However, the question of when a union of two special Lagrangian planes in \mathbb{C}^n is volume minimising is more subtle, and we will come back to this in the second part of the text, when special Lagrangians will be more at the forefront of our discussion.

Chapter 3

The A_∞ -Categories

We now discuss the A_∞ -categories through which homological mirror symmetry is formulated, based on [5, 28, 29, 30]. We refer to [31] for a discussion of higher category theory in physics. A short summary of the respective categories, for those readers who do not wish to go through the technicalities, can be given as follows.

For an algebraic variety or complex manifold \mathcal{X} , the derived category of coherent sheaves $D^b(\mathcal{X})$ consists of chain complexes of coherent sheaves, i.e. sequences

$$\dots \xrightarrow{d_{i-1}} \mathcal{E}_i \xrightarrow{d_i} \mathcal{E}_{i+1} \xrightarrow{d_{i+1}} \dots$$

with $d_{i+1} \circ d_i = 0$. The morphisms in this category are just the usual morphisms between chain complexes: termwise morphisms of sheaves which intertwine the differentials. The subtlety is that quasi-isomorphisms are inverted in $D^b(\mathcal{X})$. That is to say, whenever $f: \mathcal{E}_\bullet \rightarrow \mathcal{F}_\bullet$ induces an isomorphism between the cohomology sheaves of the respective complexes, it is an isomorphism in $D^b(\mathcal{X})$.

A Lagrangian submanifold $L \subset X$ of a Calabi-Yau manifold \mathcal{X} is called graded if it comes equipped with a function $\theta: L \rightarrow \mathbb{R}$ such that $\text{Vol}_g = \text{Re} \exp(i\theta)\Omega|_L$ where Ω is a holomorphic volume form. There is a topological obstruction to the existence of this function, called the Maslov class of L . The objects of the Fukaya category consist of graded Lagrangian submanifolds (L, θ) , together with a flat unitary connection on a line bundle $E \rightarrow L$. The morphism space of two objects (L_0, ∇_0) and (L_1, ∇_1) is obtained from $CF^*(L_0, L_1)$, the Floer chain complex of L_0 and L_1 , and the parallel transport of the connections, by viewing the vector space generated by the intersection points $L_0 \cap L_1$ as $\mathbb{C}\langle L_0 \cap L_1 \rangle = \bigoplus_{p \in L_0 \cap L_1} \text{Hom}_{\mathbb{C}}(E_{0,p}, E_{1,p})$. It turns out that the morphisms which one gets from these spaces are not associative, but they are associative up to higher homotopy. This is what an A_∞ -category captures, see the appendix. The derived Fukaya category can then be obtained from the Fukaya category through a procedure outlined in the appendices.

The physical interpretation for both of these A_∞ -categories will be discussed when we arrive at homological mirror symmetry. We also note that in the appendix B, we recall the basics of sheaf theory, particularly coherent sheaves. In the appendix D, we recall the basics of homological algebra (e.g. abelian categories, additive functors, the cohomology functor, triangulated categories, etc.), which are necessary to understand the technical details of the constructions below.

3.1 The Derived Category of Coherent Sheaves

In the appendices, we recall the definition of an abelian category, the particular example of the category of coherent sheaves, and the definition of the derived category of an abelian category. See the appendix D.

Definition 3.1.1. Let \mathcal{X} be a complex manifold or variety. Then we denote by $D^b(\mathcal{X}) := D^b(\text{Coh}(\mathcal{X}))$ the bounded derived category of coherent sheaves on \mathcal{X} .

We recall that the objects in this category are chain complexes of coherent sheaves, and we invert quasi-isomorphisms.

We also recall that the derived category is not an abelian category, although it is constructed from an abelian category. Instead, $D^b(\mathcal{X})$ is a triangulated category, with distinguished triangles (also called exact triangles) as substitutes for short exact sequences. We refer to the appendices for details, since the technicalities will not be too important for our purposes.

Example 21. There is a relatively explicit description of $D^b(\mathbb{C}\mathbb{P}^n)$. Define $\mathcal{A}_{[0,k]}$ to be the category of $\mathcal{O}_{\mathbb{C}\mathbb{P}^m}$ -modules whose objects are isomorphic to finite direct sums $\mathcal{O}(-r_1) \oplus \cdots \oplus \mathcal{O}(-r_m)$, with $0 \leq r_i \leq k$. Then a theorem of Beilinson is that

$$D^b(\mathbb{C}\mathbb{P}^n) \cong K(\mathcal{A}_{[0,n]})$$

where $K(\mathcal{A})$ is the homotopy category of bounded complexes of $\mathcal{A}_{[0,n]}$.

Example 22. When \mathcal{X} is a Riemann surface, all indecomposable coherent sheaves on it are either indecomposable vector bundles, or torsion sheaves supported at a single point $x \in \mathcal{X}$. Let \mathcal{A} denote the full subcategory of $D^b(\mathcal{X})$ whose objects are finite direct sums of indecomposable sheaves in each degree. Then in [32] it is shown that

$$D^b(\mathcal{X}) \cong \mathcal{A}$$

Let $\mathcal{X} = \mathbb{C}/\Lambda$. Torsion sheaves supported at points are just specified by a point $x \in \mathcal{X}$, a finite dimensional \mathbb{C} -vector space, and a nilpotent endomorphism. Holomorphic bundles on an elliptic curve are also fully classified, see [33]. Serre duality on an elliptic curve yields

$$\text{Hom}_{D^b(\mathcal{X})}(\mathcal{E}, \mathcal{F}) = \text{Hom}(\mathcal{E}, \mathcal{F}) \quad \text{Hom}_{D^b(\mathcal{X})}(\mathcal{E}, \mathcal{F}[1]) = \text{Ext}^1(\mathcal{E}, \mathcal{F}) \cong \text{Hom}(\mathcal{F}, \mathcal{E})^*$$

So the objects of $D^b(\mathcal{X})$ can be described as chain complexes of direct sums of vector bundles and torsion sheaves supported at a point in each degree, and the morphism spaces can be calculated quite explicitly, making $D^b(\mathcal{X})$ relatively hands-on.

3.1. THE DERIVED CATEGORY OF COHERENT SHEAVES

We now give some geometric motivation for doing this. As a first guess, maybe the derived category of coherent sheaves classifies an algebraic variety, in the same way that the chain complex of singular (co)chains classifies a simplicial complex up to homotopy equivalence. It turns out that, in some cases, the bounded derived category is indeed a complete invariant of the algebraic variety, as we mention below. Moreover, the derived category allows us to express dualities like the famous Fourier-Mukai transform.

Theorem 3.1.2 (Mukai [34]). *Let \mathcal{X} be an abelian variety and let $\widehat{\mathcal{X}}$ be its dual. Then*

$$D^b(\mathcal{X}) \cong D^b(\widehat{\mathcal{X}})$$

The isomorphism is constructed as follows. Consider the product $\mathcal{X} \times \widehat{\mathcal{X}}$, with projections π and ρ to the respective factors. Then the isomorphism is given by

$$\mathcal{E}_\bullet \mapsto R\rho_*(\pi^*\mathcal{E}_\bullet \otimes^L \mathcal{K}_\bullet)$$

where $\mathcal{E}_\bullet \in D^b(\mathcal{X})$, \otimes^L is the derived tensor product and $\mathcal{K}_\bullet \in D^b(\mathcal{X} \times \widehat{\mathcal{X}})$ is called the kernel, by analogy with the Fourier transform. The Fourier-Mukai transform has close ties to homological mirror symmetry, see [35]. We will touch on this later.

From another point of view, derived categories are important because, in certain cases, they provide a complete invariant of the algebraic variety \mathcal{X} .

Theorem 3.1.3 (Bondal-Orlov [36]). *Let \mathcal{X}_1 and \mathcal{X}_2 be smooth projective varieties. Suppose \mathcal{X}_1 has ample (anti-)canonical divisor. Then $\mathcal{X}_1 \cong \mathcal{X}_2 \iff D^b(\mathcal{X}_1) \cong D^b(\mathcal{X}_2)$ as triangulated categories.*

When the anti-canonical divisor of a variety is ample, it is called a Fano variety. A version of mirror symmetry can be formulated for such varieties, but we will not investigate that. We merely wish to illustrate that the derived category is truly an important object that we can associate to a variety, sometimes so strong that it completely classifies the variety. However, in the present case of Calabi-Yau varieties, this is not the case.

Theorem 3.1.4 (Bridgeland [37]). *If two Calabi-Yau varieties are birationally equivalent, then their derived categories are equivalent.*

Birational equivalence is much weaker than isomorphism. If two varieties are related via a sequence of blow-ups (i.e. we get isomorphic varieties after taking appropriate blow-ups), then they are birationally equivalent. Thus, for Calabi-Yau varieties, we certainly cannot view the derived category as a complete invariant. So why should we care about it, from a geometric perspective?

One such reason is given in [38], which is that we can do intersection theory of subvarieties via sheaves. We let \mathcal{X} denote a smooth projective variety of dimension n , and we let D_1 and D_2 be two divisors on \mathcal{X} . These define homology classes $[D_i] \in H_{n-2}(X, \mathbb{Z})$, where X is the smooth manifold obtained from \mathcal{X} , as in the case where \mathcal{X} denotes a

3.1. THE DERIVED CATEGORY OF COHERENT SHEAVES

complex manifold. Their Poincaré duals are $c_1(L_i)$ for two line bundles \mathcal{L}_i . As we mentioned previously, the cup product is dual to the intersection product, so the intersection class should correspond to the Poincaré dual of $c_1(L_1) \wedge c_1(L_2) \in H^4(X, \mathbb{Z})$, when the intersection is transversal. Alternatively, we can view the intersection as being $Z(s_1|_{D_2})$. This can be translated into the language of sheaves. When considering a subvariety $\mathcal{Y} \subseteq \mathcal{X}$, we will consider coherent sheaves on \mathcal{Y} as coherent sheaves on \mathcal{X} via the push-forward. The section s_1 can be viewed as a morphism of sheaves

$$\mathcal{L}_1^* \xrightarrow{s_1} \mathcal{O}_{\mathcal{X}}$$

and similarly, its restriction to D_2 gives a morphism of sheaves $\mathcal{L}_1^*|_{D_2} \xrightarrow{s_1|_{D_2}} \mathcal{O}_{D_2}$. If we consider the cokernel of this morphism, we obtain a sheaf whose support is precisely the intersection $D_1 \cap D_2$ (which we assume to be transversal for convenience), and is in fact the structure sheaf. In essence, this comes down to the following. We have a short exact sequence of sheaves

$$0 \rightarrow \mathcal{L}_1^* \xrightarrow{s_1} \mathcal{O}_{\mathcal{X}} \rightarrow \mathcal{O}_{D_1} \rightarrow 0$$

which means that we can view \mathcal{O}_{D_1} as the cohomology of the chain complex

$$\cdots \rightarrow 0 \rightarrow \mathcal{L}_1^* \xrightarrow{s_1} \mathcal{O}_{\mathcal{X}} \rightarrow 0 \rightarrow \cdots$$

We take the cochain complex and tensor it with \mathcal{O}_{D_2} to obtain

$$\cdots \rightarrow 0 \rightarrow \mathcal{L}_1^* \otimes \mathcal{O}_{D_2} = \mathcal{L}_1^*|_{D_2} \xrightarrow{s_1 \otimes \text{id}} \mathcal{O}_{\mathcal{X}} \otimes \mathcal{O}_{D_2} = \mathcal{O}_{D_2} \rightarrow 0 \rightarrow \cdots$$

The cohomology of this chain complex is again just the cokernel of the morphism of sheaves, which is the structure sheaf $\mathcal{O}_{D_1 \cap D_2}$. On the other hand, we also know that the intersection corresponds to the fibre product

$$\begin{array}{ccc} D_1 \cap D_2 & \longrightarrow & D_2 \\ \downarrow & & \downarrow t_2 \\ D_1 & \xrightarrow{t_1} & \mathcal{X} \end{array}$$

And consequently, the structure sheaf of $D_1 \cap D_2$ is just $\mathcal{O}_{D_1} \otimes \mathcal{O}_{D_2}$. The assumption that the intersection is transversal is quite important here, because otherwise the result may well be incorrect. For example, the case in which $D := D_1 = D_2$ already provides an example for which we obtain the wrong result. After all, this would yield \mathcal{O}_D as the structure sheaf of the self intersection of D , which we know is not the right way to think about self intersections; instead, we should be intersecting D with a divisor which is linearly equivalent to D , and consider its intersection with this linearly equivalent divisor. This can be thought of as saying that the intersection product is a multiplication map on homology, and $D \cap D = D$ would amount to saying that $x^2 = x$ with respect to this multiplication. So instead, we deform D so that we get a transversal intersection. Can

3.1. THE DERIVED CATEGORY OF COHERENT SHEAVES

the sheaf approach to intersection theory help us to solve this problem? Tensoring the chain complex $0 \rightarrow \mathcal{L}^* \xrightarrow{s} \mathcal{O}_X \rightarrow 0$ with \mathcal{O}_D , we get

$$0 \rightarrow \mathcal{L}^*|_D \xrightarrow{s|_D} \mathcal{O}_D \rightarrow 0$$

Of course, $s|_D = 0$ since $D = Z(s)$. As such, we are now given two pieces of information: the structure sheaf \mathcal{O}_D , as well as the line bundle $\mathcal{L}^*|_D$. The latter gives us the correct intersection, by looking at its first Chern class in $H^2(D, \mathbb{Z})$. These two pieces of information are the cohomology of the chain complex. The upshot of this observation is that the sheaf approach is actually a correct approach to intersection theory, whereas the more naive approach is not.

Longer chain complexes would appear if, instead of divisors, we took subvarieties of higher codimension. This gives us a reason why we should be interested in chain complexes of sheaves and their cohomology. The reason why we should invert quasi-isomorphisms, is that this method depends on the chosen (projective) resolution, which for our example was

$$0 \rightarrow \mathcal{L}_i^* \rightarrow \mathcal{O}_X \rightarrow \mathcal{O}_{D_i} \rightarrow 0$$

We only care about the cohomology which we obtain in this way, and so we should be working in the (bounded) derived category $D^b(\mathcal{X})$. The fact that we can work with locally free sheaves is another convenient fact of passing to the derived category, and it can be exploited in much more generality, as we describe below.

Now that we have some intuition for why we should care about the derived category $D^b(\mathcal{X})$, let us investigate some of the additional structure that it has. After all, mirror symmetry relates $D^b(\mathcal{X})$ to $D^b\text{Fuk}(X^\vee, \omega^\vee)$, through a quasi-equivalence of A_∞ -categories, but we do not yet know what the A_∞ -structure on $D^b(\mathcal{X})$ is. We will give it such a structure via D.5.6, by first constructing a different dg-category.

Definition 3.1.5. A complex of (coherent) sheaves on \mathcal{X} is called perfect if it is quasi-isomorphic to a bounded complex of locally free sheaves (i.e. vector bundles).

We will now construct a dg-category $D_\infty^b(\mathcal{X})$ whose objects are perfect complexes. Let $\mathcal{E}_\bullet, \mathcal{F}_\bullet \in D_\infty^b(\mathcal{X})$. Then we need to give $\text{Hom}_{D_\infty^b(\mathcal{X})}(\mathcal{E}_\bullet, \mathcal{F}_\bullet)$ the structure of a chain complex, which we can do as follows, as per [5]. We first define a complex at the level of \mathcal{O}_X -modules

$$\mathcal{H}om_{\mathcal{O}_X}^\bullet(\mathcal{E}_\bullet, \mathcal{F}_\bullet) \in \text{Coh}(\mathcal{X}) \quad \mathcal{H}om_{\mathcal{O}_X}^m(\mathcal{E}_\bullet, \mathcal{F}_\bullet) := \oplus_n \mathcal{H}om_{\mathcal{O}_X}(\mathcal{E}_n, \mathcal{F}_{m+n})$$

We can construct a differential D' by using the differentials on the chain complexes \mathcal{E}_\bullet and \mathcal{F}_\bullet . We define it on homogeneous elements by

$$D' : \mathcal{H}om_{\mathcal{O}_X}^n(\mathcal{E}_\bullet, \mathcal{F}_\bullet) \rightarrow \mathcal{H}om_{\mathcal{O}_X}^{n+1}(\mathcal{E}_\bullet, \mathcal{F}_\bullet) \quad (D'(f))_m := d_{\mathcal{F}_\bullet} \circ f_m - (-1)^n f_{m+1} \circ d_{\mathcal{E}_\bullet}$$

where $f_m : \mathcal{E}_m \rightarrow \mathcal{F}_{m+n}$ for f a morphism of degree n . Next, we choose an affine open cover \mathcal{U} and consider the corresponding Čech complex. We define

$$\text{Hom}_{D_\infty^b(\mathcal{X})}^k(\mathcal{E}_\bullet, \mathcal{F}_\bullet) := \bigoplus_{p+q=k} \check{C}^p(\mathcal{U}, \mathcal{H}om_{\mathcal{O}_X}^q(\mathcal{E}_\bullet, \mathcal{F}_\bullet))$$

3.1. THE DERIVED CATEGORY OF COHERENT SHEAVES

Using the standard Čech differential $D'' : \check{C}^p(\mathcal{U}, \mathcal{F}) \rightarrow \check{C}^{p+1}(\mathcal{U}, \mathcal{F})$, we can combine D' and D'' into a differential D on the total complex:

$$\begin{array}{ccccccc}
 & & \cdots & & \cdots & & \\
 & & \uparrow & & \uparrow & & \\
 \dots & \longrightarrow & \check{C}^p(\mathcal{U}, \mathcal{H}om_{\mathcal{O}_X}^{q+1}(\mathcal{E}_\bullet, \mathcal{F}_\bullet)) & \xrightarrow{D''} & \check{C}^{p+1}(\mathcal{U}, \mathcal{H}om_{\mathcal{O}_X}^{q+1}(\mathcal{E}_\bullet, \mathcal{F}_\bullet)) & \longrightarrow & \dots \\
 & & \uparrow^{D'} & & \uparrow^{D'} & & \\
 \dots & \longrightarrow & \check{C}^p(\mathcal{U}, \mathcal{H}om_{\mathcal{O}_X}^q(\mathcal{E}_\bullet, \mathcal{F}_\bullet)) & \xrightarrow{D''} & \check{C}^{p+1}(\mathcal{U}, \mathcal{H}om_{\mathcal{O}_X}^q(\mathcal{E}_\bullet, \mathcal{F}_\bullet)) & \longrightarrow & \dots \\
 & & \uparrow & & \uparrow & & \\
 & & \cdots & & \cdots & &
 \end{array}$$

The diagonals are the homogeneous degree parts of the total complex. We define a differential D on it by

$$D := D' + (-1)^p D'' : \text{Hom}_{D_\infty^b(\mathcal{X})}^k(\mathcal{E}_\bullet, \mathcal{F}_\bullet) \rightarrow \text{Hom}_{D_\infty^b(\mathcal{X})}^{k+1}(\mathcal{E}_\bullet, \mathcal{F}_\bullet)$$

This gives $D_\infty^b(\mathcal{X})$ the structure of a dg-category, i.e. an A_∞ -category for which $m_k = 0$ when $k \geq 3$. Indeed, we define $m_1(f) := (-1)^{|f|} D(f)$ and $m_2(f, g) := (-1)^{|f|} g \circ f$, where the composition is defined via the composition of morphisms of sheaves. Hence, this composition is certainly associative, which means we can indeed take the higher multiplication maps to vanish. Of course, any other affine open cover would give a quasi-isomorphic complex, and at the level of cohomology, we get

$$H^i(\text{Hom}_{D_\infty^b(\mathcal{X})}^\bullet(\mathcal{E}_\bullet, \mathcal{F}_\bullet)) = \text{Ext}^i(\mathcal{E}_\bullet, \mathcal{F}_\bullet)$$

Theorem 3.1.6. *There is an equivalence of triangulated categories $D^b(\mathcal{X}) \cong H^\bullet(D_\infty^b(\mathcal{X}))$.*

It is important that \mathcal{X} is smooth and projective, otherwise the result would not hold. To prove the result, we cite [39] which says that any coherent sheaf on a (quasi-)projective variety has a finite locally free resolution. That is, if \mathcal{F} is a coherent sheaf on \mathcal{X} , then there exists a bounded exact sequence of sheaves

$$\mathcal{E}_\bullet \rightarrow \mathcal{F} \rightarrow 0$$

and each term in \mathcal{E}_\bullet is the sheaf of sections of a vector bundle on \mathcal{X} . The cohomology of the chain complex $\mathcal{E}_\bullet \rightarrow 0$ is \mathcal{F} , and so we have an isomorphism $\mathcal{E}_\bullet \cong \mathcal{F}[0]$ in $D^b(\mathcal{X})$. Thus, each object in $D^b(\mathcal{X})$ which is concentrated in degree 0 is quasi-isomorphic to a perfect complex. Now, let

$$\begin{array}{ccc}
 \mathcal{E}_\bullet & \xrightarrow{\quad} & \mathcal{F}_\bullet \\
 & \swarrow \text{---} & \searrow \text{---} \\
 & \mathcal{G}_\bullet &
 \end{array}$$

be a distinguished triangle in $D^b(\mathcal{X})$. If two out of the three complexes are perfect, then the third is also perfect. We omit the proof, see [40]. Once we know this, it is easy to prove the following:

3.2. THE FUKAYA CATEGORY

Theorem 3.1.7. *Let \mathcal{E}_\bullet be a bounded complex of coherent sheaves on \mathcal{X} . Then \mathcal{E}_\bullet is perfect.*

The proof is by induction on the length of the complex, using distinguished triangles and the truncation of complexes. This gives us the desired equivalence $D^b(\mathcal{X}) \cong H^\bullet(D_\infty^b(\mathcal{X}))$. However, we know from [D.5.6](#) that there is a quasi-equivalence of A_∞ -categories

$$D_\infty^b(\mathcal{X}) \cong H^\bullet(D_\infty^b(\mathcal{X}))$$

which we can use to transfer the dg-structure to $D^b(\mathcal{X}) \cong H^\bullet(D_\infty^b(\mathcal{X}))$, such that the differential on $D^b(\mathcal{X})$ is identically 0. On the other hand, the m_k for $k \geq 3$ no longer have to be 0, and so the resulting A_∞ -structure on $D^b(\mathcal{X})$ is not at all trivial. In summary, we have given an interpretation of the kind of structure that the derived category encodes, and we have shown that it can be given an A_∞ -structure.

3.2 The Fukaya Category

The construction of the Fukaya category makes the need for A_∞ -structures manifest in a geometric way. Denote the Fukaya category of a Calabi-Yau manifold \mathcal{X} by $\text{Fuk}(X, \omega)$, where ω is the Kähler form. The construction can be extended to more general symplectic manifolds, but we restrict to the Calabi-Yau case.

The objects in the Fukaya category are easy to define, and they are essentially Lagrangian submanifolds - just with some extra data. We will see why the extra data is needed only when constructing the Hom-spaces.

Definition 3.2.1. The objects of $\text{Fuk}(X, \omega)$ consist of tuples (L, f, E, ∇, S_L) where

1. L is a Lagrangian submanifold.
2. $f : L \rightarrow \mathbb{R}$ is such that $\text{Re } f\Omega|_L = \text{Vol}_g$, which is called a grading.¹
3. S_L is a spin structure on L .
4. (E, ∇) is a flat unitary (i.e. structure group $U(1)$) line bundle over L .

The first two items are to ensure that the morphism spaces are well-defined, as we will see in a moment. The Fukaya category describes the intersection theory of Lagrangian submanifolds in the symplectic manifold (X, ω) . It does not make reference to the complex structure \mathcal{X} in any way. We will typically denote objects of the Fukaya category simply by L , and the other data will be implicit.² Next, we want to define the Hom-spaces. To do this, we need the Floer chain complex, which is somewhat technical to define. We refer to the [appendix C.3](#) for its construction. The main point is this: whenever we have two intersecting Lagrangians, we can perturb them via Hamiltonian

¹There is a topological obstruction to the existence of such a function, called the Maslov class. See [C.2.3](#) for details.

²One should also demand that the Lagrangian submanifolds in the Fukaya category are unobstructed, in an appropriate sense. We are not going to discuss this technicality.

3.2. THE FUKAYA CATEGORY

isotopies (see 1.1.4) to intersect transversally. We define a vector space which is generated by the intersection points. There is a natural way to define a grading on this vector space, and to define a differential, one "counts" pseudo-holomorphic disks which connect two points, and whose boundary components lie in the respective Lagrangian submanifolds. The resulting chain complex is denoted $(CF^\bullet(L_0, L_1), \partial)$ for Lagrangian submanifolds $L_0, L_1 \subseteq X$. In order for the grading to be a \mathbb{Z} -grading, we need a grading on the submanifolds. In order to get a well-defined differential when we take $CF^\bullet(L_0, L_1)$ to be a vector space over a field with characteristic other than 2, we need to equip the Lagrangians with spin structures. This is why the objects in the Fukaya category are defined as they are.

Having said that, let us take $L_0, L_1 \in \text{Fuk}(X, \omega)$ with $E_i = L_i \times \mathbb{C}$ and $\nabla_i = d$. We define

$$\text{Hom}_{\text{Fuk}(X, \omega)}(L_0, L_1) := CF^\bullet(L_0, L_1)$$


This vector spaces is already graded, and has a differential from Floer theory. To make $\text{Fuk}(X, \omega)$ into an A_∞ -category, we need higher multiplication maps which satisfy the A_∞ -relations. In the present case, the multiplication maps should be given as

$$m_k : CF^\bullet(L_{k-1}, L_k) \otimes \cdots \otimes CF^\bullet(L_0, L_1) \rightarrow CF^\bullet(L_0, L_k)[2 - k]$$

where $[2 - k]$ denotes a degree shift. We first state how this map is defined, and then explain what it means geometrically. We set

$$m_k(p_k, \dots, p_1) := \sum_{\substack{q \in L_0 \cap L_k, \\ [u] \in \pi_2(X, L_0 \cup \cdots \cup L_k): \\ \text{ind}([u]) = 2 - k}} \# \overline{\mathcal{M}}(p_1, \dots, p_k, q, [u], J) T^{\langle \omega, [u] \rangle} \cdot q$$

For $k = 1$, we recover the definition of the Floer differential. The moduli space in this equation is completely analogous to the one we defined in C.3, so we will not describe it explicitly. Instead, we describe it geometrically. Namely, for the Floer differential, we considered pseudo-holomorphic disks (with two points removed) which connected the points p and q , such that the boundary was mapped to the respective Lagrangian submanifolds. The idea here is the same. Instead of 2 punctures in the disk, we get $1 + k$ punctures. These punctures should map to the points $p_1, \dots, p_k, q \in \mathcal{X}$. The boundary of this $(k + 1)$ -punctured disk will consist of equally many line segments, and these line segments are required to map to the respective Lagrangian submanifolds L_i . For example, see the diagram below for an example of a map u which will contribute to m_2 :



3.2. THE FUKAYA CATEGORY

In this way, we get a space of maps, and we again quotient out reparameterisations. We obtain a space of dimension $\text{ind}([u]) + k - 2$. Hence, we sum over all homotopy classes of maps which yield a 0-dimensional moduli space. This space can then be compactified via Gromov compactification, and the signed count of its points (coming from the orientation induced by the spin structures) determines the higher multiplication maps.

Theorem 3.2.2. [28] *Let m_k be as above. Then the m_k satisfy the A_∞ -relations for all $k \geq 1$.*

In particular, m_2 defines an associative multiplication at the level of cohomology, but not at the level of chain complexes. When $L_0 = L_1$, this product structure coincides with the usual cup product on $HF^\bullet(L, L) = H^\bullet(L, \Lambda)$, under the right circumstances (see [28]).


Now that we have defined the Hom-spaces when $E \cong L \times \mathbb{C}$ with the trivial connection, the general case is only a mild generalisation, and it is going to involve the parallel transport of the connection ∇ on E . We note that

$$CF^\bullet(L_0, L_1) = \mathbb{C}\langle L_0 \cap L_1 \rangle \otimes_{\mathbb{C}} \Lambda$$

In other words, for each intersection point, we get a 1-dimensional complex vector space, before we pass to the Novikov field. This 1-dimensional complex vector space $\mathbb{C} \cdot p$ is interpreted as $\text{Hom}_{\mathbb{C}}(E_{0,p}, E_{1,p})$, where the E_i are the flat unitary line bundles which are part of the definition of the objects of the Fukaya category. This definition can be immediately adapted to non-trivial E_i . For the multiplication maps, we use the parallel transport of the connection. Namely, let $u \in \mathcal{M}(p_1, \dots, p_k, q, [u], J)$. Then each of the boundary strips will be mapped to some Lagrangian submanifold. Let us denote the image of the boundary strip on L_i by γ_i , which connects p_i with p_{i+1} (and we identify $p_{k+1} = q$). Then we restrict E_i to the strip γ_i , and the restricted connection gives us the parallel transport operator from A.3.2.

$$\text{Par}_{\gamma_i} : E_{i,p_i} \xrightarrow{\cong} E_{i,p_{i+1}}$$

Geometrically, this map can be pictured as in the following image:



3.2. THE FUKAYA CATEGORY

This depicts a non-trivial (real) line bundle over a Lagrangian submanifold of \mathbb{C}/Λ . The blue vectors depict the parallel transport of a given vector $v \in E_{x_0}$ along a curve $\gamma : I \rightarrow \mathbb{C}/\Lambda$ over which E may be trivialised. The section $s(t)$ in red is a parallel section of this trivialisation of γ^*E over I . Evidently, the resulting operator is $\text{Par}_\gamma = -1 \in \text{GL}(1, \mathbb{R})$.

Now, the idea is to define $m_k(a_k, \dots, a_1)$ by viewing $a_i \in \text{Hom}_{\mathbb{C}}(E_{i,p_i}, E_{i+1,p_i})$ as endomorphisms of the fibres of the line bundles. Then we can use the parallel transport operators to define

$$m_k(a_k, \dots, a_1) := \sum_{q \in L_0 \cap L_1} \left(\sum_{u \in \mathcal{M}(p_1, \dots, p_k, q, [u], J)} (-1)^{\sigma(u)} \text{Par}_{\gamma_k} \circ a_k \circ \dots \circ a_1 \circ \text{Par}_{\gamma_0} \otimes T^{\langle \omega, [u] \rangle} \right) \cdot q \quad (3.1)$$

In this expression, $(-1)^{\sigma(u)}$ is determined by the orientation of the point u in the moduli space. These multiplication maps again satisfy the A_∞ -relations.

Definition 3.2.3. Let (\mathcal{X}, ω) be a Calabi-Yau manifold. The Fukaya category $\text{Fuk}(X, \omega)$ is the A_∞ -category defined by the following data:

1. The objects defined in 3.2.1.
2. The Hom-spaces defined as

$$\text{Hom}_{\text{Fuk}(X, \omega)}(L_0, L_1) := \bigoplus_{p \in L_0 \cap L_1} \text{Hom}_{\mathbb{C}}(E_{0,p}, E_{1,p}) \otimes_{\mathbb{C}} \Lambda$$

3. The multiplication maps m_k defined in 3.1

There is a shift functor on the Fukaya category, which shifts the grading by 1. Hence, we can consider the derived Fukaya category $D^b\text{Fuk}(X, \omega)$ (see the appendix).

Example 23. We will discuss the (twisted) Fukaya category of a 2-torus in 4.4.1. Essentially, each object is isomorphic to a direct summand of a twisted complex built from the standard generators of $\pi_1(T^2)$. See [28] for details.

Chapter 4


Homological Mirror Symmetry

4.1 Mirror Symmetry and String Theory

We now aim to give an overview of the origins of mirror symmetry in string theory. This will be a high level discussion, in the sense that we omit many technicalities. We start in the 1990s when B. Greene and M. Plesser published their paper [41] titled "Duality in Calabi-Yau Moduli Space", but many more references will be given throughout the text. Their abstract reads:

"We describe a duality in the moduli space of string vacua which pairs topologically distinct Calabi-Yau manifolds and shows that they yield isomorphic conformal theories. At the level of the geometrical description, this duality interchanges the roles of Kähler and complex structure moduli and thus pairs manifolds whose Euler numbers differ by $\chi(X) \mapsto -\chi(X)$."

Let us try to give an explanation of what this means. In relativity, one considers the worldline of a particle, which is its trajectory in spacetime. When we consider strings instead of point particles, we speak of worldsheets rather than worldlines, because we are now looking at the surface which is traced out by a 1-dimensional object, i.e. a string. We can consider either open strings, diffeomorphic to the interval I , or closed strings, diffeomorphic to the circle S^1 . Let us start by considering only closed strings. In this case, worldsheets will be Riemann surfaces Σ with boundary. Interactions between strings are naturally encoded in the theory. We denote $\partial\Sigma = \partial\Sigma^- \cup \partial\Sigma^+$, the incoming and outgoing boundary components. Each component is thought of as a string. The Riemann surfaces which provide a cobordism between the components are thought of as interactions between the incoming strings and the outgoing strings, such as in the following diagram:



Such diagrams are the string theoretic analogues of Feynman diagrams in quantum field theory. For instance, the diagram above represents an interaction in which a single


4.1. MIRROR SYMMETRY AND STRING THEORY

closed string splits into two closed strings.

To get a physical theory, we need to prescribe some action functional. The action functional will be a functional on the space of maps $\text{Map}(\Sigma, X)$ for a target manifold X , which is our spacetime. This is called a non-linear sigma model (if X were a vector space, it would be a linear sigma model). The natural way to do this, is by imposing a metric on Σ and X and defining

$$S(\phi) = \int_{\Sigma} d\phi \wedge \star d\phi$$

One readily verifies that this action is invariant under conformal transformations (i.e. rescaling of the metric and diffeomorphisms), which means that the action only depends on the conformal structure on Σ . Since there is a bijection between conformal structures and complex structures on Σ , we can indeed regard Σ as a Riemann surface. Furthermore, conformal invariance allows us to shrink the boundary components to be arbitrarily small, and as such, it is typical to view the worldsheets in a closed string theory as closed Riemann surfaces, by inserting vertex operators - these correspond to the boundary string states via the operator-state correspondence. For example:



In the above picture, the marked points on \mathbb{CP}^1 denote the punctures which were "filled up" by using vertex operators. This is our crude justification for considering closed Riemann surfaces as the worldsheets for closed strings. We refer to [42, 43, 44] for a more rigorous discussion of this fact.

Next, we want to include supersymmetry, so that we will be considering superstrings.¹ From the physical perspective, there are several reasons to do this. Most importantly, it gives a way to include fermions in the theory. Supersymmetry says essentially that the Lie algebra of infinitesimal symmetries of the theory is a super Lie algebra. That is, it should be a $\mathbb{Z}/2$ -graded vector space $\mathfrak{g} = \mathfrak{g}_0 \oplus \mathfrak{g}_1$ with a Lie bracket that is also appropriately graded. When restricted to \mathfrak{g}_0 , it should yield an ordinary Lie algebra, presently the Poincaré algebra of Minkowski spacetime symmetries. However, there is a marked distinction between these local symmetries, and global symmetries. Indeed, consider for instance general relativity, in which the action functional depends on the metric. In this case, the Killing vector fields generate the global symmetries of the theory. These are the vector fields whose flow preserve the metric. They may not even exist, even though we have local symmetries given by the Poincaré algebra, which is the Lie algebra of the gauge group of general relativity. Likewise, local supersymmetry (for the present pur-

¹See [45, 43] for a mathematical perspective on supersymmetry.

4.1. MIRROR SYMMETRY AND STRING THEORY

poses, a super Lie algebra whose \mathfrak{g}_0 component is the Poincaré algebra) does not imply the existence of a global supersymmetry. In the supermanifold formalism, such a global supersymmetry would be an odd vector field (i.e. a spinor) which preserves the action. In [46] it is shown that the existence of sufficient local supersymmetry (called $N = (2, 2)$ supersymmetry) in a non-linear sigma model requires the target space to have a Kähler structure (\mathcal{X}, g) . If we additionally require the existence of a global supersymmetry (for the "effective" theory on \mathbb{R}^4), then this is precisely the requirement that there exists a parallel spinor on (\mathcal{X}, g) .² In turn (see [47]), this implies that (\mathcal{X}, g) must be a Calabi-Yau manifold. This was the primary case of interest for physicists, cf. [43]. Furthermore, Ricci flatness of the metric is required for the quantisation of the theory to remain superconformally invariant. To make the physics work out, it is required that $\dim_{\mathbb{C}} \mathcal{X} = 3$, which is how we arrive at Calabi-Yau threefolds, as well as the famous $4 + 6 = 10$ dimensions in which string theory takes place. But $\dim \mathcal{X} = 3$ is not essential for the definition of the sigma model, and the discussion below applies to Calabi-Yau manifolds more generally.

So the setup is to consider Riemann surfaces Σ , and base the theory around maps $\phi : \Sigma \rightarrow \mathbb{R}^4 \times \mathcal{X}$ where (\mathcal{X}, g) is a Calabi-Yau manifold. We will focus on maps $\phi : \Sigma \rightarrow \mathcal{X}$, which comprises one part of the theory. In fact, since the worldsheets Σ can naturally be coupled to a closed 2-form $B \in \Omega^2(X)$ via pullback, we will consider it as part of the data of a Calabi-Yau manifold. It should be interpreted as a higher gauge field, i.e. as the connection form on a trivial unitary gerbe, a higher analogue of the electric potential.³ Only its cohomology class enters into the action functional, through integration, so we view B as the representative of some class $B \in H^2(X, \mathbb{R})$. Physicists believe that every such triple (\mathcal{X}, g, B) gives rise to an $N = 2$ superconformal vertex algebra (SCVA). We are not going to define what these are explicitly, referring instead to [48]. The main point is that these objects are algebraic gadgets which encode the string theory, and they are representations of an infinite dimensional super Lie algebra called the super Virasoro algebra (all of this is discussed in the cited source). Not all $N = 2$ SCVA arise from non-linear sigma models with Calabi-Yau target space. We call the ones that do geometric.

At the algebraic level, it is quite easy to write down an involution of an $N = 2$ SCVA which yields another $N = 2$ SCVA. If two $N = 2$ SCVAs are related in this way, they are said to be mirror to each other. It should be noted that the generators which are exchanged by the mirror morphism have a distinct geometric interpretation, when the $N = 2$ SCVA comes from a Calabi-Yau threefold. In fact, the Gepner conjecture claims that every $N = 2$ SCVA with some additional property is geometric. This property is preserved by the mirror morphism, and so the mirror morphism acts on the moduli space of Calabi-

²Spinors model fermions in the standard model. In the appendix, we give a brief review of spinor bundles and spinors.

³A unitary gerbe is some higher analogue of a unitary line bundle. Flat unitary lines bundles are classified by $H^1(X, \mathbb{R}/\mathbb{Z})$, and flat unitary gerbes are classified by $H^2(X, \mathbb{R}/\mathbb{Z})$. We consider a trivial unitary gerbe, i.e. one which lies in the kernel of $H^2(X, \mathbb{R}/\mathbb{Z}) \rightarrow H^3(X, \mathbb{Z})$. This allows a lift to a class $B \in H^2(X, \mathbb{R})$ and subsequently to a closed 2-form.

4.1. MIRROR SYMMETRY AND STRING THEORY

Yau threefolds.

Definition 4.1.1. Let (\mathcal{X}, g, B) be a Calabi-Yau threefold, $\text{SCVA}(\mathcal{X}, g, B)$ the associated $N = 2$ SCVA, and let μ be the mirror morphism. Then the mirror⁴ of (\mathcal{X}, g, B) is the Calabi-Yau threefold $(\mathcal{X}^\vee, g^\vee, B^\vee)$ such that

$$\mu(\text{SCVA}(\mathcal{X}, g, B)) \cong \text{SCVA}(\mathcal{X}^\vee, g^\vee, B^\vee)$$

Unfortunately, string theory in its full generality is too difficult for us to study mathematically. In his paper "Topological Quantum Field Theory" [49], Witten discusses how one can perform a topological twisting on the original string theory, which isolates some "topological sector" of the theory. This can then be used to obtain a topological quantum field theory, which has no dynamics, i.e. the Hamiltonian is identically zero. However, because of the quantum effects, this still produces a mathematically interesting theory, and more importantly, a mathematically well-defined theory. It turns out that there are two inequivalent ways to twist a supersymmetric non-linear sigma model, and produce a topological string theory:

1. The A -model, which depends only on the complexified Kähler class $(X, B + i\omega)$
2. The B -model, which depends only on the complex structure \mathcal{X}

Witten argues in [50] that the mirror morphism exchanges the A -model on (\mathcal{X}, g, B) with the B -model on $(\mathcal{X}^\vee, g^\vee, B^\vee)$.⁵ Furthermore, the moduli spaces of the physical theories should be isomorphic in some open neighbourhood of the corresponding points in the moduli space, because deformations of the respective physical theories should coincide. This is what is meant with the quote from [41]. This statement has very non-trivial mathematical implications, particularly on Calabi-Yau threefolds, which is the case we will now specialise to. For a start, there is a relation between the Hodge diamonds of \mathcal{X} and \mathcal{X}^\vee . Recall that $h^{p,q}(\mathcal{X}) := \dim_{\mathbb{C}} H^{p,q}(\mathcal{X})$. The Hodge diamond of a complex manifold \mathcal{X} is the following diagram:

$$\begin{array}{ccccc}
 & & h^{0,0}(\mathcal{X}) & & \\
 & & h^{1,0}(\mathcal{X}) & h^{0,1}(\mathcal{X}) & \\
 & & h^{2,0}(\mathcal{X}) & h^{1,1}(\mathcal{X}) & h^{0,2}(\mathcal{X}) \\
 & h^{3,0}(\mathcal{X}) & h^{2,1}(\mathcal{X}) & h^{1,2}(\mathcal{X}) & h^{0,3}(\mathcal{X}) \\
 & & h^{3,1}(\mathcal{X}) & h^{2,2}(\mathcal{X}) & h^{1,3}(\mathcal{X}) \\
 & & h^{3,2}(\mathcal{X}) & h^{3,3}(\mathcal{X}) & \\
 & & h^{3,3}(\mathcal{X}) & &
 \end{array}$$

⁴Although we refer to "the" mirror, there is not necessarily a unique mirror manifold of a given Calabi-Yau manifold.

⁵As noted in the introduction, physicists actually conjecture that this is true more generally, not just at the level of topological string theory, but at the level of supersymmetric conformal field theories, known as type IIA and type IIB string theory. But this is not a mathematically rigorous claim.

4.1. MIRROR SYMMETRY AND STRING THEORY

On a Kähler manifold, the Hodge numbers exhibit a number of symmetries, which we mentioned in 2.4.5. In the setting of strict Calabi-Yau 3-folds, we have imposed that $h^{1,0}(\mathcal{X}) = h^{2,0}(\mathcal{X}) = 0$. Using this, together with the symmetries of the Hodge diamond for a Kähler manifold, this means that the Hodge diamond of a strict Calabi-Yau 3-fold takes the following shape:

$$\begin{array}{ccccc}
 & & & & 1 \\
 & & & & 0 & & 0 \\
 & & & 0 & h^{1,1}(\mathcal{X}) & 0 \\
 1 & & h^{2,1}(\mathcal{X}) & h^{2,1}(\mathcal{X}) & & 1 \\
 & & 0 & h^{1,1}(\mathcal{X}) & 0 \\
 & & & & 0 & & 0 \\
 & & & & & & 1
 \end{array}$$

Mirror symmetry earned its name because of the fact that, for the mirror manifold \mathcal{X}^\vee , we have $h^{1,1}(\mathcal{X}^\vee) = h^{2,1}(\mathcal{X})$ and $h^{2,1}(\mathcal{X}^\vee) = h^{1,1}(\mathcal{X})$. In other words, the Hodge diamond of \mathcal{X}^\vee is obtained by a mirror reflection along the diagonal of the Hodge diamond of \mathcal{X} . The reason for this is as follows. The complexified Kähler moduli space is $\mathcal{M}_K(X) := (H^2(X, \mathbb{R}) + i\mathfrak{K}(X))/H^2(X, \mathbb{Z})$, where $\mathfrak{K}(X)$ denotes the set of all 2-forms which may arise as Kähler classes on X . This is an open subset of $H^2(X, \mathbb{R})$, and so $\dim_{\mathbb{C}} \mathcal{M}_K(X) = h^{1,1}(\mathcal{X})$. The complex moduli space $\mathcal{M}_c(X)$ of X is defined as the space of complex structures on X modulo diffeomorphisms. The Bogomolov-Tian-Todorov theorem states that, when X admits a Calabi-Yau structure, $\mathcal{M}_c(X)$ is smooth, with tangent space given by $H^1(\mathcal{X}, T\mathcal{X})$, where $(X, J) = \mathcal{X}$. Observe that on a Calabi-Yau threefold, we have a non-degenerate pairing

$$\wedge : T\mathcal{X} \otimes \wedge^2 T\mathcal{X} \rightarrow \wedge^3 T\mathcal{X} \cong \mathcal{O}_{\mathcal{X}}$$

We obtain $T\mathcal{X} \cong \wedge^2 T^*\mathcal{X}$, implying that $H^1(\mathcal{X}, T\mathcal{X}) \cong H^1(\mathcal{X}, \Omega_{\mathcal{X}}^2) = H^{2,1}(\mathcal{X})$. It follows that $\dim_{\mathbb{C}} \mathcal{M}_c(X) = h^{2,1}(\mathcal{X})$. Thus, the assertion that some open neighbourhoods of the moduli spaces of the physical theories are isomorphic is translated into a mathematical statement about the topology of the mirror manifold, and about the existence of Calabi-Yau manifolds with certain topological properties, since $\chi(X) = -\chi(X^\vee)$.

To be slightly more precise, for certain distinguished complex structures J (large complex structure limit points, which we will not define), there should be an open neighbourhood $U \subseteq \mathcal{M}_c(X)$ of J and a holomorphic map $\mu : U \rightarrow \mathcal{M}_K(X^\vee)$ which is an isomorphism onto its image. The map μ is called the mirror map. In physics terms, the B -model of a point $p \in \mathcal{M}_c(X)$ should correspond to the A -model of the point $\mu(p) \in \mathcal{M}_K(X^\vee)$. Mathematically, this is expressed in terms of the Yukawa couplings. They determine the physical theory, so under the mirror map, the Yukawa couplings should coincide. For the A -model, the Yukawa coupling is given as follows. Let $p := B + i\omega \in \mathcal{M}_K(X)$. Then

4.1. MIRROR SYMMETRY AND STRING THEORY

we have a canonical identification $T_p\mathcal{M}_K(X) \cong H^2(X, \mathbb{C})$. The Yukawa coupling is a map

$$\mathfrak{Y}_p^{1,1} : H^2(X, \mathbb{C})^{\otimes 3} \rightarrow \mathbb{C}[[q; S]]$$

where $\mathbb{C}[[q; S]]$ is the ring of formal power series over \mathbb{C} whose exponents lie in some semigroup S . Presently, the semigroup is given by

$$S = \{Y \in H_2(X, \mathbb{Z}) \mid \int_Y \omega \geq 0 \quad \forall \omega \in \mathfrak{K}(X)\}$$

We think of $\mathfrak{Y}_p^{1,1}$ as a "quantum deformation" of the usual intersection pairing on cohomology, defined by

$$\mathfrak{Y}_p^{1,1}(\eta_1, \eta_2, \eta_3) := \eta_1 \wedge \eta_2 \wedge \eta_3 + \sum_{0 \neq Y \in H_2(X, \mathbb{Z})} \text{GW}_Y(\eta_1, \eta_2, \eta_3) \frac{q^Y}{1 - q^Y} \in \mathbb{C}[[q; S]]$$

When $H_2(X, \mathbb{Z})$ is torsion free, this may be interpreted as some formal power series, and for practical purposes it should simply be interpreted as a function on the tangent space of the moduli space. The expression $\text{GW}_Y(\eta_1, \eta_2, \eta_3)$ is the genus zero Gromov-Witten invariant of the homology class Y , on the Calabi-Yau manifold \mathcal{X} . It counts the number of holomorphic curves $\mathbb{C}\mathbb{P}^1 \rightarrow \mathcal{X}$ which represent the homology class Y , modulo reparameterisations, such that the points $0, 1, \infty \in \mathbb{C}\mathbb{P}^1$ are mapped to cycles which are Poincaré dual to the η_i . The Gromov-Witten invariants give us information about the enumerative geometry of \mathcal{X} , and mirror symmetry says that this information can be computed from the Yukawa coupling of the B -model on the mirror \mathcal{X}^\vee . As noted previously, $T_p\mathcal{M}_c(X) \cong H^1(\mathcal{X}, T\mathcal{X})$. The Yukawa coupling $\mathfrak{Y}_p^{2,1}$ for the B -model at the point q is a map

$$\mathfrak{Y}_p^{2,1} : H^1(\mathcal{X}, T\mathcal{X})^{\otimes 3} \rightarrow \mathbb{C}$$

Observe that Serre duality⁶ tells us that $H^3(\mathcal{X}, \wedge^3 T\mathcal{X}) \cong H^0(\mathcal{X}, K_{\mathcal{X}} \otimes \Omega_{\mathcal{X}}^3) = H^0(\mathcal{X}, K_{\mathcal{X}} \otimes K_{\mathcal{X}})$. Assuming a suitable normalisation on the holomorphic volume form Ω , a Calabi-Yau manifold gives us a canonical element $\Omega \otimes \Omega \in H^0(\mathcal{X}, K_{\mathcal{X}} \otimes K_{\mathcal{X}})$. Denote the Serre duality pairing by $\langle \cdot, \cdot \rangle$. Then we set

$$\mathfrak{Y}_p^{2,1}(\eta_1, \eta_2, \eta_3) := \langle \eta_1 \wedge \eta_2 \wedge \eta_3, \Omega \otimes \Omega \rangle$$

The physics of mirror symmetry suggests that under the mirror map $\mu : U \rightarrow \mathcal{M}_K(X^\vee)$, we get

$$\mu^* \mathfrak{Y}^{1,1} = \mathfrak{Y}^{2,1}$$

Conjecture 4.1.2 (The Mirror Conjecture). *Let (\mathcal{X}, g) be a Calabi-Yau threefold. Then there exists a Calabi-Yau threefold $(\mathcal{X}^\vee, g^\vee)$, and a mirror map $\mu : \mathcal{M}_c(X) \supseteq U \xrightarrow{\sim} V \subseteq \mathcal{M}_K(X^\vee)$ which makes the respective Yukawa couplings coincide.⁷*

⁶Serre duality states, for any coherent sheaf \mathcal{E} , that $H^i(\mathcal{X}, \mathcal{E}) \cong H^{n-i}(\mathcal{X}, K_{\mathcal{X}} \otimes \mathcal{E}^*)^*$.

⁷If the complex moduli space of a Calabi-Yau manifold is 0-dimensional, it cannot have a mirror in this sense, because the Kähler cone is always of dimension ≥ 1 . This is a hint that mirror symmetry holds more generally, which indeed it does, but we will not investigate that.

4.2. EXAMPLE: MIRROR SYMMETRY FOR $K3$ SURFACES

This is remarkable for several reasons. Firstly, we noted that there were "quantum corrections" in the A -model Yukawa coupling, which is why we get a formal power series. However, for the B -model Yukawa coupling, we simply get a complex number, no quantum corrections. Mathematically, this means that certain notoriously difficult computations (e.g. the Gromov-Witten invariants) can be obtained by performing much simpler calculations in the B -model of the mirror manifold \mathcal{X}^\vee . This calculation was performed for the mirror quintic by Candelas et al., in their famous paper [3]. This was done by considering a family of Calabi-Yau manifolds $\pi : \mathcal{S} \rightarrow \mathcal{B}$ for a 1-dimensional base space (considered as a subspace of the moduli space of the mirror quintic), so that the Yukawa coupling $\mathfrak{Y}^{2,1}$ becomes a function on \mathcal{B} , which is calculated by evaluating certain integrals, called period integrals. This function has a Taylor expansion, and its coefficients coincide with those of the Yukawa coupling $\mathfrak{Y}^{1,1}$ on the Kähler moduli space of the quintic. As such, the Gromov-Witten invariants for \mathcal{Q} can be read off.

There is a Gromov-Witten invariant for each degree of the map $\mathbb{C}\mathbb{P}^1 \rightarrow \mathcal{Q}$, corresponding to $H_2(\mathcal{Q}, \mathbb{Z}) \cong \mathbb{Z}$, and the degree 1, 2 and 3 Gromov-Witten invariants had been calculated by mathematicians at this point in time. The number of degree 1 and 2 curves obtained by Candelas et al. agreed with the number that was known. However, there was a discrepancy for degree 3. Was the mirror symmetry conjecture misguided? As it turned out, the computer code used by mathematicians to calculate the number of degree 3 curves on \mathcal{Q} contained an error. Once this had been corrected, the two numbers agreed. However, the physicists produced at least 10 coefficients, far more than any mathematician had been able to produce. Thus, mirror symmetry as a mathematical discipline was born, as it became clear that there are deep connections to be discovered in this area. While this is a beautiful story, the derived category of coherent sheaves has not appeared anywhere - nor has the derived Fukaya category. The story has a continuation which leads to homological mirror symmetry. We will continue telling this story after discussing mirror symmetry for $K3$ surfaces.

4.2 Example: Mirror Symmetry for $K3$ surfaces

Taking a step back from Calabi-Yau threefolds, we can consider one dimension lower to see how mirror manifests itself. In complex dimension 2, we have one strict Calabi-Yau manifold up to diffeomorphism, which is the $K3$ surface. Because of this, mirror symmetry for $K3$ surfaces is not going to give any relation between topologically distinct manifolds, in contrast to the case of Calabi-Yau threefolds.

Definition 4.2.1. A $K3$ surface \mathcal{X} a simply connected Calabi-Yau manifold of dimension $\dim \mathcal{X} = 2$.

It should come as no surprise that the quartic surface

$$\mathcal{X} = \{[z_0 : z_1 : z_2 : z_3] \in \mathbb{C}\mathbb{P}^3 \mid \sum z_i^4 = 0\}$$

4.2. EXAMPLE: MIRROR SYMMETRY FOR $K3$ SURFACES

is a $K3$ surface, just as the Fermat quintic was a Calabi-Yau threefold. Other examples of $K3$ surfaces may be obtained by resolving certain orbifolds, e.g. resolving the quotient of an abelian twofold by the natural \mathbb{Z}_2 -action $a \mapsto -a$. One can also adapt 2.4.1 to complex dimension 2, to obtain $K3$ surfaces as complete intersections in $\mathbb{C}\mathbb{P}^N$. From the description as a quartic surface, it is easy to deduce the Hodge numbers as we did for the quintic, and the only Hodge number that is interesting in this case is $h^{1,1}(\mathcal{X}) = 20$.

There is a well-known isomorphism $\mathrm{Sp}(1) \cong \mathrm{SU}(2)$. Riemannian manifolds of dimension $4n$ with $\mathrm{Hol}(g) \subseteq \mathrm{Sp}(n)$ are called hyper-Kähler manifolds, and they have interesting properties. In particular, they are Calabi-Yau since $\mathrm{Sp}(n) \subseteq \mathrm{SU}(2n)$, but we will focus on the case $n = 1$ in which the two notions coincide. Let (X, g) be a hyper-Kähler manifold and take $x \in X$, so we may identify $T_x X \cong \mathbb{R}^4 \cong \mathbb{H}$, where \mathbb{H} denotes the quaternion algebra. Then we have three complex structures on $T_x X$, denote them by I_x, J_x, K_x , corresponding to multiplication by the imaginary basis vectors. Now, $\mathrm{Hol}(g) \subseteq \mathrm{Sp}(1)$ leaves each of these vectors invariant, and so the holonomy principle A.3.5 tells us that we get three complex structures I, J, K on X . The manifold X is a Kähler manifold with respect to each of these complex structures, and they satisfy $I^2 = J^2 = K^2 = IJK = -\mathrm{id}$. We denote the respective Kähler forms by $\omega_I, \omega_J, \omega_K$, and we denote the $K3$ surface equipped with the respective complex structures by $\mathcal{X}_I, \mathcal{X}_J, \mathcal{X}_K$.

Proposition 4.2.2. *A holomorphic volume form for \mathcal{X}_I is given by $\Omega_I = \omega_J + i\omega_K$.*

Proof. First, we show that Ω_I is of type $(2, 0)$ w.r.t. I . Let $v, w \in T_x X$. Then

$$\begin{aligned} \Omega_I(Iv, w) &= \omega_J(Iv, w) + i\omega_K(Iv, w) = g(Iv, Jw) + ig(Iv, Kw) = \\ g(JIv, J^2 w) + ig(KIv, K^2 w) &= g(Kv, w) - ig(Jv, w) = -g(v, Kw) + ig(v, Jw) = \\ -\omega_K(v, w) + i\omega_J(v, w) &= i(\omega_J(v, w) + i\omega_K(v, w)) \end{aligned}$$

A similar calculation for the second argument then indeed shows $\Omega_I \in \Omega^{2,0}(X)$. Now, both ω_J and ω_K are symplectic forms on X , irrespective of the complex structure, so they are closed. But closed forms of type $(2, 0)$ are precisely the holomorphic ones, since $d = \partial + \bar{\partial}$ and $\partial\Omega_I = 0$ since $\dim \mathcal{X} = 2$. So Ω is a holomorphic 2-form, and it is nowhere vanishing since ω_J and ω_K are symplectic. \square

It is also easy to show that every hyper-Kähler manifold admits a $\mathbb{C}\mathbb{P}^1$ worth of compatible complex structures, namely by taking $a, b, c \in \mathbb{R}$ and considering $aI + bJ + cK$ with $a^2 + b^2 + c^2 = 1$. This means that the complex deformations and the Kähler deformations are not "decoupled", in the way that we described for Calabi-Yau threefolds. In some ways, this makes the analysis more intricate. However, $K3$ surfaces allow for a very explicit description of their complex moduli, which makes mirror symmetry for $K3$ surfaces rather concrete.

What makes $K3$ surfaces particularly well-behaved is that their complex moduli can be encoded through a lattice. We consider the usual bilinear form on the middle coho-

4.2. EXAMPLE: MIRROR SYMMETRY FOR $K3$ SURFACES

mology, given by the cup product:

$$\cup : H^2(X, \mathbb{Z}) \otimes H^2(X, \mathbb{Z}) \rightarrow H^4(X, \mathbb{Z}) \cong \mathbb{Z}$$

We get a lattice in the usual sense, i.e. a free abelian group together with a bilinear form. From now on, this bilinear form will just be denoted by juxtaposition of elements, or by $\langle \cdot, \cdot \rangle$ if the context requires it.

Theorem 4.2.3. *If \mathcal{X} is a $K3$ surface, then $H^2(X, \mathbb{Z})$ is an even⁸ unimodular lattice of signature $(3, 19)$.*

By the classification of unimodular lattices, there is a unique lattice up to isomorphism which satisfies the above properties, and we will denote it by $\Gamma_{3,19}$. An isomorphism of lattices $H^2(X, \mathbb{Z}) \cong \Gamma_{3,19}$ will usually be implicit.

The lattice structure together with the Hodge decomposition of $H^2(X, \mathbb{C})$, coming from the Kähler structure on \mathcal{X} , allows one to describe the complex modulus of a $K3$ surface in terms of lattice theory.

Definition 4.2.4. A Hodge structure of weight k is a lattice Γ together with a decomposition $\Gamma \otimes \mathbb{C} = \bigoplus_{p+q=k} H^{p,q}$ such that $\overline{H^{p,q}} = H^{q,p}$.

For a $K3$ surface, we take $k = 2$ and $\Gamma = H^2(X, \mathbb{Z})$, with the decomposition given by the Hodge decomposition $H^2(X, \mathbb{C}) \cong H^{2,0}(\mathcal{X}) \oplus H^{1,1}(\mathcal{X}) \oplus H^{0,2}(\mathcal{X})$. The lattice sits inside of the vector space $H^2(X, \mathbb{C})$ via the universal coefficients theorem $H^2(X, \mathbb{Z}) \otimes \mathbb{C} \cong H^2(X, \mathbb{C})$, but the Hodge decomposition varies as the complex structure on X changes.

Definition 4.2.5. An isometry of Hodge structures is a lattice isometry which preserves the (p, q) -grading.

The following theorem is the description of the complex moduli that we alluded to.

Theorem 4.2.6 (Global Torelli theorem). *Two $K3$ surfaces \mathcal{X}_1 and \mathcal{X}_2 are isomorphic as complex manifolds if and only if there is a Hodge isometry between their Hodge structures.*

This manifests itself concretely by examining the so-called period integrals. Consider a marked $K3$ surface, so that we have fixed an isomorphism $H^2(X, \mathbb{Z}) \cong \Gamma_{3,19}$. This amounts to choosing a basis $\{\gamma_i\}$ for $H_2(X, \mathbb{Z})$. We then have a map $\Phi : \mathcal{M}_{\text{cx}}(X) \rightarrow \mathbb{P}(H^2(X, \mathbb{C}))$ determined by

$$\Omega \mapsto \left[\int_{\gamma_1} \Omega : \cdots : \int_{\gamma_{22}} \Omega \right]$$

This is well defined because $\Omega_{\mathcal{X}}^2 \cong \mathbb{C}$, so a holomorphic volume form is determined up to \mathbb{C}^\times -scaling, which projectivising takes care of. The image of this map is called the period domain, which we denote \mathcal{D} , and the Torelli theorem states that $\Phi : \mathcal{M}_{\text{cx}}(X) \rightarrow \mathcal{D}$ is a local isomorphism. As such, its complex dimension is $h^{1,1}(\mathcal{X})$. The only reason why the period map is not an isomorphism onto its image, is that the complex moduli space

⁸This means that $\alpha \cup \alpha \in 2\mathbb{Z}$ for all $\alpha \in H^2(X, \mathbb{Z})$.

4.2. EXAMPLE: MIRROR SYMMETRY FOR $K3$ SURFACES

has two components. When we restrict to a single component, we essentially get that the complex moduli of a $K3$ surface is completely described by its period vector, i.e. the integrals $\int_{\gamma_i} \Omega$.

Mirror symmetry can be understood as an isomorphism between two (variations of) Hodge structures over the complexified Kähler moduli space, and the complex moduli space, respectively. This is the point of view which is presented in the classical textbook on the subject [51]. For a $K3$ surface, this takes the following form when considered pointwise in the respective moduli spaces. Consider the bilinear form on $H^\bullet(X, \mathbb{Z})$ given by

$$\langle \alpha, \beta \rangle = \langle (\alpha_0, \alpha_2, \alpha_4), (\beta_0, \beta_2, \beta_4) \rangle = \int_X \alpha_2 \cup \beta_2 - \alpha_0 \cup \beta_4 - \alpha_4 \cup \beta_0$$

It is called the Mukai pairing, which extends the cup product pairing on $H^2(X, \mathbb{Z})$. As a lattice with the Mukai pairing, $H^\bullet(X, \mathbb{Z})$ is isomorphic to $\Gamma_{3,19} \oplus H$, where H is the lattice which is called the hyperbolic plane, given by \mathbb{Z}^2 with bilinear form $e_1^2 = e_2^2 = 0$ and $e_1 e_2 = 1$. Given a $K3$ surface with complexified Kähler form $\mathbf{B} = B + i\omega$, we can obtain two different Hodge structures on $H^\bullet(X, \mathbb{Z})$: one for the A -model, and one for the B -model. For the A -model, define $\mathcal{U} = \exp(\mathbf{B}) \in H^\bullet(X, \mathbb{C})$. It is easy to verify that one has (cf. the relations satisfied by the holomorphic volume form):

$$\mathcal{U}^2 = \overline{\mathcal{U}}^2 = 0 \quad \mathcal{U}\overline{\mathcal{U}} > 0$$

Define a subspace $H_A^{2,0}(X, \mathbf{B}) \subset H^\bullet(X, \mathbb{C})$ by the complex span of \mathcal{U} , and let $H_A^{0,2}(X, \mathbf{B})$ be its conjugate. Also take $H_A^{1,1}(X, \mathbf{B}) = (H_A^{2,0}(X, \mathbf{B}) \oplus H_A^{0,2}(X, \mathbf{B}))^\perp$, where the orthogonal complement is now with respect to the Mukai pairing. Then we obtain a Hodge structure on $H^\bullet(X, \mathbb{Z})$, denoted by $H_A(X, \mathbf{B})$.

For the B -model, we also obtain a Hodge structure on $H^\bullet(X, \mathbb{Z})$. This is done in analogous fashion, by setting $H_B^{2,0}(\mathcal{X})$ to be the span of the holomorphic volume form. Notice that $H_B^{1,1}(\mathcal{X}) = H^0(X, \mathbb{C}) \oplus H^{1,1}(X, \mathbb{C}) \oplus H^4(X, \mathbb{C})$, it is not just the $(1, 1)$ part of the cohomology. Denote this Hodge structure by $H_B(\mathcal{X})$. Then there is the following theorem.

Theorem 4.2.7 ([52]). *The mirror of a $K3$ surface $(\mathcal{X}, \Omega, \mathbf{B})$ is a $K3$ surface $(\mathcal{X}^\vee, \Omega^\vee, \mathbf{B}^\vee)$ such that there is an isometry of Hodge structures:*

$$H_A(X, \mathbf{B}) \cong H_B(\mathcal{X}) \quad H_A(X^\vee, \mathbf{B}^\vee) \cong H_B(\mathcal{X})$$

Now let $\Lambda \subset H^2(X, \mathbb{Z})$ be a sublattice isomorphic to H in $\Gamma_{3,19}$. Then we can choose a basis $\{f_1, f_2\}$ for this sublattice, satisfying the relations $f_1^2 = f_2^2 = 0$ and $f_1 f_2 = 1$. We also get a decomposition $\Gamma_{3,19} = \Gamma' \oplus \Lambda$, see [53]. Define

$$\mathcal{K} = \{(\Omega, \omega) \mid \Omega \in \mathcal{D}, \quad \omega \in \Omega^\perp \subset H^2(X, \mathbb{R}), \quad \omega^2 > 0\}$$

4.3. THE HOMOLOGICAL MIRROR SYMMETRY CONJECTURE

Theorem 4.2.8 (Huybrechts [53]). *Let $(\Omega, \omega, B) \in \mathcal{K} \times H^2(X, \mathbb{R})$ with $\omega, B \in (\Gamma' \otimes \mathbb{R}) \oplus \mathbb{R}f_1$. Then the Λ -mirror K3 surface is given by*

$$\begin{aligned}\Omega^\vee &= \frac{1}{\langle \text{Re } \Omega, f_1 \rangle} \left(-\frac{1}{2}(B + i\omega + f_2)^2 f_1 + B + i\omega + f_2 \right) \\ \omega^\vee &= \frac{1}{\langle \text{Re } \Omega, f_1 \rangle} \left(\text{Im } \Omega - \langle \text{Im } \Omega, f_2 \rangle f_1 - \langle \text{Im } \Omega, B \rangle f_1 \right) \\ B^\vee &= \frac{1}{\langle \text{Re } \Omega, f_1 \rangle} \left(\text{Re } \Omega - \langle \text{Re } \Omega, f_1 \rangle f_2 - \langle \text{Re } \Omega, f_2 \rangle f_1 - \langle \text{Re } \Omega, B \rangle f_1 \right)\end{aligned}$$

Notice that this depends on a choice of sublattice $\Lambda \subset \Gamma_{3,19}$, which illustrates non-uniqueness of the mirror manifold. We refer to [52] for the string theoretic perspective on this story. It is essentially shown that the lattice which really governs the string theoretic moduli space is $\Gamma_{4,20} \cong H^*(X, \mathbb{Z})$, which is obtained as $\Gamma_{3,19} \oplus H$. The choice of Λ is then a choice of lattice that will be exchanged with $H^0(X, \mathbb{Z}) \oplus H^4(X, \mathbb{Z})$ under the mirror map.

In most cases of interest, a natural choice of Λ can be made. We are referring to those K3 surfaces which admit an elliptic fibration $\mathcal{X} \rightarrow \mathbb{C}\mathbb{P}^1$ and a section $\sigma : \mathbb{C}\mathbb{P}^1 \rightarrow \mathcal{X}$. In this case, one can take f to be the Poincaré dual of the homology class of the fibre, and s that of the section. Then $f^2 = 0$, $fs = 1$ and $s^2 = -2$. Define $f_1 = f$ and $f_2 = f + s$, so that $f_1^2 = f_2^2 = 0$ and $f_1 f_2 = 1$, meaning we can take Λ to be generated by f and s . We can then define the mirror of an elliptic K3 surface using 4.2.8, without ambiguity. In fact, if the elliptic fibration is given w.r.t. a complex structure I , one can hyper-Kähler rotate to the complex structure J , which renders the same fibration a special Lagrangian fibration, and σ a special Lagrangian section for the Kähler form ω_J . Therefore, f_1 and f_2 pair with ω_J and $\text{Im } \Omega_J$ to give 0, and so the mirror K3 surface has a relatively simple expression.

4.3 The Homological Mirror Symmetry Conjecture

Recall that when we outlined how the mirror conjecture was arrived at, we assumed that we were dealing with closed string theories. That is, all of the strings we are considering are diffeomorphic to S^1 . To be more general, we can include open strings, which are diffeomorphic to an interval I . In this case, we still obtain Riemann surfaces as worldsheets, and we can still consider a physical theory which is based around maps $\phi : \Sigma \rightarrow \mathcal{X}$ for some Calabi-Yau threefold \mathcal{X} . This time, the ends of the open strings trace out 1-dimensional objects, and we need to specify some boundary conditions in the target manifold \mathcal{X} for the ends of these strings to live on. These will be D -branes. When we are in a neighbourhood of the large volume limit of the theory, these objects have a geometric interpretation, and this will be our primary interest. Hence, from now on, we will always assume that we are in a neighbourhood of the large volume limit.

4.3. THE HOMOLOGICAL MIRROR SYMMETRY CONJECTURE

Remark 4.3.1. In the discussion that follows, we will assume that the B -field on \mathcal{X} and its mirror manifold both vanish. It should be noted that it is not necessarily true that a vanishing B -field on \mathcal{X} implies a vanishing B -field on its mirror. However, in many concrete examples this is true (see e.g. [53]). Nevertheless, it should be noted that the constructions below have to be modified in cases where the B -field is non-vanishing.

Note that the worldsheet is going to have two different types of boundaries: there are the boundaries traced out by the ends of the open strings, but also the boundary components corresponding to the incoming and outgoing strings. As before, we use conformal invariance to get rid of one type of boundary by inserting certain operators. Then we are looking to put boundary conditions on the boundary segments which are traced out by the ends of the open strings. In [5] it is determined that the appropriate boundary condition for an open string is $\phi(\partial\Sigma) \subseteq L$, where (L, E, ∇) is a submanifold $L \subseteq X$ together with a connection ∇ on a unitary line bundle E over L ,⁹ and stipulate that

$$\frac{\partial\phi^i}{\partial z} = R_j^i(\phi) \frac{\partial\phi^j}{\partial \bar{z}} + \text{Fermionic terms}$$

where R_j^i is an orthogonal matrix.

For the closed string, the $N = (2, 2)$ supersymmetry was generated by so-called supercharges, typically denoted $Q_+, Q_-, \bar{Q}_+, \bar{Q}_-$, which are elements of the aforementioned super Lie algebra. For the open string, it turns out that the full $N = (2, 2)$ supersymmetry cannot be preserved, because translational supersymmetry is broken by the presence of the boundary. To be able to define the topological twisting, which we need to get a mathematically rigorous theory, we need $N = 2$ supersymmetry. Fortunately, one can choose the boundary conditions to be time-independent, in which case time translation supersymmetry can be maintained. More concretely, it means we can take a pair of linear combinations of the supercharges above, which generate an $N = 2$ supersymmetry algebra. There are two different choices of pairs that one can make: these are called A -type supersymmetry, and B -type supersymmetry. They allow for the topological twisting of a supersymmetric sigma model for the open string to obtain the topological A -model and the topological B -model, respectively.

We have just attempted to summarise a very intricate topic in half a page, with the goal being to inform the reader about the main ideas that lead up to the open topological string models. To find out about the details of these constructions, we refer the reader to [5, 4, 49, 54, 42, 55, 56]. We will now instead move on to discuss the main objects of interest to us.

Remark 4.3.2. We stress that we will assume that we are in a neighbourhood of the large volume limit of \mathcal{X} . There is no general understanding of D -branes in the more general setting.

⁹This picture naturally extends to higher rank bundles, which is called "brane stacking".

4.3. THE HOMOLOGICAL MIRROR SYMMETRY CONJECTURE

The analysis in [55] reveals, under the above assumption, which additional conditions are imposed on the D -branes in these models. A D -brane in the A -model (resp. B -model) will be called an A -brane (resp. a B -brane). Then it is shown that an A -brane consists of:

1. A Lagrangian submanifold $L \subset X$.
2. A grading f on L .
3. A spin structure S_L on L .
4. A flat unitary line bundle (E, ∇) on L .

In other words, A -branes are the objects of the Fukaya category.¹⁰ On the other hand, a B -brane consists of¹¹:

1. A complex submanifold $\mathcal{Y} \subseteq \mathcal{X}$.
2. A smooth vector bundle $E \rightarrow Y$.
3. A connection ∇ on E such that F_∇ is of type $(1, 1)$.

As discussed in [8], there is a bijection

$$(E \rightarrow Y, \nabla) \text{ with } F_\nabla \in \Omega^{1,1}(Y) \iff \text{Holomorphic vector bundles } \mathcal{E} \rightarrow \mathcal{Y}$$

In other words, B -branes are complex submanifolds together with holomorphic vector bundles. If $\iota: \mathcal{Y} \rightarrow \mathcal{X}$ is the inclusion, then we get a coherent sheaf $\iota_* \mathcal{E}$ on \mathcal{X} . The support of this sheaf is \mathcal{Y} , so we see that the preliminary B -branes are objects in $\text{Coh}(\mathcal{X})$.

Next, we wish to turn the respective collections of branes into categories. This seems like a reasonable thing to do, geometrically: we view strings as being elements of $\text{Hom}(A, B)$ for some boundary conditions A and B . If we then additionally have a string in $\text{Hom}(B, C)$, we should be able to concatenate these strings, given matching boundary conditions, and obtain an element of $\text{Hom}(A, C)$, i.e an open string spanning from A to C . This is the basic motivation for the following constructions.

The Category of A -branes Let (L_i, E_i, ∇_i) for $i = 0, 1$ be a pair of A -branes, which will be abbreviated simply to L_i . What is the Hilbert space of open string states, for strings beginning on L_0 and ending on L_1 ? We assume that the intersection $L_0 \cap L_1$ is transversal, which can be achieved by using a Hamiltonian perturbation, if necessary. Those A -branes which are related by a Hamiltonian isotopy should be considered as isomorphic A -branes. Morphisms from L_0 to L_1 are interpreted as open strings beginning on L_0 and ending on L_1 . The space of morphisms is $\text{Hom}(L_0, L_1) := CF^*(L_0, L_1)$, which can be turned into a chain complex $(\text{Hom}(L_0, L_1), Q) = (CF^*(L_0, L_1), \partial)$ where Q is the BRST


¹⁰There is an "unobstructedness" criterion for A -branes, akin to the one we mentioned for objects in the Fukaya category. We omit it here as well.

¹¹This is a preliminary definition that we will refine below.

4.3. THE HOMOLOGICAL MIRROR SYMMETRY CONJECTURE

operator, presently given by the Floer differential.¹² The cohomology of this complex (i.e. the Floer cohomology) is the "true" Hilbert space of open string states from L_0 to L_1 . The action of taking cohomology can be interpreted as quotienting out gauge equivalences. That is, if two states differ by a Q -exact term, then they are gauge equivalent, so we should really consider them as equal in the Hilbert space. Elements of the Hilbert space correspond to certain operators, under the operator-state correspondence. Physicists would like these operators to form an algebra, at the level of Q -cohomology.

For technical reasons explained in [5], one wants to preserve the information contained in the chain complexes. We know the solution: instead of passing to cohomology, we consider A_∞ -categories up to quasi-equivalence. We take $m_1 = Q$, and the higher multiplication maps m_k are determined by the $(k+1)$ -point correlation functions¹³ of the theory and non-degenerate forms on the morphism spaces, which can be obtained from the geometry of the Calabi-Yau manifold. Defining these forms rigorously is the subject of much technical analysis, for the A -side. Given $a_i \in \text{Hom}(L_i, L_{i+1})$ we consider the $(k+1)$ -point function, denoted $\langle a_0 \dots a_k \rangle$, and we define $m(a_k, \dots, a_1)$ by $\langle m_k(a_k, \dots, a_1), a_0 \rangle = \langle a_k \dots a_0 \rangle$, where $\langle \cdot, \cdot \rangle$ denotes the non-degenerate bilinear form on $\text{Hom}(L_0, L_k)$. Remarkably, one can show (see [56]) that these m_k satisfy the A_∞ -relations, and in fact yield the Fukaya category (up to some technicalities, such as changing the field from \mathbb{C} to Λ). The cohomology category $H^*(\text{Fuk}(X, \omega))$ has a grading on its morphism spaces. The degree of an element in the Hilbert space is called the ghost number. Mathematically, these ghost numbers are obtained via the Maslov index from C.2. The pseudo-holomorphic disks which define the Floer differential (i.e. the BRST operator) are interpreted as the process of quantum tunneling from one intersection point to another. The higher order multiplication maps have a similar geometric interpretation through quantum tunneling. For instance, consider $m_2 : CF^*(L_1, L_2) \otimes CF^*(L_0, L_1) \rightarrow CF^*(L_0, L_2)$ which is derived from the 3-point function. Its contributions come from maps $u : \mathbb{D} \setminus \{q, p_1, p_2\} \rightarrow X$ as in the following diagram:



On the left hand side of the figure, we have the so-called pair of pants (for open strings), which physically represents an interaction in which the string state q splits into p_1 and p_2 . One can evaluate the 3-point function of the theory on these states, once they have been mapped into \mathcal{X} , and this contributes to the multiplication map m_2 . Geometri-

¹²As the notation suggests, this operator Q comes from the supercharges.

¹³We can think of these as giving the vacuum expectation values of certain observables in the quantum theory.

4.3. THE HOMOLOGICAL MIRROR SYMMETRY CONJECTURE

cally, this can be seen as the conformal equivalence between the pair of pants and the thrice-punctured disk, which gives us the contribution to m_2 in the sense of the Fukaya category, see 3.2. Taking the derived Fukaya category can be interpreted physically as allowing certain states to bind together or decay into each other, we refer to [5] for details.

The Category of B -branes Next, we wish to turn B -branes into a category. The first thing to note is that our original definition does not yield enough B -branes, particularly if we wish mirror symmetry to be true. The way to solve this is by considering the ghost number, which in the A -model required the Maslov index of an intersection point $p \in L_0 \cap L_1$ of two Lagrangians. To eventually include more objects into our category of B -branes, we will (paradoxically) first restrict our attention to locally free sheaves on \mathcal{X} , i.e. holomorphic vector bundles. The intricacies of the morphisms between two B -branes will reveal that by analysing this case, we have in fact already included the more general types of B -branes as well. One can show that the BRST operator Q for the B -model may be identified with $\bar{\partial}$, which leads to the Hilbert space of open string states between \mathcal{E} and \mathcal{F} , namely $\text{Hom}(\mathcal{E}, \mathcal{F}) = \oplus_q H^{0,q}(\mathcal{X}, \text{Hom}(\mathcal{E}, \mathcal{F})) = \oplus_q \text{Ext}^q(\mathcal{E}, \mathcal{F})$. This space already comes equipped with a natural integer grading, so it is tempting to say that the integer q is the ghost number. However, given the lengths one needs to go to on the A -model side, let us instead assign an integer $\mu(\mathcal{E})$ to every holomorphic vector bundle on \mathcal{X} . Then we will say that an element of $H^{0,q}(\mathcal{X}, \text{Hom}(\mathcal{E}, \mathcal{F}))$ has ghost number $q + \mu(\mathcal{F}) - \mu(\mathcal{E})$, by analogy with the Maslov index. Then we can consider $\mathcal{E}_\bullet = \oplus_k \mathcal{E}_k$ where \mathcal{E}_k is a holomorphic vector bundle with $\mu(\mathcal{E}_k) = k$.

To enlarge the class of objects we are considering, we will deform the objects. It is ghost number one operators which correspond to deformations of the theory, i.e. elements of $\text{Ext}^q(\mathcal{E}_k, \mathcal{E}_{k-q+1})$. Taking $q = 1$ corresponds to $\text{Ext}^1(\mathcal{E}_k, \mathcal{E}_k)$ which are sheaf deformations of \mathcal{E}_k , as discussed in 5.3.1. Taking $q = 0$, we get morphisms $d_k : \mathcal{E}_k \rightarrow \mathcal{E}_{k+1}$. We denote $d = \sum_k d_k$. Under the operator-state correspondence, these open string states yield operators W , which can be used to deform the action of the theory by adding a boundary term $S \mapsto S + \int_{\partial\Sigma} W$. As it turns out (see again [55]), to keep the condition that the BRST operator of the deformed theory squares to zero (which is required for physical consistency), we must have $d^2 = 0$. In other words, (\mathcal{E}_\bullet, d) is a complex of locally free sheaves, and this is really a better way to think about a B -brane (for $\mathcal{Y} = \mathcal{X}$), because we cannot assume that the B -model inherently knows the "right" ghost number, while the A -model does not. Thus, imposing this "ignorance" on the B -model, we are naturally led to consider complexes of locally free sheaves. We still include holomorphic vector bundles $\mathcal{E} \rightarrow \mathcal{X}$ as B -branes, since these are complexes in a trivial way. The Hilbert space of open string states spanning from \mathcal{E}_\bullet to \mathcal{F}_\bullet is then given by $\oplus_q \text{Ext}^q(\mathcal{E}_\bullet, \mathcal{F}_\bullet)$, and this time, the integer q really is the ghost number. However, once again, rather than passing to cohomology, we will take the category of B -branes to be $D_\infty^b(\mathcal{X})$ (the differential graded category of perfect complexes). That is, we only obtain the Hilbert spaces of open string states after taking the cohomology category $H^\bullet(D_\infty^b(\mathcal{X}))$.

As we know, $H^\bullet(D_\infty^b(\mathcal{X})) \cong D^b(\mathcal{X})$. So the question is: how did we end up in the

4.3. THE HOMOLOGICAL MIRROR SYMMETRY CONJECTURE

derived category, if we started out by considering complexes of vector bundles? Mathematically, we already know this. Physically, the answer to this question is given by the notion of an isomorphism between the branes, which will tell us that we need to invert quasi-isomorphisms. In this way, it turns out that the category of B -branes really is $H^*(D_\infty^b(\mathcal{X})) = D^b(\mathcal{X})$. Suppose that $\mathcal{E}_\bullet \rightarrow \mathcal{E}'_\bullet$ is a quasi-isomorphism, and suppose that $\mathcal{E}'_\bullet \rightarrow \mathcal{I}_\bullet$ is an injective resolution (i.e. a quasi-isomorphism with a chain complex of injective objects). The latter quasi-isomorphism may lie in the category of chain complexes of \mathcal{O}_X -modules, for instance. Then $\mathcal{E}_\bullet \rightarrow \mathcal{I}_\bullet$ is also an injective resolution, and so for all chain complexes \mathcal{F}_\bullet of (locally free) sheaves, we get

$$\mathrm{Ext}^q(\mathcal{E}_\bullet, \mathcal{F}_\bullet) \cong \mathrm{Ext}^q(\mathcal{E}'_\bullet, \mathcal{F}_\bullet) \quad \mathrm{Ext}^q(\mathcal{F}_\bullet, \mathcal{E}_\bullet) \cong \mathrm{Ext}^q(\mathcal{F}_\bullet, \mathcal{E}'_\bullet)$$

Thus, the resulting Hilbert spaces are isomorphic with respect to an arbitrary B -brane \mathcal{F}_\bullet , and we should consider \mathcal{E}_\bullet and \mathcal{E}'_\bullet isomorphic as B -branes because they are physically indistinguishable. That is, we should invert quasi-isomorphisms. In conclusion, the category of B -branes is really $D^b(\mathcal{X})$. Its A_∞ -structure is given by the dg-structure on $D_\infty^b(\mathcal{X})$ and the minimal model [D.5.6](#).

This also remedies the issue we caused by restricting to locally free sheaves on \mathcal{X} . Namely, we now have complexes of arbitrary sheaves, which are supported on submanifolds $\iota: \mathcal{Y} \hookrightarrow \mathcal{X}$. Given a locally free sheaf $\mathcal{E} \rightarrow \mathcal{Y}$, it is tempting to think of $\iota_*\mathcal{E}$ as being the corresponding element in the category of B -branes. This is not true, and this result is called the Freed-Witten anomaly. It states that when \mathcal{Y} does not admit a Spin-structure, the bundle over \mathcal{Y} which corresponds to $\iota_*\mathcal{E}$ as a B -brane, is not in fact \mathcal{E} itself, but a twisted version (i.e. tensored with a line bundle). Let us not be too concerned about this, as it does not detract from the fact that by considering the locally free sheaves on \mathcal{X} itself, we have in fact also obtained the B -branes which are supported on proper submanifolds. See [\[57\]](#) for more on this. As for the A -model, the higher multiplication maps m_k on the morphisms in $D^b(\mathcal{X})$ are determined by the k -point correlation functions of the theory and the non-degenerate forms on the Hilbert spaces, see above and also [\[58\]](#). In this case, the non-degenerate forms come from Serre duality on $D^b(\mathcal{X})$.

Given that the mirror conjecture was based on an isomorphism between the A -model with target space (\mathcal{X}, g) , and the B -model with target space $(\mathcal{X}^\vee, g^\vee)$, their respective categories of boundary conditions should then also be isomorphic. Thus, the mathematical statement of this conjecture becomes:

Conjecture 4.3.3 (Kontsevich). *Let (\mathcal{X}, g) and $(\mathcal{X}^\vee, g^\vee)$ be a mirror pair of Calabi-Yau threefolds. Then there is a quasi-equivalence of A_∞ -categories*

$$D^b(\mathcal{X}) \cong D^b \mathrm{Fuk}(X^\vee, \omega^\vee) \quad D^b(\mathcal{X}^\vee) \cong D^b \mathrm{Fuk}(X, \omega)$$

Shockingly, the idea of D -branes did not yet exist when Kontsevich formulated his conjecture. Furthermore, it is not yet known whether the homological mirror symmetry conjecture (as it is formulated in [0.0.1](#)) implies the mirror conjecture, but it is suspected

4.4. EXAMPLE: HOMOLOGICAL MIRROR SYMMETRY FOR COMPLEX TORI

that it does, since open string theories include more data than closed string theories. Homological mirror symmetry is known to hold more generally, namely for Calabi-Yau hypersurfaces in $\mathbb{C}\mathbb{P}^N$. The case $N = 2$ corresponds to an elliptic curve, which was originally treated in [59]. The case $N = 3$ is the quartic surface, for which Seidel proved homological mirror symmetry in [60]. The case $N \geq 4$ was covered by Sheridan in [61]. This result covers the quintic threefolds, for which mirror symmetry was originally established.

One thing that should immediately stand out to us, is the following. The derived Fukaya category (i.e. the A -model) is very difficult to work with. There are many technicalities to be addressed and it is unpleasant to attempt to compute things directly in this setting. On the other hand, the derived category of coherent sheaves (i.e. the B -model) is much more pleasant to work with. This is the realm of algebraic geometry, in which there are methods to calculate things concretely. As we mentioned before, the derived equivalence of categories typically does not arise from an equivalence of the underlying categories. Specifically, there is no equivalence of categories between $\text{Fuk}(X, \omega)$ (or its twisted version) and $\text{Coh}(\mathcal{X}^\vee)$. However, one might hope that we can still translate between the symplectic geometry of \mathcal{X} , and the algebraic geometry on the mirror \mathcal{X}^\vee , as was done when physicists computed the Gromov-Witten invariants of the quintic via period integrals on the mirror manifold. Especially when we have something which is defined on the heart of the t -structure on $D^b(\mathcal{X}^\vee)$ (see D.4.4) and can be extended to the entire derived category, we might hope that this has an analogue on the symplectic side. This is part of the motivation of the Thomas-Yau conjecture, which we will discuss in the next part of the text.

4.4 Example: Homological Mirror Symmetry for Complex Tori

As noted, homological mirror symmetry was shown to hold for Calabi-Yau hypersurfaces in $\mathbb{C}\mathbb{P}^N$. This is the most general class of Calabi-Yau manifolds for which homological mirror symmetry is known to hold, but it does not contain those Calabi-Yau manifolds which are easiest to work with: complex tori (except the elliptic curve). In fact, Abouzaid showed that homological mirror symmetry holds on 4-tori (see [62]), but the general case remains more elusive. In spite of this, most objects of interest can be written down quite explicitly, and so complex tori provide a nice class of examples to study, as per [63].


4.4.1 The Elliptic Curve

We start by considering an elliptic curve $\mathcal{X} = \mathbb{C}/\Lambda$ where Λ is some lattice. In this setting, it is quite evident how to define the mirror morphism at the level of objects, because this is one of the few cases in which the entire Fukaya category is well-understood. The

4.4. EXAMPLE: HOMOLOGICAL MIRROR SYMMETRY FOR COMPLEX TORI

Calabi-Yau metric on a complex torus is the natural flat metric, induced by the Euclidean metric on \mathbb{C} . Taking smooth coordinates x, y , the metric is just $A(dx \otimes dx + dy \otimes dy)$ for some $A \in \mathbb{R}_{>0}$. The complex structure acts via $\frac{\partial}{\partial x} \mapsto \frac{\partial}{\partial y}$ and $\frac{\partial}{\partial y} \mapsto -\frac{\partial}{\partial x}$. As such, we get $dx \mapsto -dy$ and $dy \mapsto dx$ under the pullback. This means the Kähler form, defined as $\omega(v, w) = g(v, Jw)$, is given by $A(-dx \otimes dy + dy \otimes dx) = Ady \wedge dx = \frac{A}{2i} dz \wedge d\bar{z}$. We choose the holomorphic volume form to be $\Omega = \sqrt{A}dz$, so that $\frac{1}{2i}\Omega \wedge \bar{\Omega} = \omega$. These are the standard conventions. A complexified Kähler form is then given as $\mathbf{B} = B + i\omega$, and we write its cohomology class as $\rho = B + iA$ (because $\int_X \omega = A$). The mirror elliptic curve $\mathbb{C}/\tilde{\Lambda}$ is given by $\tilde{\tau} = \rho$ and $\tilde{\rho} = \tau$. But what does this mean for homological mirror symmetry?

As a first step, we should identify the Lagrangian submanifolds - but this is trivial in complex dimension 1, they are just the submanifolds of dimension 1. We will be interested in closed submanifolds, and so we look for embeddings $\gamma : S^1 \rightarrow \mathbb{C}/\Lambda$. We write the image as $\gamma(S^1) = L_\gamma$. Now that we know what the Lagrangians look like, we need to determine which ones admit a grading. That is, we need to identify the Lagrangians whose Maslov class vanishes. We distinguish between two cases: when γ is null-homotopic, and when it is not. Suppose first that γ is null-homotopic, say a circle in a fundamental domain $U \subset \mathbb{C}$. This determines a map $f : L_\gamma \rightarrow \text{LGr}(1)$ by sending a point $t \in L_\gamma$ to $T_t L_\gamma \in \text{LGr}(1)$. Geometrically, this is the following map:



The arrows indicate the direction in which we move in $\text{LGr}(1)$, as we go around L_γ counterclockwise. We can see that f winds $\text{LGr}(1)$ twice, so the Maslov class of L_γ is non-trivial. As such, these Lagrangian submanifolds cannot be graded.

On the other hand, we have Lagrangian submanifolds which represent non-trivial loops. In \mathbb{C} , we may perturb these to straight lines (with rational slope) without altering the homotopy class of f . It is clear that these straight lines correspond to the constant map $f : \mathbb{R} \rightarrow \text{LGr}(1)$ under the canonical identification $T_t L_\gamma \subset T_t \mathbb{C} \cong \mathbb{C}$. Therefore, all the homotopically non-trivial Lagrangian submanifolds admit a grading.

So it is quite clear what the Lagrangian submanifolds of the elliptic curve look like. In

4.4. EXAMPLE: HOMOLOGICAL MIRROR SYMMETRY FOR COMPLEX TORI

fact, there is a specific class of Lagrangian submanifolds which is distinguished, namely the geodesics. In our discussion leading up to the Thomas-Yau conjecture, we will see that geodesics are the first incarnation of special Lagrangian submanifolds, and these objects are really of fundamental importance. We can in fact define the Fukaya category of the elliptic curves by only considering geodesics as objects, and excluding the more general Lagrangian submanifolds, even if they are graded.

Before we do so, we need to equip each L_γ with a complex line bundle and a flat unitary connection. Every complex line bundle over $L_\gamma \cong S^1$ is necessarily trivial, which means that a flat connection on it is completely specified by $\nabla = d - 2\pi i \beta dx$ (where we assume for convenience that the x -coordinate parameterises L_γ). Taking our objects to be geodesics, then, any object in the Fukaya category is isomorphic to an object which is completely specified by three numbers (y_0, α, β) , where y_0 is the y -coordinate of the intersection of $\gamma(t)$ with the y -axis in \mathbb{C} , α is the grading¹⁴ of $\gamma(t)$, and β specifies the flat connection ∇ . Establishing mirror symmetry at the level of objects is not too difficult now (with the B -field turned off - although turning it on would only require mild adaptations). Denote $(L, E, \nabla) := L_{(y_0, \alpha, \beta)}$. Then homological mirror symmetry at the level of objects is, essentially, the map

$$L_{(y_0, \alpha, \beta)} \mapsto (\mu_{-iy_0 + \beta}^* \mathcal{L}) \otimes \mathcal{L}^{k-1} := \mathcal{L}(y_0, \alpha, \beta)$$

for some distinguished holomorphic line bundle $\mathcal{L} = \mathcal{O}(p_0)$ of degree 1. Here, $\mu_z : \mathbb{C}/\tilde{\Lambda} \rightarrow \mathbb{C}/\tilde{\Lambda}$ is the translation map, i.e. the group law on the mirror elliptic curve. Technically, one would also have to specify where the skyscraper sheaves are sent, and consider twisted objects in the Fukaya category, but the above correspondence is the main intuition. Higher rank vector bundles also appear, as the transforms of special Lagrangians with non-integral slope. This does not say anything about the mirror functor at the level of morphisms, and doing so would be significantly more involved, as it is most conveniently expressed through the theory of θ -functions. We refer to [59] for the original proof.

4.4.2 Interlude: The Ströminger-Yau-Zaslow Conjecture

We now want to consider the case $\mathcal{X} = \mathbb{C}^n/\Lambda$. First, we need to establish what its mirror manifold is. There is no obvious way of doing so, as there was for $n = 1$. But there is a more geometric approach to mirror symmetry which, in the present case at least, allows for an explicit construction of a mirror manifold. We will outline the procedure for complex tori, and then explain how it is motivated by homological mirror symmetry, and what the more general picture is. We follow [64, 63].

We start by considering a real symplectic torus $X = V/\Lambda$ where V is some real vector space of dimension $2n$, together with a complexified Kähler form $\mathbf{B} = B + i\omega$. We take a

¹⁴The slope is to be thought of as a complex phase, with rational tangency, and hence can be written as $\exp(\pi i \alpha)$. A choice of grading is a choice of α .

4.4. EXAMPLE: HOMOLOGICAL MIRROR SYMMETRY FOR COMPLEX TORI

linear Lagrangian subspace $L \subset V$, and assume that $L \cap \Lambda \otimes \mathbb{R} = L$. Then we may obtain a Lagrangian fibration $\pi : V/\Lambda \rightarrow V/(L + \Lambda)$. Consider the map

$$\Phi : V \oplus V^* \rightarrow \text{Hom}_{\mathbb{R}}(V, \mathbb{C}) \quad (w, \eta) \mapsto (v \mapsto \mathbf{B}(v, w) + \eta(v))$$

Suppose that $\Phi(w, \eta) = 0$. Then $\text{Im } \Phi(w, \eta)(v) = \omega(v, w) = 0$ for all $w \in V$, and hence $w = 0$. As such, Φ is an injective map between vector spaces of equal dimension $2n$, so it is an isomorphism. Since $\Phi(L \oplus L^\perp) = \text{Hom}_{\mathbb{R}}(V/L, \mathbb{C})$, we get an isomorphism

$$V/L \oplus L^* \cong \text{Hom}_{\mathbb{R}}(L, \mathbb{C})$$

which gives a complex structure to $V^\vee = V/L \oplus L^*$ in a natural way, with a lattice $\Lambda^\vee = \Lambda/(\Lambda \cap L) \oplus (\Lambda \cap L)^\perp$. The mirror manifold to $(V/\Lambda, \mathbf{B})$ is then given by the complex torus $\mathcal{X}^\vee = V^\vee/\Lambda^\vee$. But why is this a reasonable construction? We will outline the geometric argument for this below.

We can view the symplectic torus $(V/\Lambda, \omega)$ as the quotient $T^*M/T_{\mathbb{Z}}^*M$ where $M = V/(L + \Lambda)$ is the base space, i.e. a real n -torus. We have denoted by $T_{\mathbb{Z}}^*M$ a (constant) lattice in T^*M . Thus, $T^*M/T_{\mathbb{Z}}^*M$ amounts to taking the quotient of \mathbb{R}^n by a lattice in each fibre, resulting in an n -torus fibration V/Λ over M . The choice of L can be viewed as the choice of cotangent fibres of M , which gives a splitting

$$\Lambda = (\Lambda \cap L) \oplus (\Lambda \cap L)^\perp := \Lambda_f \oplus \Lambda_b$$

The quotient L/Λ_f gives the fibre. The Lagrangian fibration $\pi : X = V/\Lambda \rightarrow V/(L + \Lambda)$ from above is now just $T^*M/T_{\mathbb{Z}}^*M \rightarrow M$ w.r.t. the canonical symplectic structure on T^*M . The mirror manifold was obtained as

$$X^\vee = V^\vee/\Lambda^\vee = \frac{V/L \oplus L^*}{(\Lambda/\Lambda_f) \oplus \Lambda_b}$$

Recall that $L^\perp \cong L^*$ naturally, via the symplectic form, which is why we can view the second summand as a lattice in L^* . It is clear that $(V/L)/(\Lambda/\Lambda_f) = M$, so we in fact get a torus fibration $\pi^\vee : X^\vee \rightarrow M$. The difference lies in the fact that the fibres are now $L^*/(\Lambda \cap L)^\perp$, which is the dual torus of the fibre of $\pi : X \rightarrow M$. Using the isomorphism Φ constructed from \mathbf{B} together with the Euclidean metric, X^\vee acquires a complex structure \mathcal{X}^\vee .

Definition 4.4.1. The pair $(\pi : X \rightarrow M, \mathbf{B})$ and $\pi^\vee : \mathcal{X}^\vee \rightarrow M$ is called an SYZ dual fibration.

The SYZ construction from [19] is based on the following idea. We have the skyscraper sheaves \mathcal{O}_x for $x \in \mathcal{X}$. The space of isomorphism classes of skyscraper sheaves is just the manifold \mathcal{X} , since such skyscraper sheaves are given by specifying a single point. Skyscraper sheaves are BPS branes¹⁵ for type IIA string theory. Under mirror symmetry,

¹⁵We will discuss BPS branes in greater detail in the second part of the text. For now: we can think of a BPS brane as a D -brane for the full string theory which preserves some supersymmetry. They are the D -branes that appear in the topological string theories, but which also arise in the genuine string theory, and so they are also called physical A -branes (resp. physical B -branes).

4.4. EXAMPLE: HOMOLOGICAL MIRROR SYMMETRY FOR COMPLEX TORI

the type IIA model on \mathcal{X} is exchanged with the type IIB model on \mathcal{X}^\vee . A skyscraper sheaf is mapped to some BPS brane for the type IIB string theory on \mathcal{X}^\vee , which is a special Lagrangian submanifold $L \subset X^\vee$ together with a flat unitary line bundle on it. These branes can be deformed into nearby BPS branes, and the space which parameterises these deformations is the moduli space of the brane. Hence, there is a BPS brane in \mathcal{X}^\vee whose moduli space is \mathcal{X} . It turns out that deformations of special Lagrangian submanifolds are given by harmonic 1-forms on it, so we must have $b_1(L) = n$ for the dimensions to match, and the authors determine L to be a special Lagrangian n -torus, see also [65] for a more rigorous explanation.

Given a torus, say $\mathbb{T} = W/\Gamma$ for some real vector space W and lattice Γ , we may view $\Gamma = H_1(\mathbb{T}, \mathbb{Z})$, and so $\Gamma^* = H^1(\mathbb{T}, \mathbb{Z})$. Then $W^* = H^1(\mathbb{T}, \mathbb{R})$ and one has $W^*/\Gamma^* = H^1(\mathbb{T}, \mathbb{R})/\Gamma^* \cong H^1(\mathbb{T}, \mathbb{R}/\mathbb{Z})$. But $H^1(\mathbb{T}, \mathbb{R}/\mathbb{Z}) = \text{Hom}(\pi_1(\mathbb{T}), \mathbb{R}/\mathbb{Z})$. Hence, W^*/Γ^* , the dual torus, may be viewed as the moduli space of flat unitary line bundles on \mathbb{T} . Because of this, the moduli space of flat unitary line bundles on any fixed Lagrangian torus L is given by $H^1(L, \mathbb{R}/\mathbb{Z}) \cong T^n$ (one may have to allow singularities, if L is immersed but not embedded). Say the moduli space of a given special Lagrangian torus $L \subset X^\vee$ is M . We argued that the moduli space of the full BPS brane is X , so we need to include deformations of the flat unitary line bundles, if we want to obtain X from M . By the preceding discussion, these deformations fiber as n -tori over M . The conclusion is that X may be expressed as a T^n -fibration $\pi : X \rightarrow M$, where M is the moduli space associated to a special Lagrangian n -torus in X^\vee and the fibres of π are the moduli spaces of flat unitary line bundles. Furthermore, SYZ argue that each fibre $\pi^{-1}(m) \subset X$ is itself a special Lagrangian torus in X , so that $\pi : X \rightarrow M$ is a special Lagrangian fibration.

Definition 4.4.2. Let $\pi : X \rightarrow M$ be a special Lagrangian fibration of a Calabi-Yau manifold, such that the Calabi-Yau metric is flat along the fibres. Let $M_0 = M \setminus M_{\text{sing}}$ be the complement of the singular locus of M , and $X_0 = \pi^{-1}(M_0)$. Let X_0^\vee be the dual torus fibration over M_0 , obtained using the flat metric on each fibre. Then an SYZ mirror of $\pi : X \rightarrow M$ is a compactification of X_0^\vee .

The above definition is purely at a topological level. We used the Calabi-Yau structure on X to define a special Lagrangian torus fibration, but we did not say how to obtain a Calabi-Yau structure on X^\vee . Doing so would be considerably more involved, and there is currently no mathematically rigorous statement on how to obtain the Calabi-Yau metric on the mirror manifold in general. The semi-flat case, i.e. when $M = M_0$, does allow for a description of the Calabi-Yau structure on the mirror manifold, see [5]. This is precisely the construction we gave above for complex tori.

Remark 4.4.3. In [66] it is shown that Calabi-Yau manifolds in complex dimension ≤ 3 which satisfy some regularity assumptions have a well-behaved dual torus fibration. It is also shown that the dual torus fibration of the quintic 3-fold coincides with a model of the mirror quintic, providing strong evidence in favour of the SYZ approach. It serves as a testing ground for ideas, for example in [67] where the ideas of homological mirror symmetry are tested for an SYZ dual pair, and are indeed shown to hold. We will

4.4. EXAMPLE: HOMOLOGICAL MIRROR SYMMETRY FOR COMPLEX TORI

comment more on this later.

The main appeal of the SYZ approach is that it gives an intrinsic way to construct a mirror manifold X^\vee at the topological level (and conjecturally for the full Calabi-Yau structure). The first step is to find a special Lagrangian torus fibration. This was easy to do for a torus, but more generally it is a non-trivial matter, and this is where the Thomas-Yau conjecture may be of use. Before we go on to discuss the Thomas-Yau conjecture, we will discuss homological mirror symmetry for complex tori.

4.4.3 Homological Mirror Symmetry and Complex Tori

Having discussed the SYZ approach to mirror symmetry, we know how to construct a mirror of a symplectic torus ($X = V/\Lambda, \mathbf{B} = B + i\omega$). We choose a linear Lagrangian subspace $L \subset V$, and then the mirror torus is given by

$$\mathcal{X}^\vee = \frac{V/L \oplus L^*}{(\Lambda \cap L) \oplus (\Lambda \cap L)^\perp}$$

together with the flat metric, and the complex structure is determined by \mathbf{B} . We would like to get a feeling for homological mirror symmetry on these spaces, and this is done by using the smooth Fourier-Mukai transform, which we will now explain. For simplicity, we assume that $\mathbf{B} = i\omega$, so the B -field is turned off.

As noted previously, given a torus $\mathbb{T} = W/\Gamma$, the dual torus $\mathbb{T}^* = W^*/\Gamma^*$ may be viewed as the moduli space of unitary line bundles on it, so we denote its points by pairs (E, ∇) with $E \rightarrow \mathbb{T}$ a flat unitary line bundle. There is a canonical line bundle $\mathcal{P} \rightarrow \mathbb{T} \times \mathbb{T}^*$ together with a connection $\nabla^{\mathcal{P}}$ (which is not flat). The restriction $\mathcal{P}|_{\mathbb{T} \times \{(E, \nabla)\}}$ is given by E , and $\nabla^{\mathcal{P}}|_{\mathbb{T} \times \{(E, \nabla)\}} = \nabla$. Strictly speaking this is not a definition, but the intuition is clear. Next we take the complex torus $\pi : \mathcal{X} \rightarrow B$ and its mirror $\pi^\vee : \mathcal{X}^\vee \rightarrow B$. Apply the construction to the Lagrangian fibres of π . Then we get a line bundle with connection $\mathcal{P}_B \rightarrow X \times_B X^\vee$.

Definition 4.4.4. Let $\iota : L \rightarrow X$ be a compact Lagrangian submanifold which is transversal to the fibres of π (but not necessarily a section), together with a flat line bundle (E, ∇) on it. Let $p : L \times_B X^\vee \rightarrow X^\vee$ and $q : L \times_B X^\vee \rightarrow L$ be the natural projection maps. Then the smooth Fourier-Mukai transform of (L, E, ∇) is a vector bundle on X^\vee

$$\mathcal{F}(L, E, \nabla) = p_* \left((\iota \times \text{id})^* \mathcal{P}_B \otimes q^* E \right)$$

together with its natural connection.¹⁶

So we get a vector bundle with connection $\mathcal{F}(L, E, \nabla)$ on X^\vee . The rank of $\mathcal{F}(L, E, \nabla)$ is the cardinality of L in a generic fibre of π . The following is proved in [68].

¹⁶Since we use the pushforward p_* , it is not immediate that we get a connection. However, p is a proper, unramified covering map, and for this reason, the pushforward of a connection along p is well-defined.

4.5. THE B -FIELD IN HOMOLOGICAL MIRROR SYMMETRY

Theorem 4.4.5. *The curvature of the connection on $\mathcal{F}(L, E, \nabla)$ has vanishing $(0, 2)$ -component w.r.t. the complex structure \mathcal{X}^\vee . Hence, $\mathcal{F}(L, E, \nabla)$ defines a holomorphic bundle on \mathcal{X}^\vee .*

The authors also prove that in the case of the elliptic curve, every holomorphic vector bundle on \mathbb{C}/Λ can be obtained in this way, so this construction generalises what we discussed for the elliptic curve previously. This sets up the mirror functor for a large class of objects in the Fukaya category. For the morphisms, we refer to [63], where it is also shown that Hamiltonian isotopic A -branes yield isomorphic holomorphic bundles on \mathcal{X}^\vee . Of course, the discussion at the level of morphisms is once again much more intricate and technical, which is why we choose to omit it.

4.5 The B -field in Homological Mirror Symmetry

In the discussion of homological mirror symmetry that we presented above, we essentially neglected the B -field, in the sense that we assumed $B = 0$ on the Calabi-Yau manifold, as well as its mirror. In [48], it is argued that homological mirror symmetry does not work without including the B -field. This makes sense, given the fundamental role that the B -field plays in mirror symmetry itself: if we would only consider the Kähler cone, there could never be a local isomorphism between it, and the complex moduli space. Let alone a holomorphic one. So one should include the B -field when talking about homological mirror symmetry as well. The B -field is really a cohomology class $[B] \in H^2(X, \mathbb{R}/\mathbb{Z})$ which lies in the kernel of $H^2(X, \mathbb{R}/\mathbb{Z}) \rightarrow H^3(X, \mathbb{Z})$ and can hence be lifted to $[B] \in H^2(X, \mathbb{R})$ and then represented by a 2-form $B \in \Omega^2(X)$, which we have been calling the B -field. In the same way that $H^1(X, \mathbb{R}/\mathbb{Z})$ classifies flat $U(1)$ -bundles, the cohomology group $H^2(X, \mathbb{R}/\mathbb{Z})$ classifies so-called flat $U(1)$ -gerbes. Point particles acquire an electric charge from pulling back a connection 1-form from a $U(1)$ -bundle to their worldline. Strings acquire an electric charge from pulling back a connection 2-form on a $U(1)$ -gerbe to their worldsheet. In the case of a trivial $U(1)$ -bundle, the connection may be represented as $\nabla = d + A$, so the 1-form A determines the (flat) connection. For a trivial $U(1)$ -gerbe, the 2-form B characterises a connection, and this is what the B -field represents: a higher analogue of the electric potential.

The A -model The category of A -branes does not change very much when the B -field is turned on. Instead of considering graded (unobstructed) Lagrangians with flat unitary line bundles on them, one considers graded Lagrangians with unitary line bundles such that

$$F_\nabla = 2\pi i B|_L$$

The B -model The category of B -branes changes quite substantially when the B -field is turned on. Instead of considering coherent sheaves on \mathcal{X} , one considers B -twisted sheaves on \mathcal{X} . Consider the canonical inclusion of sheaves $\mathbb{R}/\mathbb{Z} \hookrightarrow \mathcal{O}_{\mathcal{X}}^\times$. Then we may also view B as a class in $B \in H^2(\mathcal{X}, \mathcal{O}_{\mathcal{X}}^\times)$. Such a class may be represented by a 2-cocycle

4.5. THE B -FIELD IN HOMOLOGICAL MIRROR SYMMETRY

on \mathcal{X} , i.e. an open cover $\{U_i\}$ of \mathcal{X} together with sections $B_{ijk} \in \Gamma(U_i \cap U_j \cap U_k, \mathcal{O}_{\mathcal{X}}^\times)$ which represent B . That is, we require

$$\delta(B)_{ijkl} := B_{jkl}|_{U_{ijkl}} - B_{ikl}|_{U_{ijkl}} + B_{ijl}|_{U_{ijkl}} - B_{ijk}|_{U_{ijkl}} = 0$$

so that the chosen sections indeed yield a cohomology class, which should be B .

Definition 4.5.1. Let $B \in H^2(\mathcal{X}, \mathcal{O}_{\mathcal{X}}^\times)$ be represented by a 2-cocycle B . Then a B -twisted sheaf on \mathcal{X} is a collection $\{(\mathcal{E}_i)_{i \in I}, (\varphi_{ij})_{i, j \in I}\}$ where \mathcal{E}_i is a sheaf on $U_i \subseteq X$ and $\varphi_{ij} : \mathcal{E}_i|_{U_{ij}} \rightarrow \mathcal{E}_j|_{U_{ij}}$ is an isomorphism of sheaves, satisfying

$$\varphi_{ij} \circ \varphi_{jk} \circ \varphi_{ki} = B_{ijk} \text{id}_{\mathcal{E}_i}$$

One can show that this definition does not depend on the choice of 2-cocycle which is chosen to represent B . Like coherent sheaves on \mathcal{X} , twisted sheaves on \mathcal{X} form an abelian category, and so one can consider the derived category $D^b(\mathcal{X}, B)$ of B -twisted sheaves on \mathcal{X} .

In their paper [48], Kapustin and Orlov consider complex tori. These are specified by four pieces of data: (Λ, J, g, B) . Here, Λ is a lattice in a real vector space V so that $X = V/\Lambda$ yields a torus, J is a complex structure, g is a flat metric and $B \in H^2(X, \mathbb{R}/\mathbb{Z})$ is a class which lies in the kernel of the map $H^2(X, \mathbb{R}/\mathbb{Z}) \rightarrow H^3(X, \mathbb{Z})$.

From the above data, they explicitly construct the $N = 2$ SCVA of the complex torus. They then deduce a criterion under which two such $N = 2$ SCVA are isomorphic. If one supposes that two complex tori $(\Lambda_1, J_1, g_1, B_1)$ and $(\Lambda_2, J_2, g_2, B_2)$ are both mirror to a given complex torus $(\Lambda^\vee, I^\vee, g^\vee, B^\vee)$, then the $N = 2$ SCVAs of these two tori are isomorphic, by definition. Kapustin and Orlov prove that if two $N = 2$ SCVAs are isomorphic, then there exists an isomorphism of lattices from $\Lambda_1 \oplus \Lambda_1^*$ to $\Lambda_2 \oplus \Lambda_2^*$ which intertwines some complex structures on the complex tori $V/\Lambda_1 \times V^*/\Lambda_1^*$ and $V/\Lambda_2 \times V^*/\Lambda_2^*$. Homological mirror symmetry would imply that $D^b(\mathcal{X}_1) \cong D^b(\mathcal{X}_2)$, since \mathcal{X}_1 and \mathcal{X}_2 are both mirror to \mathcal{X}^\vee . There is a known condition under which such an equivalence of categories holds for two algebraic tori, and this condition is different from the one that Kapustin and Orlov deduce from an isomorphism of the $N = 2$ SCVAs. So generically, there will be no isomorphism between $D^b(\mathcal{X}_1)$ and $D^b(\mathcal{X}_2)$. But they prove the following.

Theorem 4.5.2 (Kapustin and Orlov [48]). *Let $(\Lambda_1, J_1, g_1, B_1)$ and $(\Lambda_2, J_2, g_2, B_2)$ define two algebraic tori such that their $N = 2$ SCVAs are isomorphic. Then*

$$D^b(\mathcal{X}_1, B_1) \cong D^b(\mathcal{X}_2, B_2)$$

Because of this theorem, they suggest to modify the homological mirror symmetry conjecture as follows, in the presence of the B -field (necessary only if $h^{2,0}(\mathcal{X}) \neq 0$).

Conjecture 4.5.3 (Homological Mirror Symmetry with B -field). *Let (\mathcal{X}, ω, B) and $(\mathcal{X}^\vee, \omega^\vee, B^\vee)$ be a mirror pair of Calabi-Yau manifolds. Then there exists a quasi-equivalence of A_∞ -categories*

$$D^b(\mathcal{X}, B) \cong D^b \text{Fuk}(X^\vee, \omega^\vee, B^\vee) \quad D^b(\mathcal{X}^\vee, B) \cong D^b \text{Fuk}(X, \omega, B)$$

Part II

The Thomas-Yau Conjecture

Chapter 5

Donaldson-Thomas Invariants

Donaldson-Thomas invariants are numbers that one can extract from a Calabi-Yau threefold. They are obtained by doing intersection theory on the moduli space of stable sheaves, of some given Calabi-Yau variety \mathcal{X} , depending on a choice of Kähler class. They have an interpretation in terms of string theory, as counting certain branes in type IIA string theory, called BPS branes. These BPS branes will be of great importance to us, as they are the central motivation for the Thomas-Yau conjecture. References for this chapter are [29, 14, 69, 70], as well as [8] for the final section.

The Donaldson-Thomas invariants themselves will, at first, not be directly related to the Thomas-Yau conjecture, but they appear in the table that Thomas gives at the start of his paper (which we have replicated in 8.2), which we aim to explain. However, in turn, the Thomas-Yau conjecture inspired further developments of Donaldson-Thomas theory which we will return to later and therefore we will start this second part of the text by discussing Donaldson-Thomas invariants, to be interpreted in the context of string theory later on.

5.1 Stability of Sheaves

Let \mathcal{E} be a coherent sheaf on \mathcal{X} , and let \mathcal{E}_x denote the stalk at $x \in \mathcal{X}$. We recall that $\text{Supp}(\mathcal{E}) = \{x \in \mathcal{X} \mid \mathcal{E}_x \neq 0\}$. When \mathcal{X} is coherent, this is a closed subset, and thus defines a subscheme of \mathcal{X} , which has a dimension.

Definition 5.1.1. The dimension of a coherent sheaf \mathcal{E} on \mathcal{X} is defined as $\dim \mathcal{E} = \dim \text{Supp}(\mathcal{E})$. A coherent sheaf \mathcal{E} is called pure if $\dim \mathcal{F} = \dim \mathcal{E}$ for all subsheaves $\mathcal{F} \subseteq \mathcal{E}$.

If we denote the irreducible components of the support by \mathcal{Y}_k , then these define a homology class

$$[\mathcal{E}] := \sum_k \text{Length}(\mathcal{E}_{\eta_k}) [\mathcal{Y}_k] \in H_{2 \dim \mathcal{E}}(X, \mathbb{Z})$$

where η_k is the unique generic point of the irreducible component \mathcal{Y}_k . We use the notation $[\mathcal{E}]$ to refer either to this class, or its Poincaré dual.

Every coherent sheaf has a determinant line bundle, and a Chern character, which generalise those of vector bundles. Define the determinant line bundle of \mathcal{E} by taking a locally free resolution $\mathcal{E}_\bullet \rightarrow \mathcal{E} \rightarrow 0$. Then we set

$$\det(\mathcal{E}) = \det(\mathcal{E}_\bullet) = \otimes_k \det(\mathcal{E}_k)^{(-1)^k} \in \text{Pic}(\mathcal{X})$$

5.1. STABILITY OF SHEAVES

The determinant line bundle of a locally free sheaf is just its top exterior power. Similarly, we define

$$\text{ch}(\mathcal{E}) := \oplus_k (-1)^k \text{ch}(\mathcal{E}_k) \in H^\bullet(X, \mathbb{Q})$$

Let $n = \dim \mathcal{X}$, $d = \dim \mathcal{E}$ and $c = n - d$. The following result tells us that the Chern character of a coherent sheaf encodes topological information about its support.

Proposition 5.1.2. *For a coherent sheaf \mathcal{E} on \mathcal{X} , we have $\text{ch}_k(\mathcal{E}) = 0$ if $k < c$, and $\text{ch}_k(\mathcal{E}) = [\mathcal{E}]$ if $k = c$.*

Recall that $\chi(\mathcal{X}, \mathcal{E}) = \sum_k (-1)^k \dim_{\mathbb{C}} H^k(\mathcal{X}, \mathcal{E})$. A version of the Grothendieck-Riemann-Roch theorem can that be stated as follows.

Theorem 5.1.3. *Let \mathcal{X} be a smooth projective variety and $\mathcal{E} \in \text{Coh}(\mathcal{X})$. Then*

$$\chi(\mathcal{X}, \mathcal{E}) = \int_{\mathcal{X}} \text{ch}(\mathcal{E}) td(\mathcal{X})$$

We will use this in a moment. First we consider the Hilbert polynomial. We fix an ample line bundle $\mathcal{O}_{\mathcal{X}}(1)$ on \mathcal{X} (obtained as the pullback of $\mathcal{O}(1)$ for some embedding $\mathcal{X} \rightarrow \mathbb{C}\mathbb{P}^N$). Denote by $\mathcal{E}(m) = \mathcal{E} \otimes \mathcal{O}_{\mathcal{X}}(m)$. The Kähler class of \mathcal{X} is $c_1(\mathcal{O}_{\mathcal{X}}(1))$ and will thus be denoted by ω .

Definition 5.1.4. The Hilbert polynomial $P(\mathcal{E})$ is defined by the function

$$m \mapsto \chi(\mathcal{X}, \mathcal{E} \otimes \mathcal{O}_{\mathcal{X}}(m))$$

One can show that this indeed defines a polynomial $P(\mathcal{E}, m) \in \mathbb{Q}[m]$, see [29].

Proposition 5.1.5. *The Hilbert polynomial $P(\mathcal{E})$ can be uniquely written as*

$$P(\mathcal{E}, m) = \sum_{k=0}^{\dim \mathcal{E}} \frac{\alpha_k(\mathcal{E})}{k!} m^k$$

for some integers $\alpha_k(\mathcal{E})$.

Some of these integers have a clear topological meaning to them. For example, $\alpha_0(\mathcal{E})$ is obtained by evaluating at $m = 0$, which yields $\chi(\mathcal{X}, \mathcal{E})$. More interesting is the coefficient of the leading term $\alpha_d(\mathcal{E})$, where $d = \dim \mathcal{E}$. By the Grothendieck-Riemann-Roch theorem and the fact that $\text{ch}(\mathcal{E} \otimes \mathcal{O}_{\mathcal{X}}(m)) = \text{ch}(\mathcal{E}) \text{ch}(\mathcal{O}_{\mathcal{X}}(1))^m$, we find that

$$\alpha_d(\mathcal{E}) = \langle [\mathcal{E}], \omega^d \rangle$$

For instance, taking $\mathcal{E} = \mathcal{O}_{\mathcal{X}}$ yields

$$\alpha_n(\mathcal{O}_{\mathcal{X}}) = \int_{\mathcal{X}} \omega^n = \text{deg}(\mathcal{X}) \quad \alpha_{n-1}(\mathcal{O}_{\mathcal{X}}) = -\frac{1}{2} \langle K_{\mathcal{X}}, \omega^{n-1} \rangle$$

5.2. THE MODULI SPACE

where $K_{\mathcal{X}}$ is the homology class of the canonical divisor of \mathcal{X} . The rank of a sheaf is defined as $\text{rank}(\mathcal{E}) = \alpha_n(\mathcal{E})/\alpha_n(\mathcal{O}_{\mathcal{X}})$. The degree of \mathcal{E} is defined as

$$\text{deg}(\mathcal{E}) = \alpha_{n-1}(\mathcal{E}) - \text{rank}(\mathcal{E})\alpha_{n-1}(\mathcal{O}_{\mathcal{X}})$$

All of this is to say that the Hilbert polynomial encodes certain discrete invariants of the sheaf. This information is generally not equivalent to that of the Chern character. One can determine the Hilbert polynomial from the Chern character, but not vice versa.

We define $p(\mathcal{E}, m) = P(\mathcal{E}, m)/\alpha_d(\mathcal{E})$ and recall that there is a natural ordering on the set $\mathbb{Q}[m]$. We say $f \leq g$ if and only if $f(m) \leq g(m)$ for m sufficiently large, and $f < g$ if and only if $f(m) < g(m)$ for m sufficiently large.

Definition 5.1.6. Let \mathcal{E} be a coherent sheaf on \mathcal{X} . Then \mathcal{E} is called semi-stable if it is pure, and for all proper subsheaves $\mathcal{F} \subset \mathcal{E}$, we have $p(\mathcal{F}) \leq p(\mathcal{E})$. If the inequality is strict, then \mathcal{E} is called stable.

Now we know what it means for a sheaf to be (semi-)stable. Next we want to understand what it means to have a moduli space of stable sheaves.

5.2 The Moduli Space

To understand what a moduli space is, we first want to establish what it means to "deform" a sheaf. In essence, this can be interpreted as varying the holomorphic structure on some smooth vector bundle. But one has to be a lot more careful, particularly if one wants a compact moduli space. Let $\mathcal{S} = \mathcal{X} \times \mathcal{B}$ for some fixed scheme \mathcal{X} (which will be a Calabi-Yau threefold for our purposes), while the scheme \mathcal{B} may vary.

Definition 5.2.1. Suppose that \mathcal{E} is a coherent $\mathcal{O}_{\mathcal{S}}$ -module. Then the collection $\{\mathcal{E}|_{\pi^{-1}(b)} := \mathcal{E}_b\}$ where b is a closed point of \mathcal{B} is called a flat family of coherent sheaves on \mathcal{X} if \mathcal{E} is flat over \mathcal{B} .

This is the "correct" definition to ensure that the family of sheaves \mathcal{E}_b vary along the base space \mathcal{B} in a nice way. The definition for a module to be flat over a ring is not illuminating. It is best to phrase this condition in some equivalent way which makes its use more apparent. We fix a line bundle $\mathcal{O}_{\mathcal{S}}(1)$ on \mathcal{S} , whose restriction to each fibre of π is ample. This is essentially specifying an embedding $\mathcal{S} = \mathcal{X} \times \mathcal{B} \rightarrow \mathbb{C}\mathbb{P}^N \times \mathcal{B}$.

Theorem 5.2.2. [29] *Suppose that \mathcal{B} is reduced. Then the following are equivalent:*

1. \mathcal{E} is flat over \mathcal{B} .
2. The Hilbert polynomial $P(\mathcal{E}_b)$ is a locally constant function on \mathcal{B} .

If \mathcal{B} is not reduced, then $1 \implies 2$.

The fact that the Hilbert polynomial becomes a locally constant function for flat families gives us a good reason to believe that this is indeed the right definition, in the sense

5.2. THE MODULI SPACE

that the sheaves vary "continuously" along the base \mathcal{B} in some appropriate sense, because the discrete invariants do not jump. We can think of a flat family of sheaves on \mathcal{X} as a deformation of some given sheaf, say \mathcal{E}_{b_0} for some base point $b_0 \in \mathcal{B}$. For example, if $\mathcal{E}|_{b_0}$ is (the sheaf of sections of) a vector bundle on \mathcal{X} , then such a family might correspond to different holomorphic structures on the underlying smooth vector bundle E . The following result tells us that (excluding the case of the empty set) this is a reasonably good intuition to keep in mind, as we are really considering families of vector bundles, which "degenerate" along closed subsets on \mathcal{B} to sheaves which are less well-behaved (we remind the reader that Zariski open subsets in irreducible varieties are dense).

Proposition 5.2.3. [29] *Let \mathcal{E} be a flat family of coherent sheaves over \mathcal{B} . Then the set $\{b \in \mathcal{B} \mid \mathcal{E}|_{\mathcal{X}_b} \text{ is locally free}\}$ is open in \mathcal{B} .*

The moduli space of sheaves is interpreted as a scheme whose points parameterise isomorphism classes of sheaves with fixed topological data, such as the Chern character. Inside of this space sits an open subset which parameterises the holomorphic structures on a given vector bundle with the corresponding Chern character (again, up to isomorphism). Having established the goal, let us outline how this is done more formally. Recall the following definition, in which the functor $h_{\mathcal{M}} : \text{Sch}_{\mathbb{C}} \rightarrow \text{Set}$ is defined by $S \mapsto \text{Mor}(S, \mathcal{M})$.

Definition 5.2.4. Let $M : \text{Sch}_{\mathbb{C}} \rightarrow \text{Set}$ be a contravariant functor. Then a scheme \mathcal{M} is said to represent M if there is an isomorphism of functors $M \cong h_{\mathcal{M}}$.

To define the moduli space of sheaves, we wish to define a suitable functor. If this functor can be represented by a scheme \mathcal{M} , then we call \mathcal{M} the fine moduli space of sheaves (with some fixed discrete invariants). So what should the functor be? First, we fix the data of a smooth, irreducible and projective variety \mathcal{X} , together with an embedding $\mathcal{X} \rightarrow \mathbb{C}\mathbb{P}^N$ and an element $\eta \in H^{\text{ev}}(X, \mathbb{Q})$ such that $\eta_0 > 0$, which will be the Chern character of the sheaves we want to parameterise. The condition $\eta_0 > 0$ rules out the case of torsion sheaves. We define a functor $M_{\eta}^{\text{ss}}(\mathcal{X}) : \text{Sch}_{\mathbb{C}} \rightarrow \text{Set}$, which sends \mathcal{B} to the set of semi-stable coherent sheaves on \mathcal{X} parameterised by \mathcal{B} whose Chern character is η , modulo S -equivalence (see B.2.2 for a definition).

This functor is typically not representable, in which case a fine moduli space does not exist. We look for a slightly weaker notion, which is that of a coarse moduli space.

Definition 5.2.5. A coarse moduli space for a moduli functor M is a scheme \mathcal{M} together with a natural transformation $M \rightarrow h_{\mathcal{M}}$, such that

1. For every natural transformation $M \rightarrow h_S$, there exists a unique natural transformation $h_{\mathcal{M}} \rightarrow h_S$ which makes the obvious diagram commute.
2. There is a bijection $M(\text{Spec}(\mathbb{C})) \rightarrow h_{\mathcal{M}}(\text{Spec}(\mathbb{C}))$.

We can describe the difference between a coarse moduli space and a fine moduli space in more geometric terms. If \mathcal{M} is a fine moduli space, then every family parame-

5.3. DONALDSON-THOMAS INVARIANTS

terised by \mathcal{B} gives rise to a unique morphism $\mathcal{B} \rightarrow \mathcal{M}$, which is the definition of $h_{\mathcal{M}} \cong M$. Consider $h_{\mathcal{M}}(\mathcal{M}) = \text{Mor}(\mathcal{M}, \mathcal{M})$, which contains the identity map, and hence defines a tautological family $\mathcal{F} \rightarrow \mathcal{M}$. The family parameterised by \mathcal{B} is then obtained as the pull-back of the tautological family by the morphism $\mathcal{B} \rightarrow \mathcal{M}$.

A coarse moduli space, on the other hand, also determines a morphism $\mathcal{B} \rightarrow \mathcal{M}$, but this time, there may not be a universal family. However, the second condition in the definition still says that two objects correspond to the same point in the moduli space if and only if they are isomorphic. This is a more reasonable criterion to ask for. And indeed, in [29], the following is proved.

Theorem 5.2.6. *The functor $M_{\eta}^{ss}(\mathcal{X})$ has a coarse moduli space $\mathcal{M}_{\eta}^{ss}(\mathcal{X})$ which is a projective \mathbb{C} -scheme. There is an open subset $\mathcal{M}_{\eta}^s(\mathcal{X})$ which is a coarse moduli space of stable sheaves.*

There is a reason why we restrict to semi-stable sheaves: this is required to use the machinery of geometric invariant theory, which is how the moduli space $\mathcal{M}_{\eta}^{ss}(\mathcal{X})$ is constructed (see [29]). But that is not a satisfactory reason. A better reason is as follows: we want the moduli space to be a scheme of finite type, but if we consider families of unstable sheaves, then we get unbounded families of sheaves, such as $\{\mathcal{O}(-n) \oplus \mathcal{O}(n)\}$ over $\mathbb{C}\mathbb{P}^1$ for all $n \geq 0$.

Another reason is perhaps more conceptual, and is related to the automorphisms of the objects we are parameterising.

Proposition 5.2.7. *Let \mathcal{E} be a stable sheaf on \mathcal{X} . Then $\text{End}(\mathcal{E}) \cong \mathbb{C}$.*

Since the identity morphism is always an automorphism, this means that stable sheaves have the minimal automorphism size that we could ask for, i.e. \mathbb{C}^{\times} . Non-trivial automorphisms of the objects being parameterised are problematic because they give rise to non-trivial families whose fibres are all isomorphic. Such a family cannot correspond to a constant map $\mathcal{B} \rightarrow \mathcal{M}$, because this map already corresponds to the trivial family. This immediately rules out the existence of a fine moduli space. Because the automorphisms of stable sheaves are as small as they can be, this can be accounted for, so that we can still hope for a coarse moduli space.

5.3 Donaldson-Thomas Invariants

The Donaldson-Thomas invariants, originally defined in [14], are obtained from the moduli space $\mathcal{M}_{\eta}^s(\mathcal{X})$, using a virtual fundamental class. First, we need to establish some preliminary results.

Theorem 5.3.1. *Let $\mathcal{E} \in \mathcal{M}_{\eta}^s(\mathcal{X})$. Then $T_{\mathcal{E}}\mathcal{M}_{\eta}^s(\mathcal{X}) \cong \text{Ext}^1(\mathcal{E}, \mathcal{E})$.*

At the level of holomorphic vector bundles, this is not a mystery. If \mathcal{E} is a holomorphic vector bundle, then $\text{Ext}^1(\mathcal{E}, \mathcal{E}) = H^1(\mathcal{X}, \text{End}(\mathcal{E}))$. Its elements may be viewed as

5.3. DONALDSON-THOMAS INVARIANTS

endomorphism-valued $(0,1)$ -forms deforming the $\bar{\partial}$ -operator. One has to take cohomology because $\bar{\partial}$ -exact terms are the result of complex gauge transformations. These are infinitesimal deformations of the holomorphic structure, because there is a bijection between holomorphic structures on E , and $\bar{\partial}$ -operators squaring to 0 up to gauge equivalence. We elaborate on this in the next section.

Tangent vectors in the moduli space correspond to infinitesimal deformations of the sheaf. Not all infinitesimal deformations can be integrated to give a deformation, there are obstruction classes. That is, every tangent vector yields an element in $\text{Ext}^k(\mathcal{E}, \mathcal{E})$ for $k \geq 2$. These spaces are called the obstruction spaces, and the elements in them determined by a tangent vector are called the obstruction classes. The tangent vector can be integrated if and only if each obstruction class vanishes.¹

We now move on to state the main result given in [14], which leads to the definition of the Donaldson-Thomas invariants.

Theorem 5.3.2 (Thomas). *Let \mathcal{X} be a smooth projective algebraic threefold, and let $\mathcal{M}_\eta^s(\mathcal{X})$ be a moduli space of sheaves on \mathcal{X} . Suppose the dimension of $\text{Ext}^3(\mathcal{E}, \mathcal{E})$ is constant for $\mathcal{E} \in \mathcal{M}_\eta^s(\mathcal{X})$. Then the groups $\text{Ext}^1(\mathcal{E}, \mathcal{E})$ and $\text{Ext}^2(\mathcal{E}, \mathcal{E})$ govern a perfect obstruction theory of virtual dimension*

$$vdim \mathcal{M}_\eta^s(\mathcal{X}) = \dim \text{Ext}^1(\mathcal{E}, \mathcal{E}) - \dim \text{Ext}^2(\mathcal{E}, \mathcal{E})$$

Let us address the notion of a virtual dimension. Intuitively, one can think about this as follows. Suppose that $\mathcal{E} \rightarrow \mathcal{X}$ is a holomorphic vector bundle (where \mathcal{X} is some complex manifold, not necessarily an algebraic threefold), and $\mathcal{Y} \subseteq \mathcal{X}$ is given as $Z(s)$ for some $s \in H^0(\mathcal{X}, \mathcal{E})$. For a generic section which intersects the zero section transversally, we would get $\dim \mathcal{Y} = \dim \mathcal{X} - \text{rank } \mathcal{E}$. But the intersection does not have to be transversal. Using the machinery of virtual cycles, however, there is a way to construct a homology class $[Y]^{\text{vir}} \in H_{2vdim Y}(X, \mathbb{Z})$ which behaves as if the section s were transversal. This homology class is called the virtual fundamental class. So what are the conditions under which a moduli space admits a virtual fundamental class? An answer is given by [71].

Theorem 5.3.3. *Suppose \mathcal{M} is a moduli space with a perfect obstruction theory. Then it admits a virtual fundamental class $[\mathcal{M}]^{\text{vir}} \in H_{2vdim \mathcal{M}}(\mathcal{M}, \mathbb{Z})$.*

So this is what Thomas's thesis established: that the moduli spaces of sheaves on a smooth projective algebraic threefold carry a perfect obstruction theory, and hence allow for the definition of a virtual fundamental class. We will not comment on the notion of a perfect obstruction theory, as it is a highly technical one. We refer to [71]. Calabi-Yau threefolds are of particular interest, for the following reason.

¹Intuitively, we can picture this as follows: if all the obstruction classes of the tangent vectors at a point vanish, then the deformations are unobstructed, which corresponds to a smooth point in the moduli space. If the deformations are obstructed, this corresponds to a singular point.

5.3. DONALDSON-THOMAS INVARIANTS

Corollary 5.3.4. *Let \mathcal{X} be a Calabi-Yau threefold. Then $\text{vdim } \mathcal{M}_\eta^s(\mathcal{X}) = 0$.*

Proof. We see that $\text{Ext}^0(\mathcal{E}, \mathcal{E}) = \text{Hom}(\mathcal{E}, \mathcal{E}) = \text{End}(\mathcal{E}) = \mathbb{C}$. By Serre duality

$$\text{Ext}^q(\mathcal{E}, \mathcal{E}) \cong \text{Ext}^{3-q}(\mathcal{E}, \mathcal{E} \otimes K_X)^* \cong \text{Ext}^{3-q}(\mathcal{E}, \mathcal{E})^*$$

so the dimension of $\text{Ext}^3(\mathcal{E}, \mathcal{E})$ is constant when \mathcal{X} is a Calabi-Yau threefold. Furthermore, by the same argument, $\text{Ext}^1(\mathcal{E}, \mathcal{E}) \cong \text{Ext}^2(\mathcal{E}, \mathcal{E})^*$ which implies that the virtual dimension of $\mathcal{M}_\eta^s(\mathcal{X})$ is zero. \square

This finally brings us to the definition of the Donaldson-Thomas invariants. We assume that semi-stability implies stability. This happens for instance whenever the rank and degree are coprime. In this case, the moduli space is compact, so we can integrate over it.

Definition 5.3.5. Let \mathcal{X} be a Calabi-Yau threefold, and $\eta \in H^{\text{ev}}(X, \mathbb{Q})$ with $\eta_0 > 0$. Let \mathcal{M} be the corresponding moduli space of stable sheaves. Then the Donaldson-Thomas invariant $\text{DT}_\eta(\mathcal{X})$ is defined as

$$\text{DT}_\eta(\mathcal{X}) := \langle [M]^{\text{vir}}, 1 \rangle = \int_{M^{\text{vir}}} 1$$

In much of the literature, however, Donaldson-Thomas invariants are something more specific. Namely, one looks at the moduli space of rank 1 torsion free sheaves satisfying $\det(\mathcal{L}) \cong \mathcal{O}_X$. These are ideal sheaves for curves $C \subset X$. By fixing

$$\eta = (1, 0, -\beta, -n)$$

one fixes a (co)homology class $[C] = \beta$ and $\chi(\mathcal{O}_C) = n$. The moduli space of these objects sits inside of the one whose construction we outlined above. It is a projective scheme with a symmetric obstruction theory, so that we again get a virtual fundamental class of dimension 0. The invariants obtained by integrating 1 over these virtual fundamental cycles are also called the Donaldson-Thomas invariants $\text{DT}_{\beta, n}(\mathcal{X})$ in the literature. They are assembled into a generating function

$$Z_\beta^{\text{DT}}(q) = \sum_n \text{DT}_{\beta, n}(\mathcal{X}) q^n$$

This is typically normalised by dividing by zero dimensional subschemes, i.e. setting

$$Z_\beta^{\text{red}}(q) = Z_\beta^{\text{DT}}(q) / Z_0^{\text{DT}}(q)$$

In this way, Donaldson-Thomas theory becomes a theory of counting curves on a Calabi-Yau threefold. Furthermore, it has been shown in [72] that Donaldson-Thomas invariants for higher ranks may be deduced from this, at least on certain Calabi-Yau manifolds such as the quintic threefold, so it appears that interpreting Donaldson-Thomas theory as a curve counting theory is really the right thing to do.

5.4. GAUGE THEORY AND HOLOMORPHIC BUNDLES

We encountered another curve counting theory, namely Gromov-Witten theory, where one instead considers the moduli space of (stable) maps from a curve into the Calabi-Yau manifold. A deep conjecture, known as the MNOP (Maulik-Nekrasov-Okounkov-Pandharipande) conjecture asserts that these two theories actually coincide. That is, the generating function $Z_\beta^{\text{red}}(q)$ should coincide with the corresponding generating function for the Gromov-Witten invariants, after making a change of variables. There is an extensive list of examples for which the MNOP conjecture has been shown to hold. We refer the reader to [70] for more details. By [72], Gromov-Witten theory also determines the higher rank Donaldson-Thomas invariants, which is quite surprising. After all, the curves which appear in Gromov-Witten theory are interpreted as closed strings in the target space \mathcal{X} , where the Donaldson-Thomas invariants count certain "stable" boundary conditions for open strings.

5.4 Gauge Theory and Holomorphic Bundles

We can also view the story outlined above from a gauge theoretic perspective. We will need some results from [8] for this, as well as for the next chapters. Let (\mathcal{E}, h) be a Hermitian holomorphic vector bundle on a complex manifold \mathcal{X} . A connection ∇ on E splits into

$$\nabla = \nabla^{1,0} \oplus \nabla^{0,1} : \Gamma(X, E) \rightarrow \Omega^{1,0}(X, E) \oplus \Omega^{0,1}(X, E)$$

Definition 5.4.1. A connection ∇ on a holomorphic bundle $\mathcal{E} \rightarrow \mathcal{X}$ is called compatible with the holomorphic structure if $\bar{\partial}_{\mathcal{E}} = \nabla^{0,1}$. A connection ∇ on a Hermitian bundle (E, h) is called unitary if for all $s_1, s_2 \in \Gamma(X, E)$,

$$d\langle s_1, s_2 \rangle = \langle \nabla s_1, s_2 \rangle + \langle s_1, \nabla s_2 \rangle$$

Given a Hermitian metric on a holomorphic bundle, there is a unique connection which is both compatible with the metric and the holomorphic structure. We call this connection the Chern connection. It can be defined locally by choosing an orthonormal frame, and representing h as a matrix of functions. Then the local connection 1-form is defined as $A = h^{-1}\partial h$. This is a $(1, 0)$ -form and so the $(0, 1)$ component of $\nabla = d + A$ is just $\bar{\partial}$, as desired. It is also clear that the curvature form F_∇ is of type $(1, 1)$. We will use the following two results later on in the text.

Theorem 5.4.2 ([8]). *Let (\mathcal{E}, h) be a Hermitian holomorphic vector bundle on \mathcal{X} . Then a unitary connection on E is the Chern connection of (\mathcal{E}, h) if and only if F_∇ is of type $(1, 1)$.*

Theorem 5.4.3 ([8]). *Let (E, h) be a Hermitian vector bundle on a complex manifold \mathcal{X} . There is a bijective correspondence between holomorphic structures on E , and unitary connections ∇ on E which satisfy $(\nabla^{0,1})^2 = 0$. Given such a connection, it is the Chern connection of the holomorphic structure on (E, h) .*

5.4. GAUGE THEORY AND HOLOMORPHIC BUNDLES

Because of this, we get a bijective correspondence between holomorphic structures on E , and unitary connections with curvature of type $(1, 1)$.

Definition 5.4.4. A partial connection is a linear operator $\bar{\partial}: \Gamma(X, E) \rightarrow \Omega^{0,1}(X, E)$ which satisfies the Leibniz rule.

Similar to connections, partial connections form an infinite dimensional affine space modelled on $\Omega^{0,1}(X, \text{End}(E))$. Note that the unitary connection is determined by its partial connection. For the local one forms, $A = \alpha - \alpha^\dagger$, where \dagger is the Hermitian conjugate.

Theorem 5.4.5 ([8]). *A partial connection on a complex vector bundle over a complex manifold is the Dolbeault operator of a holomorphic structure if and only if $\bar{\partial}^2 = 0$.*

Isomorphic holomorphic structures on E correspond to (partial) connections which are related by a (unitary) gauge transformation. Because of this, we can interpret the moduli space of holomorphic structures on a given vector bundle E in two ways. Either, we consider the space of "flat" partial unitary connections and quotient by unitary gauge transformations, or we consider the space of "flat" partial connections and quotient by complex gauge transformations, instead of just unitary ones. The latter brings us to the motivation for Donaldson-Thomas invariants.

Thomas called these invariants "holomorphic Casson invariants", for the following reason. On a simply connected smooth 3-manifold X , one can consider the Chern-Simons action functional on the space of connections \mathcal{A} of a given vector bundle $E \rightarrow X$. Fixing some connection ∇_0 as the origin of the affine space, we may identify $\mathcal{A} \cong \Omega^1(X, \text{End}(E))$. The Chern-Simons action functional is defined as

$$S(A) := \int_X \text{tr}(A \wedge d_{A_0} A + \frac{2}{3} A^3)$$

The functional itself is not gauge invariant, but its differential is. Indeed, it is readily seen to be

$$dS_A: T_A \mathcal{A} \rightarrow \mathbb{R} \quad a \mapsto \int_X \text{tr}(a \wedge F_A)$$

which is gauge invariant. The zeroes of the differential give flat connections in the moduli space of connections $\mathcal{M} = \mathcal{A}/\mathcal{G}$. Counting these zeroes of the differential in \mathcal{M} gives the Casson invariant. In their paper [73], Donaldson and Thomas discuss how Calabi-Yau manifolds are the natural complex analogue of oriented manifolds, since they have a trivial canonical bundle. This motivates them to search for an analogue of the Casson invariant for Calabi-Yau threefolds. Given a simply connected Calabi-Yau threefold, we can take a holomorphic volume form $\Omega \in \Omega^{3,0}(X)$. We want to pair this with some $(0, 3)$ -form so that we can integrate over X . If we are looking for an analogue of the Casson invariant, it makes sense to consider vector bundles $E \rightarrow X$ and define $\mathcal{A}^{0,1} = \Omega^{0,1}(X, \text{End}(E))$, which we can identify with the space of $\bar{\partial}$ -operators on E , after fixing some basepoint $\bar{\partial}_0$. One then defines a functional on $\mathcal{A}^{0,1}$ by

$$S_{\mathbb{C}}(A) = \int_X \text{tr}(A \wedge \bar{\partial}_0 A + \frac{2}{3} A^3) \wedge \Omega$$

5.4. GAUGE THEORY AND HOLOMORPHIC BUNDLES

Once again, the differential descends to $\mathcal{M}^{0,1} = \mathcal{A}^{0,1}/\mathcal{G}_{\mathbb{C}}$, but this time, the differential is given by

$$dS_{\mathbb{C}} : T_A \mathcal{A}^{0,1} \rightarrow \mathbb{C} \quad a \mapsto \int_X \text{tr}(a \wedge F_A^{0,2}) \wedge \Omega$$

Thus, its zeroes correspond to $\bar{\partial}$ -operators which square to 0 - these are in bijective correspondence with holomorphic structures on E . This gives a solid foundation to believe that the moduli space $\mathcal{M}(X, E)$ of holomorphic structures on E has (virtual) dimension 0, since it corresponds to the critical points of some functional.

Donaldson had developed the analytic techniques to do intersection theory on the moduli space of connections using the Yang-Mills functional on 4-manifolds (which is referred to as Donaldson theory, see [8]), instead of the holomorphic Chern-Simons functional on Calabi-Yau threefolds. These analytic techniques could not be carried over to the world of Calabi-Yau threefolds in any straightforward manner, so instead, Thomas worked within the more rigid framework of algebraic geometry to formulate his result, which culminated in Donaldson-Thomas invariants.

Chapter 6

The Kobayashi-Hitchin Correspondence

The Kobayashi-Hitchin correspondence (also known as the Donaldson-Uhlenbeck-Yau theorem) will relate stability in algebraic geometry, to the existence of a solution to a certain partial differential equation. The result holds for any compact Kähler manifold, and so whenever $\mathcal{Y} \subset \mathcal{X}$ is a compact complex submanifold of a Calabi-Yau manifold, the discussion below applies to \mathcal{Y} . This will be quite important for us.

6.1 μ -Stability

We recall that we discussed the notion of stable coherent sheaves in 5.1. There, we presented one definition of stability. There are other notions of stability, such as μ -stability (also known as slope stability) which we now define. Recall that the rank of a coherent sheaf can be defined as $\text{rank}(\mathcal{E}) = \alpha_n(\mathcal{E})/\alpha_n(\mathcal{O}_{\mathcal{X}})$, where α_k are the coefficients of the Hilbert polynomial for the respective sheaves.

Proposition 6.1.1. *Suppose that \mathcal{E} is the sheaf of sections of a holomorphic vector bundle of rank r on a Kähler manifold (\mathcal{X}, ω) . Then $\text{rank}(\mathcal{E}) = r$, where the left hand side is the rank as defined by the Hilbert polynomial.*

Proof. The Grothendieck-Riemann-Roch theorem yields

$$\chi(\mathcal{X}, \mathcal{E}(m)) = \frac{r m^n}{n!} \text{deg}(\mathcal{X}) + \frac{m^{n-1}}{(n-1)!} \int_{\mathcal{X}} \omega^{n-1} \wedge c_1(\mathcal{E}) + \text{Lower degree terms}$$

It follows from the definition, and the fact that $\alpha_n(\mathcal{O}_{\mathcal{X}}) = \text{deg}(\mathcal{X})$, that $\text{rank}(\mathcal{E}) = r$. \square

Consequently, we also find $\text{deg}(\mathcal{E}) = \int_{\mathcal{X}} \omega^{n-1} \wedge c_1(\mathcal{E})$, so we take this to be the definition of the degree of a vector bundle, and we have the notion of μ -stability for vector bundles. Define $\mu(\mathcal{E}) = \text{deg}(\mathcal{E})/\text{rank}(\mathcal{E})$.

Definition 6.1.2. Let \mathcal{E} be a vector bundle over \mathcal{X} . Then \mathcal{E} is called μ -semi-stable if, for all proper sub-sheaves $\mathcal{F} \subset \mathcal{E}$, it holds that

$$\mu(\mathcal{F}) \leq \mu(\mathcal{E})$$

If the inequality is strict, then \mathcal{E} is called μ -stable.

6.2. THE HERMITIAN-YANG-MILLS EQUATION

One can verify that every μ -stable vector bundle is also stable in the sense of 5.1, so every μ -stable vector bundle corresponds to a point in the appropriate moduli space of stable sheaves. It is these μ -stable vector bundles which will be related to some partial differential equation, known as the Hermitian-Yang-Mills equation.

6.2 The Hermitian-Yang-Mills Equation

The Hermitian-Yang-Mills (HYM) equation is a partial differential equation which involves the curvature of a connection on a Hermitian vector bundle over a Kähler manifold. In essence, its solutions correspond to connections with the most "convenient" curvature properties. For instance, we would prefer to have flat connections. But flat connections on a vector bundle do not always exist, for example if $c_1(E) \neq 0$. Define $F_A \cdot \omega \in \Gamma(X, \text{End}(E))$ by $F_A \wedge \omega^{n-1} = (F_A \cdot \omega)\omega^n$. This corresponds to orthogonal projection in $\Omega^2(X, \text{End}(E))$ onto the span of $\omega \otimes \text{id}_E$. The inner product on $\Omega^2(X, \text{End}(E))$ is induced by the Kähler metric on \mathcal{X} and the Hermitian metric on $\text{End}(E)$.

Definition 6.2.1. Let A be a connection on a Hermitian vector bundle E over a Kähler manifold \mathcal{X} . Then we say that A has constant central curvature if

$$F_A \cdot \omega = \lambda(E)\text{id}_E$$

for some constant¹ $\lambda(E)$.

If $\lambda(E) = 0$, i.e. when the vector bundle has degree 0, then a connection has constant central curvature if and only if it is flat. However, when E does not admit a flat connection, constant central curvature connections are those which are projectively flat, which is the next best thing that one could hope for.

Definition 6.2.2. Let $E \rightarrow X$ be a smooth vector bundle with Hermitian metric over a compact Kähler manifold \mathcal{X} . Let A be a unitary connection on E with curvature F_A . Then A satisfies the HYM equation if

$$\begin{cases} F_A \cdot \omega = \lambda(E)\text{id}_E \\ F_A^{0,2} = 0 \end{cases}$$

The first condition says that we are "minimising" the curvature, and the second condition says that the $\bar{\partial}$ -operator determined by the $(0,1)$ -part of the covariant derivative of A

$$\nabla^{0,1} : \Gamma(X, E) \rightarrow \Omega^{0,1}(X, E)$$

is the Dolbeault operator of some holomorphic structure \mathcal{E} on E . The name is derived from Yang-Mills theory, combined with the theory of complex Hermitian geometry. We

¹This constant is $\frac{\pi i}{3\text{Vol}(X)}\mu(E)$, for a threefold.

6.3. HERMITIAN-YANG-MILLS AS SYMPLECTIC QUOTIENT

refer to the paper by Donaldson [74] as well as the book [8] by Donaldson and Kronheimer for its relation to Yang-Mills theory, in which they relate stability of bundles to Yang-Mills connections on complex surfaces. Presently, however, we are interested in complex threefolds, for which the following was proved by Uhlenbeck and Yau in [7].

Theorem 6.2.3 (Kobayashi-Hitchin Correspondence). *Let \mathcal{E} be a Hermitian holomorphic vector bundle over a compact Kähler manifold. Then \mathcal{E} is μ -(poly)stable² if and only if \mathcal{E} admits a HYM connection. This is the unique HYM connection in the orbit of the Chern connection.*

Remark 6.2.4. Recall that the Chern connection ∇ on a Hermitian holomorphic vector bundle $\mathcal{E} \rightarrow \mathcal{X}$ is defined uniquely by the following properties:

1. $\nabla^{0,1} = \bar{\partial}_{\mathcal{E}}$
2. $d\langle s_1, s_2 \rangle = \langle \nabla s_1, s_2 \rangle + \langle s_1, \nabla s_2 \rangle$

This beautiful theorem relates something which can be defined algebraically, to something which requires the analysis of infinite dimensional vector spaces, partial differential equations and Hermitian metrics. This is quite a remarkable result, which gives us yet another insight into why complex geometry is a fascinating playground in which to combine algebraic and analytic methods. Another example of such interplay is, of course, mirror symmetry. Is the Kobayashi-Hitchin correspondence related to mirror symmetry in any tangible way? The Thomas-Yau conjecture asserts that the answer to this question is affirmative, as we will explain when discussing the conjecture.

6.3 Hermitian-Yang-Mills as Symplectic Quotient

The Kobayashi-Hitchin correspondence can also be viewed as a version of the Kempf-Ness theorem, but now in infinite dimensions. We first recall the Kempf-Ness theorem in finite dimensions, which states the following. Suppose that $\mathcal{X} \subseteq \mathbb{C}\mathbb{P}^N$ is a complex manifold, with a complex reductive Lie group G acting on \mathcal{X} holomorphically. Suppose that K is the maximal compact subgroup (e.g. $SU(n) \subset SL(n, \mathbb{C})$, which is not itself a complex Lie group, only a smooth Lie group) of G acting on the symplectic manifold (X, ω) by symplectomorphisms, as well as preserving the complex structure. Suppose this action admits a momentum map $\mu : X \rightarrow \mathfrak{g}^*$.³ Let $\xi \in \mathfrak{g}^*$ be a central element. Then we can consider two quotients: the symplectic reduction $X // K := \mu^{-1}(\xi)/K$, or the GIT quotient \mathcal{X}/G . The Kempf-Ness theorem asserts that $X // K \cong \mathcal{X}/G$ is an isomorphism.

We explain an important feature of the theorem above, for which we first need to recall some geometric invariant theory. The GIT quotient \mathcal{X}/G requires a linearisation of

²Polystable means that the vector bundle is a direct sum of stable vector bundles of the same rank. We can exclude these by requiring the connection to be irreducible.

³A momentum map for the group action is a map such that $\iota_{\rho(\xi)}\omega = \langle \mu, \xi \rangle$ for all $\xi \in \mathfrak{g}$, where $\rho(\xi)_x := \frac{d}{dt}|_{t=0} \exp(t\xi) \cdot x$ is the fundamental vector field of ξ .

6.3. HERMITIAN-YANG-MILLS AS SYMPLECTIC QUOTIENT

the action, to an action of G on $\mathcal{O}_{\mathcal{X}}(-1)$. The latter is the restriction of the tautological line bundle $\mathbb{C}^{N+1} \supset \mathcal{O}(-1) \rightarrow \mathbb{C}\mathbb{P}^N$. Choose a lift \tilde{x} of a point $x \in X$ to the fibre $\pi^{-1}(x)$. Then the point x is said to be semi-stable if and only if the closure of $G \cdot \tilde{x}$ does not contain the origin. There is a surjective morphism $\mathcal{X}^{ss} \rightarrow \mathcal{X}/G$ from the semi-stable points to the GIT quotient. Thus, in some sense, the GIT quotient only sees those points which are semi-stable. On the other hand, we have the symplectic quotient $X // K$. The Kempf-Ness theorem then implies that the momentum map allows us to select distinguished representatives of the (semi-)stable points. In particular, the momentum map informs us about the "right" notion of stability. This perspective is important for the Thomas-Yau conjecture, but more immediately, for what we are about to demonstrate.

We now move on to the infinite dimensional case. Let $E \rightarrow X$ be a Hermitian vector bundle on a complex manifold \mathcal{X} , and consider the space of unitary connections \mathcal{A} , which we identify with $\Omega^1(X, \mathfrak{u}(r))$, the 1-forms with values in the adjoint bundle of the $U(r)$ -bundle over X . We will consider Kähler manifolds (\mathcal{X}, ω) .

Proposition 6.3.1. *Define $\vartheta : T_A \mathcal{A} \times T_A \mathcal{A} \rightarrow \mathbb{R}$ by*

$$(a, b) \mapsto \int_X \text{tr}(a \wedge b) \wedge \omega^{n-1}$$

Then ϑ is a symplectic form on the infinite dimensional manifold \mathcal{A} .

For the proof, we use some Hodge theory. We recall that the Hermitian metric on E together with the Riemannian metric on X give us an inner product $\langle \cdot, \cdot \rangle$ on $\Omega^k(X, \mathfrak{u}(r))$, defined using the Hodge star operator $\star : \Omega^k(X, \mathfrak{u}(r)) \rightarrow \Omega^{n-k}(X, \mathfrak{u}(r))$ by

$$(a, b) \mapsto \int_X \text{tr}(a \wedge \star b)$$

Here, the trace map is induced by the Hermitian metric, and we have

$$\text{tr} : \Omega^k(X, \mathfrak{u}(r)) \times \Omega^l(X, \mathfrak{u}(r)) \rightarrow \Omega^{k+l}(X)$$

which is non-degenerate. The operator $L : \Omega^k(X, \mathfrak{u}(r)) \rightarrow \Omega^{k+2}(X, \mathfrak{u}(r))$ defined by $a \mapsto a \wedge \omega$ is called the Lefschetz operator, and its adjoint with respect to the inner product is denoted by $\Lambda : \Omega^k(X, \mathfrak{u}(r)) \rightarrow \Omega^{k-2}(X, \mathfrak{u}(r))$. Both L and Λ are isometries with respect to the inner product.

Proof. The form ϑ is evidently skew-symmetric and bilinear. It is defined independently of $A \in \mathcal{A}$, hence closed. For non-degeneracy, suppose that $0 \neq a \in T_A \mathcal{A}$. Consider the 1-form $b := \Lambda^{n-1} \circ \star a \in \Omega^1(X, \mathfrak{u}(r))$. Then

$$\begin{aligned} \vartheta(a, b) &= \int_X \text{tr}(a \wedge b) \wedge \omega^{n-1} = \int_X \text{tr}(a \wedge \Lambda^{n-1} \circ \star a \wedge \omega^{n-1}) = \int_X \text{tr}(a \wedge L^{n-1} \circ \Lambda^{n-1} \circ \star a) = \\ &= \langle a, L^{n-1} \circ \Lambda^{n-1} a \rangle = \langle \Lambda^{n-1} a, \Lambda^{n-1} a \rangle = \|a\|^2 > 0 \end{aligned}$$

where we used that Λ is an isometry w.r.t. $\langle \cdot, \cdot \rangle$. □

6.3. HERMITIAN-YANG-MILLS AS SYMPLECTIC QUOTIENT

There is an infinite dimensional Lie group acting on \mathcal{A} , namely the group of unitary bundle automorphisms \mathcal{G} . It acts on connections ∇ as $(\varphi \cdot \nabla)s = \varphi \circ \nabla \varphi^{-1}(s)$.

Proposition 6.3.2. *The group \mathcal{G} acts on (\mathcal{A}, ϑ) by symplectomorphisms.*

Proof. First, we need to calculate the differential of the group action $\Theta^\varphi : \mathcal{A} \mapsto \mathcal{A}$, for $\varphi \in \mathcal{G}$. Using the transformation rule of connection 1-form and taking the curve $A + ta$ in \mathcal{A} , we obtain

$$d\Theta_A^\varphi(a) = \left. \frac{d}{dt} \right|_{t=0} \varphi \cdot (A + ta) = \left. \frac{d}{dt} \right|_{t=0} \text{Ad}_\varphi A + t \text{Ad}_\varphi a + \varphi d\varphi^{-1} = \text{Ad}_{\varphi^{-1}} a \in \Omega^1(X, \mathfrak{u}(r))$$

We conclude that $d\Theta_A^\varphi = \text{Ad}_{\varphi^{-1}}$. Since the trace is invariant under conjugation, it follows that Θ^φ is a symplectomorphism of (\mathcal{A}, ϑ) . \square

Denote by $\text{Lie}(\mathcal{G}) = \Gamma(X, \mathfrak{u}(r)) = \Omega^0(X, \mathfrak{u}(r))$ the Lie algebra of \mathcal{G} . Using the trace map and integration again, we make the identification $\text{Lie}(\mathcal{G})^* = \Omega^{2n}(X, \mathfrak{u}(r))$. Then a momentum map for the action of \mathcal{G} is going to be a map $\mu : \mathcal{A} \rightarrow \Omega^{2n}(X, \mathfrak{u}(r))$ such that $d\langle \mu, \xi \rangle = \iota_{\rho(\xi)} \vartheta$ for all $\xi \in \Omega^0(X, \mathfrak{u}(r))$, where $\rho(\xi)$ denotes the fundamental vector field of ξ on \mathcal{A} .

Proposition 6.3.3. *The symplectic action of \mathcal{G} on (\mathcal{A}, ϑ) has momentum map*

$$\mu : A \mapsto F_A \wedge \omega^{n-1}$$

Proof. First we compute the derivative of μ at $A \in \mathcal{A}$, taking the curve $A + ta$. Observe that $F_{A+ta} = d(A + ta) + \frac{1}{2}[A + ta, A + ta] = F_A + t(da + [A, a]) + t^2[a, a]$. Then

$$d\mu_A(a) = \left. \frac{d}{dt} \right|_{t=0} \mu(A + ta) = \left. \frac{d}{dt} \right|_{t=0} F_{A+ta} \wedge \omega^{n-1} = (da + [A, a]) \wedge \omega^{n-1} = d_A a \wedge \omega^{n-1}$$

Next we want to know the fundamental vector field of $\xi \in \Omega^0(X, \mathfrak{u}(r))$. By definition, its value at a point $A \in \mathcal{A}$ is

$$\left. \frac{d}{dt} \right|_{t=0} \exp(t\xi) \cdot A = \left. \frac{d}{dt} \right|_{t=0} \exp(t\xi) \circ A \circ \exp(-t\xi) + \exp(t\xi) d \exp(-t\xi) = \xi \circ A - A \circ \xi - d\xi = -d_A \xi$$

Now we need to check that $\vartheta_A(\rho(\xi), a) = d\langle \mu, \xi \rangle_A(a)$. We use integration by parts and closedness of ω to find

$$\begin{aligned} d\langle \mu, \xi \rangle_A(a) &= \left. \frac{d}{dt} \right|_{t=0} \int_X \text{tr}(F_{A+ta} \wedge \omega^{n-1} \wedge \xi) = \int_X \text{tr}(d_A a \wedge \omega^{n-1} \wedge \xi) = \\ &= \int_X d \text{tr}(a \wedge \omega^{n-1} \wedge \xi) - \int_X \text{tr}(a \wedge \omega^{n-1} \wedge d_A \xi) = - \int_X \text{tr}(a \wedge d_A \xi) \wedge \omega^{n-1} = \vartheta_A(a, \rho(\xi)) \end{aligned}$$

\square

6.3. HERMITIAN-YANG-MILLS AS SYMPLECTIC QUOTIENT

Furthermore, we can give \mathcal{A} the structure of an infinite dimensional Kähler manifold, since the complex structure on X induces an almost complex structure on \mathcal{A} . The latter is an affine space, so the almost complex structure is automatically integrable. Inside of this infinite dimensional Kähler manifold $(\mathcal{A}, \vartheta, J)$ we can consider the submanifold (in fact, it may have singularities - let us not worry about this) which consists of connections whose curvature is of type $(1, 1)$. Denote this subspace by $\mathcal{A}^{(1,1)}$. Because the connections are unitary, the connection is completely determined by its $\bar{\partial}$ -part. We also have a complex gauge group $\mathcal{G}_{\mathbb{C}}$ which acts on these operators via $\bar{\partial} \mapsto \varphi \circ \bar{\partial} \circ \varphi^{-1}$. We may view $\mathcal{G}_{\mathbb{C}}$ as the complexification of \mathcal{G} .

Given the data of the infinite dimensional Kähler manifold with a Hamiltonian Lie group action, we can take the formal symplectic quotient

$$\mathcal{A}^{(1,1)} // \mathcal{G} := \mu^{-1}(\omega^n \otimes \lambda(E)\text{id}_E) / \mathcal{G}$$

These correspond precisely to those gauge orbits which contain a HYM connection. The Kobayashi-Hitchin correspondence may then be phrased as saying that the notion of μ -stability is the right notion for the infinite dimensional Kempf-Ness theorem to hold, i.e.

$$\mathcal{A}^{(1,1)} // \mathcal{G} \cong \mathcal{A}^{(1,1)} / \mathcal{G}_{\mathbb{C}}$$

where the latter is interpreted as the GIT quotient (whose definition relies on the notion of stability). This is a good formal picture to keep in mind. It does not constitute a proof of anything⁴, since many technicalities have to be worked out when dealing with infinite dimensional spaces, as well as the singularities of $\mathcal{A}^{(1,1)}$. In [8], the analytic framework to deal with these problems is presented for algebraic surfaces.

⁴Except that the virtual dimension of the moduli space is $H^1(\mathcal{X}, \text{End}(\mathcal{E}))$, which readily follows from the description $\mathcal{M} = \mathcal{A}^{(1,1)} / \mathcal{G}_{\mathbb{C}}$. But we already knew this.

Chapter 7

Hitchin Systems and $P = W$ Phenomena

In this chapter, we apply the same ideas of infinite dimensional symplectic reduction to obtain some different moduli spaces, namely the so-called Higgs moduli space and the de Rham moduli space. Each of them is a non-compact hyper-Kähler manifold, and it will turn out that we can view one as a hyper-Kähler rotation of the other. It turns out that the former admits a fibration by holomorphic Lagrangian tori, provided by an algebraically completely integrable system. This is a special Lagrangian torus fibration for the rotated complex structure, i.e. the de Rham moduli space, and one can show that (for bundles with structure group $GL(r, \mathbb{C})$) the de Rham moduli space is mirror to itself.

There is also a third moduli space, called the Betti moduli space. It is an affine variety, which is analytically isomorphic to the de Rham moduli space, but not algebraically so. Mirror symmetry does not care about whether the isomorphism is algebraic or analytic, so we may just as well say that the Betti moduli space is mirror to the de Rham moduli space. But the Betti moduli space carries a certain filtration on its cohomology ring, stemming from the fact that it is an algebraic variety. This filtration is preserved by algebraic maps, but not by analytic maps. The question then becomes: what happens to this filtration under the analytic isomorphism between the two spaces? From the perspective of mirror symmetry, the special Lagrangian torus fibration of the de Rham moduli space is an important geometric property, and the $P = W$ conjecture states that it induces a filtration which coincides with the filtration on the cohomology of the Betti moduli space. In this way, mirror symmetry gives rise to an interesting mathematical conjecture, which has in fact been verified (for structure group $G = GL(r, \mathbb{C})$) in [75].

7.1 Mirror Symmetry from Hitchin Systems

Mirror symmetry is exhibited by one particularly beautiful class of examples, namely by various Hitchin systems. These are algebraically completely integrable systems, which are generally quite rare to find. Hitchin constructed a whole class of these, depending on a choice of Riemann surface and a complex reductive algebraic group G . Such a group has a so-called Langlands dual group G^\vee , and it turns out that certain moduli spaces associated to G -bundles and G^\vee -bundles over the chosen Riemann surface exhibit mirror symmetry. These moduli spaces can be obtained through infinite dimensional quotients, in the same way that one obtains the moduli space of HYM equations. This way, mirror symmetry can be used to relate areas within mathematics such as the geometric Langlands correspondence and the duality theory of Lie groups.

7.1. MIRROR SYMMETRY FROM HITCHIN SYSTEMS

Definition 7.1.1. Let (\mathcal{X}, g) be a Kähler manifold, and $G_{\mathbb{C}}$ a complex reductive algebraic group. A $G_{\mathbb{C}}$ -Higgs bundle on \mathcal{X} is a pair (\mathcal{P}, ϕ) where \mathcal{P} is a holomorphic $G_{\mathbb{C}}$ -bundle, and $\phi \in \Gamma(\mathcal{X}, T^* \mathcal{X}^{1,0} \otimes \text{Ad}(\mathcal{P}))$ such that

$$\begin{cases} \bar{\partial}\phi = 0 \\ \phi \wedge \phi = 0 \end{cases}$$

A Higgs bundle is called semi-stable if, for every sub-bundle $\mathcal{F} \subset \mathcal{E}$ such that $\phi(\mathcal{F}) \subset \Omega_{\mathcal{X}}^1 \otimes \mathcal{F}$, we have $\mu(\mathcal{F}) \leq \mu(\mathcal{E})$, and stable if the inequality is strict.

On a Riemann surface, the condition $\phi \wedge \phi = 0$ is vacuous, and this is the case we will consider first. Hitchin studied these moduli spaces in [76], because rank 2 Higgs bundles appeared as solutions to Hitchin's equations which he encountered a year prior, see [77]. This is an interesting story in and of itself, since Hitchin's equations are the dimensional reduction of the Yang-Mills equations to 2 dimensions, but we shall not digress. We will be considering $G_{\mathbb{C}} = \text{GL}(r, \mathbb{C})$ so that we can just look at vector bundles for simplicity, but the story carries over for any complex reductive algebraic group. We fix a fixed smooth vector bundle $E \rightarrow \Sigma$ over a Riemann surface Σ , which is equivalent to specifying its rank r and degree d . Denote the moduli space of holomorphic structures on E by $\mathcal{M}(r, d)$. As we have seen, its tangent space is $H^{0,1}(\Sigma, \text{End}(E))$. Dual to this space is $H^{1,0}(\Sigma, \text{End}(E))$, via the integration map

$$(\phi, A) \mapsto \int_{\Sigma} \text{tr}(\phi \wedge A)$$

In other words, as our notation suggests, the space of Higgs fields ϕ for a given complex structure $\bar{\partial} = \bar{\partial}_0 + A$ on E is precisely $T_A^* \mathcal{M}(r, d)$. In similar vein to the previous section, we can consider the affine space $\mathcal{A} \times \Omega^{1,0}(\Sigma, \text{End}(E)) = T^* \mathcal{A}$. It carries a natural integrable complex structure, with respect to which the map

$$((\alpha_1, \beta_1), (\alpha_2, \beta_2)) \mapsto i \int_{\Sigma} \text{Tr}(\alpha_1 \wedge \beta_2 - \alpha_2 \wedge \beta_1)$$

is a holomorphic symplectic form. It carries a natural action of $\mathcal{G}_{\mathbb{C}}$, namely

$$\varphi \cdot (\bar{\partial}, \phi) = (\varphi \circ \bar{\partial} \circ \varphi^{-1}, \varphi \circ \phi \circ \varphi^{-1})$$

Proposition 7.1.2. *The action of $\mathcal{G}_{\mathbb{C}}$ on $T^* \mathcal{A}$ is symplectic, with moment map*

$$\mu(\bar{\partial}, \phi) = i \int_{\Sigma} \text{Tr}(\xi \wedge \bar{\partial}\phi)$$

where $\xi \in \Gamma(\Sigma, \text{End}(E))$.

The proof is similar to the one we presented when discussing the Kobayashi-Hitchin correspondence, so we leave it as an exercise to the reader.

Definition 7.1.3. The moduli space of (stable) Higgs bundles of rank r and degree d is defined by

$$\mathcal{M}_H(r, d) = T^* \mathcal{A} // \mathcal{G}_{\mathbb{C}}$$

7.1. MIRROR SYMMETRY FROM HITCHIN SYSTEMS

A crucial insight is now that $T^*\mathcal{A}$ is in fact a hyper-Kähler manifold. In this case, it is often possible to instead take a hyper-Kähler quotient $\mathcal{X}/G_{\mathbb{C}} = X //_{\hbar} G$, which gives a hyper-Kähler structure on the quotient. This is presently the case, and we can give $\mathcal{M}_H(r, d)$ the structure of a hyper-Kähler manifold. To do this, we equip E with a Hermitian connection and consider pairs (A, Φ) where $\nabla = \nabla_0 + A$ is a unitary connection on E , and $\Phi \in \Omega^1(\Sigma, u(E))$. The space of such pairs, which we denote \mathcal{H} is an affine space over $\Omega^1(\Sigma, u(E)) \oplus \Omega^1(\Sigma, u(E))$. We see that it is isomorphic to $T^*\mathcal{A}$ by $(\bar{\partial}, \phi) \mapsto (\partial + \bar{\partial}, \phi - \phi^\dagger)$, where ϕ^\dagger denotes the Hermitian conjugate.

It carries three natural complex structures, namely

$$\begin{aligned} I(\alpha, \beta) &= (\star\alpha, -\star\beta) \\ J(\alpha, \beta) &= (-\beta, \alpha) \\ K(\alpha, \beta) &= I \circ J(\alpha, \beta) \end{aligned}$$

Of course the Hodge star is defined w.r.t. the conformal structure on Σ , and the Hermitian metric on E . Furthermore, \mathcal{H} carries three distinct symplectic forms which are holomorphic w.r.t. the respective complex structures:

$$\begin{aligned} \omega_I((\alpha_1, \beta_1), (\alpha_2, \beta_2)) &= \int_{\Sigma} \text{tr}(-\alpha_1 \wedge \alpha_2 + \beta_1 \wedge \beta_2) \\ \omega_J((\alpha_1, \beta_1), (\alpha_2, \beta_2)) &= \int_{\Sigma} \text{tr}(\beta_1 \wedge \star\alpha_2 - \alpha_1 \wedge \star\beta_2) \\ \omega_K((\alpha_1, \beta_1), (\alpha_2, \beta_2)) &= \int_{\Sigma} \text{tr}(\beta_1 \wedge \alpha_2 + \alpha_1 \wedge \beta_2) \end{aligned}$$

The unitary gauge group \mathcal{G} acts on \mathcal{H} by conjugation on both factors, as $\mathcal{G}_{\mathbb{C}}$ does on $T^*\mathcal{A}$.

Theorem 7.1.4. *The metric g on \mathcal{H} defined by*

$$g((\alpha_1, \beta_1), (\alpha_2, \beta_2)) = - \int_{\Sigma} \text{tr}(\alpha_1 \wedge \star\alpha_2 + \beta_1 \wedge \star\beta_2)$$

is a hyper-Kähler metric on \mathcal{H} with respect to I, J, K , and $\omega_I, \omega_J, \omega_K$ are the corresponding symplectic forms. The group \mathcal{G} acts by symplectomorphisms w.r.t. each symplectic form, and carries a hyper-Kähler moment map defined by

$$\begin{aligned} \mu_1(A, \Phi) &= -F_A + \Phi \wedge \Phi - 2\pi i \mu(E)\omega \\ (\mu_2 + i\mu_3)(A, \Phi) &= 2i\bar{\partial}\phi \end{aligned}$$

where ω is the Kähler form on Σ and we used $\mathcal{H} \cong T^\mathcal{A}$.*

The equations $F_A - \Phi \wedge \Phi = -2\pi i \mu(E)\omega$ and $\bar{\partial}\phi = 0$ are precisely Hitchin's equations, and we have a diffeomorphism $\mathcal{M}_H(r, d) \cong \mathcal{H} //_{\hbar} \mathcal{G}$ (after restricting to reducible connections), giving $\mathcal{M}_H(r, d)$ the structure of a hyper-Kähler manifold, and $(\mathcal{H} //_{\hbar} \mathcal{G}, I) \cong$

7.1. MIRROR SYMMETRY FROM HITCHIN SYSTEMS

$\mathcal{M}_H(r, d)$ as complex manifolds. (Actually this is not quite true, because of the difference between the notions of stability. What is true, is that $\mathcal{M}_H(r, d) \subset \mathcal{H} //_{\hbar} \mathcal{G}$ as a dense open subset. But from now on, we define $\mathcal{M}_H(r, d) := \mathcal{H} //_{\hbar} \mathcal{G}$ instead).

So now, we are in a situation where the moduli space of stable Higgs bundles $\mathcal{M}_H(r, d)$ is itself a (non-compact) hyper-Kähler manifold. But Hitchin showed that even more is true. One can give $\mathcal{M}_H(r, d)$ the structure of an algebraically completely integrable system. We do this following [78]. Consider the space $\mathcal{B} = \bigoplus_{k=1}^r H^0(\Sigma, K_{\Sigma}^{\otimes k})$. We can compute the dimension of this space using the Riemann-Roch theorem, which states

$$\dim H^0(\Sigma, \mathcal{F}) - \dim H^1(\Sigma, \mathcal{F}) = \int_{\Sigma} \text{ch}(\mathcal{F}) \text{Td}(\Sigma) = \deg(\mathcal{F}) - g + 1$$

Take $\mathcal{F} = K_{\Sigma}^{\otimes k}$. Serre duality states that

$$H^1(\Sigma, K_{\Sigma}^{\otimes k}) = H^0(\Sigma, (K_{\Sigma}^{\otimes k})^* \otimes K_{\Sigma})^* = H^0(\Sigma, K_{\Sigma}^{1-k})^*$$

We assume that $g > 1$, in which case $\deg(K_{\Sigma}) = 2g - 2 > 0$, implying that $\dim H^0(\Sigma, K_{\Sigma}^{\otimes 1-k}) = 0$ whenever $k > 1$. Using Riemann-Roch again, we also find

$$\dim H^0(\Sigma, K_{\Sigma}^{\otimes k}) = \deg(K_{\Sigma}^{\otimes k}) - g + 1 = k(2g - 2) - g + 1 = (2k - 1)(g - 1)$$

It follows that $\dim_{\mathbb{C}} \mathcal{B} = g + \sum_{k=1}^r (2k - 1)(g - 1) = 1 - r^2(1 - g)$, where the extra g comes from $H^0(\Sigma, K_{\Sigma})$ which we had to consider separately. As noted, $\mathcal{M}_H(r, d) = T^* \mathcal{M}(r, d)$ (morally), and we saw that $T_{\mathcal{E}} \mathcal{M}(r, d) = H^1(\Sigma, \mathcal{E} \text{nd}(\mathcal{E}))$. Now, $\text{ch}(\mathcal{E} \text{nd}(\mathcal{E})) = \text{ch}(\mathcal{E}) \text{ch}(\mathcal{E}^*) = (r + c_1(\mathcal{E}))(r - c_1(\mathcal{E})) = r^2$, so we get

$$\dim \mathcal{M}_H(r, d) = 2 \dim \mathcal{M}(r, d) = 2 \left(1 - \int_{\Sigma} r^2 \text{Td}(\Sigma) \right) = 2(1 - r^2(1 - g))$$

where we used that $\dim H^0(\Sigma, \mathcal{E} \text{nd}(\mathcal{E})) = 1$ because \mathcal{E} is stable by assumption. In conclusion, $\dim \mathcal{B} = \frac{1}{2} \dim \mathcal{M}_H(r, d)$. Then Hitchin defines a holomorphic map

$$\begin{aligned} f : \mathcal{M}_H(r, d) &\rightarrow \mathcal{B} \\ (\mathcal{E}, \phi) &\mapsto (\text{tr}(\phi), \dots, \text{tr}(\wedge^r \phi)) \end{aligned}$$

where we view $\wedge^k \phi : \wedge^k \mathcal{E} \rightarrow \wedge^k (\mathcal{E} \otimes K_{\Sigma}) = (\wedge^k \mathcal{E}) \otimes K_{\Sigma}^{\otimes k}$, and take the natural trace map, so that $\det(\lambda - \phi) = \lambda^r + \sum_{k=1}^r \text{tr}(\wedge^k \phi) \lambda^{r-k}$. In other words, the Hitchin map f sends ϕ to the coefficients of its characteristic polynomial. After choosing a basis in $H^0(\Sigma, K_{\Sigma}^{\otimes k})$, we identify $\mathcal{B} = \mathbb{C}^m$ where $m = 1 - r^2(1 - g)$. We can then separate f into its components $f_j = \pi_j \circ f : \mathcal{M}_H(r, d) \rightarrow \mathbb{C}^m \rightarrow \mathbb{C}$ where π_j is projection onto the j -th coordinate.

Definition 7.1.5. Take $b \in \mathcal{B}$. Define

$$\mathcal{C}_b = \{(z, w) \in T^* \Sigma \mid \det(w - \phi(z)) = 0\}$$

where z is a local holomorphic coordinate on Σ , and w a holomorphic coordinate for the fibre of $T^* \Sigma$. Then \mathcal{C}_b is called the spectral curve of $b \in \mathcal{B}$.

7.1. MIRROR SYMMETRY FROM HITCHIN SYSTEMS

Evidently, the characteristic polynomial will generically have r distinct solutions, so that we get an r -sheeted ramified covering map $\pi : \mathcal{C}_b \rightarrow \Sigma$. Hitchin then proves the following.

Theorem 7.1.6. *The fibre $f^{-1}(b)$ of the Hitchin map is $\text{Jac}(\mathcal{C}_b)$, the Jacobian variety of the spectral curve.*

More generally, there is a singular locus, the discriminant locus, in the base \mathcal{B} over which the spectral curve is singular, and one has to be somewhat more careful. But for a dense open subset in \mathcal{B} , the generic fibre is a projective complex torus. In summary:

Theorem 7.1.7 (Hitchin [76]). *The moduli space of stable Higgs bundles $\mathcal{M}_H(r, d)$ is a hyper-Kähler manifold, which admits the structure of an algebraically completely integrable system $f : \mathcal{M}_H(r, d) \rightarrow \mathcal{B}$. The fibres of f are the Jacobian varieties of the spectral curves $\mathcal{C}_b \subset T^*\Sigma$, which are Lagrangian w.r.t. the holomorphic symplectic form $\vartheta_I = \omega_J + i\omega_K$.*

Proof. It remains to show that the fibres are Lagrangian. To see this, fix $1 \leq k \leq r$ and take $a \in \Omega^{0,1}(\Sigma, K_\Sigma^{\otimes(1-k)})$. Define $f_a(\bar{\partial}, \phi) = \int_\Sigma \text{tr}(\phi^k) \wedge a : T^*\mathcal{A} \rightarrow \mathbb{C}$. Then

$$\begin{aligned} df_a(\alpha, \beta) &= \int_\Sigma \text{tr}(\dot{\phi} \wedge \phi^{k-1}) \wedge a = \\ &= i \int_\Sigma \text{tr}(\dot{\phi} \wedge (-ia) \wedge \phi^{k-1}) = \vartheta_I((-ia\phi^{k-1}, 0), (\alpha, \beta)) \end{aligned}$$

where we have taken a curve with $\dot{\gamma}(0) = (\alpha, \beta)$ to compute the differential. We conclude that the Hamiltonian vector field of df_a is $(-ia\phi^{k-1}, 0)$, which are all tangent to the fibres since the second component, which is the component in $T\mathcal{B}$, vanishes. For this reason, all the f_a Poisson commute for various a , also after descending to the quotient $\mathcal{M}_H(r, d)$. We can choose sufficiently many a_i such that $\{df_{a_i}\}_{i \in I}$ span $T^*\mathcal{B}$, because of the non-degenerate pairing between $H^{0,1}(\mathcal{X}, K^{\otimes(1-k)})$ and $H^0(\mathcal{X}, K^{\otimes k})$. By non-degeneracy of Ω_I , the corresponding Hamiltonian vector field span the tangent space to the fibre, which leads one to conclude that $\varphi_I|_{f^{-1}(b)} = 0$ for all $b \in \mathcal{B}$. Hence, the fibres are Lagrangian. \square

Now, the fact that the Lagrangian fibres are holomorphically embedded w.r.t. to the complex structure I means that they are calibrated by ω_I . The fact that they are Lagrangian w.r.t. ϑ_I means that, for each fibre $f^{-1}(b) = L_b$, we have $\omega_J|_{L_b} = \omega_K|_{L_b} = 0$. Because of this, we have the following corollary.

Corollary 7.1.8. *The Hitchin map $f : \mathcal{M}_H(r, d) \rightarrow \mathcal{B}$ is a special Lagrangian fibration for the Calabi-Yau manifold $(M_H(r, d), \omega_K, \Omega_K = (\omega_I + i\omega_J)^{n/2})$, where $n = \dim_{\mathbb{C}} \mathcal{M}_H(r, d)$.*

So what interpretation do the complex structures J and K on $\mathcal{M}_H(r, d)$ have, if any? Fortunately for us, it has a very nice geometric interpretation. Recall that J acts by $J(\alpha, \beta) = (-\beta, \alpha)$. This corresponds to multiplying $\alpha + i\beta$ by i , and so the space $(T^*\mathcal{A}, J)$ consists of all complex connections on E , not unitary ones. We can split a complex connection ∇ into its unitary and self-adjoint part, writing $\nabla = D + i\Phi$. Then $\mathcal{G}_{\mathbb{C}}$ acts by

7.1. MIRROR SYMMETRY FROM HITCHIN SYSTEMS

complex gauge transformations on the space of complex connections. Its momentum map is

$$\mu(D, \Phi) = \mu_3 + i\mu_1 = D\Phi + i(-F_D + \Phi \wedge \Phi - 2\pi i\mu(E)\omega)$$

The space $\mu^{-1}(0)$ consists of those complex connections satisfying $F_D + iD\Phi - \Phi \wedge \Phi = F_{\nabla} = -2\pi i\mu(E)\omega$. In other words, those complex connections with constant central curvature. We consider the case $d = 0$, meaning we obtain the space of flat complex connections. In other words, $(M_H(r, d), J)$ is the moduli space of flat connections on E . Its topology and smooth structure comes from infinite dimensional symplectic reduction, as we have been discussing. Denote this moduli space by $\mathcal{M}_{dR}(r, d)$. Recall that a flat connection on E corresponds to a representation $\rho : \pi_1(\Sigma) \rightarrow \mathrm{GL}(r, \mathbb{C})$. Connections are gauge equivalent if and only if the representations are conjugate. Associated to the general linear group is its representation variety, which is an affine variety. Using geometric invariant theory, one can look at the moduli space of these representations (i.e. the conjugacy classes of representations) which is an affine algebraic variety associated to a ring of invariant polynomials. After endowing it with the analytic topology instead of the Zariski topology, we get a non-compact algebraic complex manifold denoted $\mathcal{M}_B(r, 0)$ (a minor adaptation generalises to any $d \neq 0$). It is not at all clear if $\mathcal{M}_{dR}(r, 0)$ and $\mathcal{M}_B(r, 0)$ are diffeomorphic, since their constructions are very different. All we know is that there is a natural bijection between them, which may or may not preserve the possibly very distinct topologies and smooth structures.

Theorem 7.1.9 (Non-abelian Hodge Correspondence [79]). *Let \mathcal{X} be a projective Kähler manifold. Then there are diffeomorphisms*

$$M_H(r, d) \cong M_{dR}(r, d) \cong M_B(r, d)$$

The Riemann-Hilbert correspondence gives an analytic isomorphism $\mathcal{M}_{dR}(r, d) \cong \mathcal{M}_B(r, d)$. Note that we are not dealing with projective manifolds now, so analytic isomorphisms are not necessarily algebraic, and indeed, this isomorphism is not an algebraic one. In conclusion, we have a non-compact hyper-Kähler manifold $(M_H(r, d), g)$ for which $(M_H(r, d), g, I) \cong \mathcal{M}_H(r, d)$ and $(M_H(r, d), g, J) \cong \mathcal{M}_{dR}(r, d) \cong \mathcal{M}_B(r, d)$, but the final isomorphism is only complex analytic, and not algebraic. We have outlined the intuition for the case when $\mathcal{X} = \Sigma$ is a Riemann surface, but remarkably this theorem holds much more generally. Of course, one has to impose the additional condition $\phi \wedge \phi = 0$ on the Higgs field in higher dimensions.

Convention 2. From now on, we will write $\mathcal{M}_H := \mathcal{M}_H(1, 0)$ and $\mathcal{M}_B := \mathcal{M}_B(1, 0)$.

Example 24. Let $\mathcal{X} = \mathbb{C}^n / \Lambda$ be an abelian variety with a principal polarisation, so that $\mathrm{Jac}(\mathcal{X}) \cong \mathcal{X}$. The moduli space of degree 0 line bundles on \mathcal{X} is precisely $\mathrm{Jac}(\mathcal{X})$, and its tangent space is $H^{0,1}(\mathcal{X})$ as we noted previously. The condition $\phi \wedge \phi = 0$ is always satisfied for line bundles. Therefore,

$$T^* \mathcal{M} = \mathcal{M}_H = \mathrm{Jac}(\mathcal{X}) \times H^{1,0}(\mathcal{X}) \cong \mathcal{X} \times \mathbb{C}^n$$

7.1. MIRROR SYMMETRY FROM HITCHIN SYSTEMS

The Hitchin map is the projection $\mathcal{X} \times \mathbb{C}^n \rightarrow \mathbb{C}^n$. Now we want to find $\mathcal{M}_H(r, 0)$ for $r > 0$. Suppose (\mathcal{E}, ϕ) is a semistable Higgs bundle on \mathcal{X} . It is S -equivalent to the direct sum of the sub-quotients of its Jordan-Hölder filtration, so we may take a direct sum $(\mathcal{L}_1 \oplus \cdots \oplus \mathcal{L}_r, \phi|_{\mathcal{L}_1} \oplus \cdots \oplus \phi|_{\mathcal{L}_r})$ to represent its isomorphism class. Since $\phi|_{\mathcal{L}_i} \in \text{Hom}(\mathcal{L}_i, \mathcal{L}_i \otimes \Omega_{\mathcal{X}}^1)$ and $\Omega_{\mathcal{X}}^1 \cong \mathcal{O}_{\mathcal{X}}^{\oplus n}$, we can identify $(\phi|_{\mathcal{L}_1}, \dots, \phi|_{\mathcal{L}_r}) \in \mathbb{C}^n$. There is an obvious map

$$\begin{aligned} \mathcal{M}_H(r, 0) &\rightarrow \text{Sym}^r(\text{Jac}(\mathcal{X}) \times \mathbb{C}^n) \\ (\mathcal{E}, \phi) &\mapsto (\mathcal{L}_1, \dots, \mathcal{L}_r, \phi|_{\mathcal{L}_1}, \dots, \phi|_{\mathcal{L}_r}) \end{aligned}$$

which likewise has an obvious inverse. In [15], it is shown that this is indeed an isomorphism of varieties. We conclude that $\mathcal{M}_H(r, 0) \cong \text{Sym}^r(\mathcal{X} \times \mathbb{C}^n)$.

Return briefly to the case $r = 1$. Then the Betti moduli space is $\text{Hom}(\Lambda, \mathbb{C}^\times)$, since \mathbb{C}^\times is abelian so conjugation acts trivially. We have diffeomorphisms

$$\text{Hom}(\Lambda, \mathbb{C}^\times) \cong \text{Hom}(\Lambda, S^1) \times \text{Hom}(\Lambda, \mathbb{R}) \cong \mathcal{X} \times H^{1,0}(X, \mathbb{R}) \cong \mathcal{M}_H$$

The first isomorphism is polar decomposition, the second isomorphism is the Kobayashi-Hitchin correspondence (since an element of $\text{Hom}(\Lambda, S^1)$ is a flat unitary connection), and the Hodge decomposition theorem $\eta = \phi + \bar{\phi} \mapsto \phi$ (since $\text{Hom}(\Lambda, \mathbb{R}) = H^1(X, \mathbb{R})$). Now, when $r > 1$ we are just looking for tuples of $2n$ invertible matrices which pairwise commute, and geometric invariant theory yields that they additionally need to be simultaneously diagonalisable. It follows that $\mathcal{M}_B(r, 0) \cong \text{Sym}^r((\mathbb{C}^\times)^{2n})$, and we again get a diffeomorphism

$$\mathcal{M}_H(r, 0) \cong \mathcal{M}_B(r, 0)$$

Note that $\mathcal{M}_B(r, 0)$ is an affine variety, so it does not contain any projective subvarieties. Hence, it is clear that $\mathcal{M}_H(r, d)$ and $\mathcal{M}_B(r, d)$ are not biholomorphic, since $\mathcal{M}_H(r, d)$ contains many projective subvarieties - the fibres of the Hitchin map.

Remark 7.1.10. The Kobayashi-Hitchin correspondence can be viewed as a corollary to the non-abelian Hodge correspondence. Indeed, if $\phi = 0$ then the condition of stability for a Higgs bundle is just the same as the underlying holomorphic vector bundle being stable. The subset of those representations of $\pi_1(X)$ which are unitary can be shown to correspond to the set of Higgs bundles with $\phi = 0$. This yields the Kobayashi-Hitchin correspondence for $d = 0$, and arbitrary degree follows by making small modifications.

Now, we can also consider moduli spaces of $G_{\mathbb{C}}$ -Higgs bundles for any complex reductive algebraic group $G_{\mathbb{C}}$. In this case, the Higgs field becomes a holomorphic section $\phi \in H^0(\Sigma, \text{Ad}(\mathcal{P}) \otimes K_{\Sigma})$, and the topological type is fixed by an element $d \in \pi_1(G_{\mathbb{C}})$. The resulting constructions are highly analogous and can be found in Hitchin's original paper.

7.2. THE $P = W$ PHENOMENON

We denote the respective moduli spaces by $\mathcal{M}_H(G_{\mathbb{C}}, d)$, $\mathcal{M}_{dR}(G_{\mathbb{C}}, d)$ and $\mathcal{M}_B(G_{\mathbb{C}}, d)$. For a complex reductive algebraic group $G_{\mathbb{C}}$, we can consider its root datum $(\Gamma^*, \Delta, \Gamma_*, \Delta^\vee)$ where Γ^* is the lattice of characters of a maximal torus, Γ_* is the dual lattice, Δ the roots of the Lie algebra, and Δ^\vee the coroots. There is an involution on the set of complex reductive algebraic groups, defined by

$$(\Gamma^*, \Delta, \Gamma_*, \Delta^\vee) \mapsto (\Gamma_*, \Delta^\vee, \Gamma^*, \Delta)$$

The image of a group $G_{\mathbb{C}}$ under this involution is called its Langlands dual, denoted by $G_{\mathbb{C}}^\vee$.

Example 25. Some examples of Langlands dual groups:

1. $\mathrm{GL}(r, \mathbb{C})^\vee = \mathrm{GL}(r, \mathbb{C})$
2. $\mathrm{SL}(r, \mathbb{C})^\vee = \mathrm{PGL}(r, \mathbb{C})$
3. $\mathrm{SO}(2n+1, \mathbb{C})^\vee = \mathrm{Sp}(2n, \mathbb{C})$
4. $\mathrm{SO}(2n, \mathbb{C})^\vee = \mathrm{SO}(2n, \mathbb{C})$

Theorem 7.1.11 ([80]). *Let $G_{\mathbb{C}}$ be a complex reductive group and $G_{\mathbb{C}}^\vee$ its Langlands dual. Let $f : \mathcal{M}_{dR}(G_{\mathbb{C}}, d) \rightarrow B$ and $f^\vee : \mathcal{M}_{dR}(G_{\mathbb{C}}^\vee, d) \rightarrow B^\vee$ be the respective special Lagrangian fibrations obtained from the Hitchin system, for some Riemann surface. Then $B = B^\vee$ and the fibrations are an SYZ mirror pair.¹*

The Langlands dual group of $\mathrm{GL}(r, \mathbb{C})$ is itself, which means the corresponding special Lagrangian fibration is mirror to itself. Recall also that we have a complex analytic isomorphism $\mathcal{M}_{dR}(r, d) \cong \mathcal{M}_B(r, d)$, so $\mathcal{M}_B(r, d)$ is mirror to $\mathcal{M}_{dR}(r, d)$. The question is: can the Betti moduli space "see" the special Lagrangian fibration of the de Rham moduli space in some way? The $P = W$ conjecture claims that it can, through a certain filtration on its cohomology which we explore next.

7.2 The $P = W$ Phenomenon

Every complex algebraic variety comes equipped with its weight filtration, which is an abstraction of the Hodge decomposition on $H^\bullet(X, \mathbb{Q})$ that also works for non-projective varieties. In particular, the Betti moduli space has a canonical weight filtration associated to it, denote it by $W_\bullet(H^\bullet(M_B(r, d), \mathbb{Q}))$. We will give a precise definition in a moment. The special Lagrangian fibration $\mathcal{M}_{dR}(r, d) \rightarrow B$ also induces a filtration on cohomology, through the perverse filtration. This is a filtration which can be associated to a proper morphism between irreducible smooth quasi-projective varieties. We can hyper-Kähler rotate our special Lagrangian fibration to return to the Hitchin map, which sat-

¹Each of these moduli spaces also has a natural B -field on it, coming from a $U(1)$ -gerbe.

7.2. THE $P = W$ PHENOMENON

isfies these conditions. As such, we get a filtration $P_\bullet(H^\bullet(M_H(r, d), \mathbb{Q}))$. Again, a precise definition will be given in a moment.

Conjecture 7.2.1 (The $P = W$ conjecture [16]). *Under $M_H(r, d) = M_{dR}(r, d) \cong M_B(r, d)$, we have*

$$P_k(H^m(M_H(r, d), \mathbb{Q})) = W_{2k}(H^m(M_B(r, d), \mathbb{Q})) = W_{2k+1}(H^m(M_B(r, d), \mathbb{Q}))$$

One can really think of this as a refined version of topological mirror symmetry (i.e. at the level of Hodge numbers) for non-projective varieties. The $P = W$ conjecture has been shown to hold for $G = \mathrm{GL}(r, \mathbb{C})$ in [75]. Other structure groups and their Langlands duals give rise to a similar statement, as was addressed in the original paper, where they prove their conjecture for certain low rank groups and their duals. In [15], it was observed that the $P = W$ phenomenon holds for the Higgs and Betti moduli spaces associated to abelian varieties of any dimension. This is the example that we will investigate further, as it allows for some explicit computations.

To treat these examples, we need to familiarise ourselves with the respective cohomology filtrations, starting with Deligne's weight filtration.

Definition 7.2.2. A pure Hodge structure of weight k is an abelian group $H_{\mathbb{Z}}$ of finite rank, together with a filtration

$$H_{\mathbb{Z}} \otimes \mathbb{C} = F^0 \supseteq F^1 \supseteq \dots \supseteq F^{k+1} = 0$$

such that $F^p \cap \overline{F^q} = 0$ and $F^p \oplus \overline{F^q} = H_{\mathbb{Z}} \otimes \mathbb{C}$ for all $p + q = k + 1$.

Example 26. Take any compact Kähler manifold \mathcal{X} and let $H_{\mathbb{Z}} = H^k(X, \mathbb{Z})$ and $F^p = \oplus_{i \geq p} H^{i, k-i}$. This yields a pure Hodge structure of weight k .

Theorem 7.2.3 (Deligne). *Let \mathcal{X} be a quasi-projective \mathbb{C} -variety. Then every cohomology group $H^m(X, \mathbb{Q})$ carries a natural weight filtration*

$$0 = W_{-1}H^m(X, \mathbb{Q}) \subseteq W_0(H^m(X, \mathbb{Q})) \subseteq \dots \subseteq W_{2m}H^m(X, \mathbb{Q}) = H^m(X, \mathbb{Q})$$

and a Hodge filtration

$$H^m(X, \mathbb{C}) = F^0 H^m(X, \mathbb{C}) \supseteq \dots \supseteq F^m H^m(X, \mathbb{C}) \supseteq F^{m+1} H^m(X, \mathbb{C}) = 0$$

such that the sub-quotients $\mathrm{Gr}_k^W(H^m(X, \mathbb{C})) = W_k H^m(X, \mathbb{C}) / W_{k-1} H^m(X, \mathbb{C})$ carry a pure Hodge structure of weight k induced by F^\bullet . This is called a mixed Hodge structure. It satisfies the following properties:

1. If \mathcal{X} is projective and smooth, then

$$0 = W_{m-1}H^m(X, \mathbb{Q}) \subseteq W_m H^m(X, \mathbb{Q}) = H^m(X, \mathbb{Q})$$

7.2. THE $P = W$ PHENOMENON

2. If \mathcal{X} is smooth (but not necessarily projective) then

$$0 = W_{-1}H^m(X, \mathbb{Q}) \subseteq \cdots \subseteq W_{2m}H^m(X, \mathbb{Q}) = H^m(X, \mathbb{Q})$$

and for any smooth compactification $\iota: \mathcal{X} \rightarrow \overline{\mathcal{X}}$ one has

$$W_m H^m(X, \mathbb{Q}) = \iota^* H^m(\overline{\mathcal{X}}, \mathbb{Q})$$

3. For any algebraic map $f: \mathcal{X} \rightarrow \mathcal{Y}$, one has

$$f^*(W_k H^m(Y, \mathbb{Q})) \subseteq W_k H^m(X, \mathbb{Q})$$

4. The weight filtration is compatible with the cup product, so we may write

$$W_k H^\bullet(X, \mathbb{Q})$$

5. The weight filtration is compatible with the Künneth formula

$$W_k H^\bullet(X \times Y, \mathbb{Q}) \cong \bigoplus_{i+j=k} W_i H^\bullet(X, \mathbb{Q}) \otimes W_j H^\bullet(Y, \mathbb{Q})$$

The theorem above is rather lengthy, but it mostly just expresses the naturality of the weight filtration with respect to natural operations such as compactifications, pullback, products, etc. The weight filtration of Deligne's theorem will be denoted by $W_\bullet H^\bullet(X, \mathbb{Q})$.

Example 27. Let's consider $\mathcal{X} = \mathbb{C}^\times$. There is a smooth compactification $\iota: \mathcal{X} \rightarrow \mathbb{CP}^1$, and $H^1(\mathbb{CP}^1, \mathbb{Q}) = 0$. It follows from 2) that $W_1 H^1(X, \mathbb{Q}) = 0$. This is sufficient to determine the entire weight filtration of \mathcal{X} , since \mathcal{X} is smooth:

$$\begin{aligned} W_0 H^0(X, \mathbb{Q}) &= H^0(X, \mathbb{Q}) = \mathbb{Q} \\ W_0 H^1(X, \mathbb{Q}) &= W_1 H^1(X, \mathbb{Q}) = 0 & W_2 H^1(X, \mathbb{Q}) &= H^1(X, \mathbb{Q}) = \mathbb{Q} \\ W_4 H^2(X, \mathbb{Q}) &= H^2(X, \mathbb{Q}) = 0 \end{aligned}$$

By the Künneth formula, we can now compute the weight filtration for $\mathcal{M}_B(r, 0) = (\mathbb{C}^\times)^{2n}$, and it is given essentially by the binomial theorem. That is, the $2k$ -th step of the weight filtration consists of the first k summands of the cohomology of T^{2n} , which is the exterior algebra:

$$W_{2k} H^m((\mathbb{C}^\times)^{2n}, \mathbb{Q}) = \bigoplus_{j=0}^k \wedge^j \mathbb{Q}^{2n} = W_{2k+1} H^m((\mathbb{C}^\times)^{2n}, \mathbb{Q})$$

We define $H^{k;p,q}(\mathcal{X}) := \mathrm{Gr}_F^p \mathrm{Gr}_{p+q}^W H^k(X, \mathbb{C})$.

7.2. THE $P = W$ PHENOMENON

Definition 7.2.4. The Deligne-Hodge numbers are $h^{k;p,q}(\mathcal{X}) = \dim_{\mathbb{C}} H^{k;p,q}(\mathcal{X})$. The Deligne-Hodge polynomial is defined as

$$\mathfrak{W}_{\mathcal{X}}(t, u, v) = \sum_{k,p,q \geq 0} h^{k;p,q}(\mathcal{X}) t^k u^p v^q$$

The Deligne-Hodge polynomial can be treated as a refinement of the Hodge numbers of a compact Kähler manifold, which satisfies additional convenient properties and relations. For example, it follows from Deligne's theorem that the polynomials are multiplicative with respect to Cartesian products.

Example 28. Take an abelian variety \mathcal{X} . We saw in 24 that the corresponding Higgs moduli space of rank 1 and degree 0 is $\mathcal{M}_H \cong \mathcal{X} \times \mathbb{C}^n$. Its Deligne-Hodge polynomial is easy to deduce, since this is a Cartesian product, and \mathcal{X} is smooth and projective. By equipping an abelian variety with a flat metric, we can compute the harmonic forms, as they are the constant ones. Thus, we can compute the Deligne-Hodge polynomial as the product of the Deligne-Hodge polynomial of an elliptic curve, which is $\mathfrak{W}(t, u, v) = 1 + tu + tv + t^2 uv = (1 + tu)(1 + tv)$. We also have $\mathfrak{W}_{\mathbb{C}} = 1$. Therefore,

$$\mathfrak{W}_{\mathcal{X} \times \mathbb{C}^n}(t, u, v) = (1 + tu)^n (1 + tv)^n$$

We saw that the Betti moduli space is $\mathcal{M}_B = (\mathbb{C}^\times)^{2n}$. We have $H^0(\mathbb{C}^\times, \mathbb{C}) \cong H^1(\mathbb{C}^\times, \mathbb{C}) \cong \mathbb{C}$, and all other cohomology groups vanish. Since \mathbb{C}^\times is smooth, we get $h^{0;0,0}(\mathcal{X}) = h^{1;1,1}(\mathcal{X}) = 1$, with all other Hodge numbers vanishing, so $\mathfrak{W}_{\mathbb{C}^\times} = 1 + tuv$ and

$$\mathfrak{W}_{(\mathbb{C}^\times)^{2n}} = (1 + tuv)^{2n}$$

We deduce that $H^\bullet(\mathcal{M}_B, \mathbb{Q}) = \oplus_p H^{p,p}((\mathbb{C}^\times)^{2n}, \mathbb{Q})$, so in fact we can simply define the Hodge numbers of \mathcal{M}_B by $h^{p,p}(\mathcal{M}_B) = h^{p;p,p}(\mathcal{M}_B)$, and all other Hodge numbers vanishing.

Clearly, $\mathfrak{W}_{\mathcal{M}_H} \neq \mathfrak{W}_{\mathcal{M}_B}$ in spite of the diffeomorphism between the two spaces. Instead, the $P = W$ conjecture says that the perverse filtration of \mathcal{M}_H should coincide with the weight filtration of \mathcal{M}_B , so let us see how this perverse filtration is defined.

Definition 7.2.5. Let \mathcal{Y} be a quasi-projective variety equipped with a flag $\mathcal{Y}_0 \subset \dots \subset \mathcal{Y}_d = \mathcal{Y}$ (generic with $\dim \mathcal{Y}_k = k$), and let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a morphism of varieties. Then the perverse P_\bullet filtration on $H^\bullet(\mathcal{X}, \mathbb{Q})$ is defined by

$$P_{m-k-1} H^m(\mathcal{X}, \mathbb{Q}) = \ker(H^m(\mathcal{X}, \mathbb{Q}) \rightarrow H^m(f^{-1}(Y_k), \mathbb{Q}))$$

Remark 7.2.6. Actually the above definition is a massive shortcut. The perverse filtration can be defined in much more general context, which involves significantly more machinery as well, such as in the references we have given [16, 75].

7.2. THE $P = W$ PHENOMENON

Definition 7.2.7. Let $P \cdot H^\bullet(X, \mathbb{Q})$ be a perverse filtration. Then the perverse Hodge numbers $h_{\mathfrak{P}}^{p,q}(\mathcal{X})$ are defined as

$$h_{\mathfrak{P}}^{p,q}(\mathcal{X}) = \dim_{\mathbb{Q}} \text{Gr}_p^P H^q(X, \mathbb{Q}) = \dim_{\mathbb{Q}} P_p H^q(X, \mathbb{Q}) / P_{p-1} H^q(X, \mathbb{Q})$$

Our goal in this section will be to check if it is true that

$$h_{\mathfrak{P}}^{p,q}(\mathcal{M}_H(r, 0)) = h^{p,q}(\mathcal{M}_B(r, 0))$$

Example 29. We are interested in the perverse filtration associated to the Hitchin map $f : \mathcal{M}_H(r, 0) \cong \text{Sym}^r(\mathcal{X} \times \mathbb{C}^n) \rightarrow \mathbb{C}^r$. We start with the case $r = 1$, i.e. $f : \mathcal{X} \times \mathbb{C}^n \rightarrow \mathbb{C}^n$ the projection map. Homotopy invariance of cohomology yields

$$P_k H^\bullet(X \times \mathbb{C}^n, \mathbb{Q}) \cong P_k H^\bullet(X, \mathbb{Q})$$

So the perverse filtration on \mathcal{M}_H is induced by the constant map $c : X \rightarrow \{*\}$. It follows from the definition that

$$P_k H^m(\mathcal{M}_H, \mathbb{Q}) = P_k H^m(X, \mathbb{Q}) = \begin{cases} H^m(X, \mathbb{Q}) & \text{if } k \geq m \\ 0 & \text{if } k < m \end{cases}$$

Equivalently, we have

$$P_k H^\bullet(\mathcal{M}_H, \mathbb{Q}) = \bigoplus_{j=0}^k \wedge^j \mathbb{Q}^{2n}$$

which shows that the $P = W$ conjecture holds for \mathcal{X} and $r = 1$, as per [15]. To find $h_p^{p,q}(\mathcal{M}_H)$, we need the dimension of

$$\text{Gr}_p^P H^{p+q}(\mathcal{M}_H, \mathbb{Q}) = P_p H^{p+q}(\mathcal{M}_H, \mathbb{Q}) / P_{p-1} H^{p+q}(\mathcal{M}_H, \mathbb{Q})$$

So fix $p, q \in \mathbb{N}$ and let $k = p + q$. From the description of the perverse filtration we see that this quotient is non-trivial if and only if $H^k(X, \mathbb{Q}) \neq 0$ and $p = k$, in which case it is $H^k(X, \mathbb{Q})$. We conclude that

$$h_{\mathfrak{P}}^{p,q}(\mathcal{M}_H) = \begin{cases} b_p(X) & \text{if } p = q \\ 0 & \text{else} \end{cases}$$

The generating function is $\mathfrak{P}_{\mathcal{M}_H}(u, v) = \sum_{p,q} h_{\mathfrak{P}}^{p,q}(\mathcal{M}_H) u^p v^q = (1 + uv)^{2n}$. We conclude that

$$\mathfrak{W}_{\mathcal{M}_B}(1, u, v) = \mathfrak{P}_{\mathcal{M}_H}(uv, v)$$

Thus, the Hodge numbers of the perverse filtration on $H^\bullet(\mathcal{M}_H, \mathbb{Q})$ match those of the weight filtration on $H^\bullet(\mathcal{M}_B, \mathbb{Q})$:

$$h_{\mathfrak{P}}^{p,q}(\mathcal{M}_H) = h^{p,q}(\mathcal{M}_B)$$

7.3. ENUMERATIVE GEOMETRY AND $P = W$

Now, using the fact that $\mathcal{M}_H(r, 0) = \text{Sym}^r(\mathcal{M}_H)$ and $\mathcal{M}_B(r, 0) = \text{Sym}^r(\mathcal{M}_B)$, we want to deduce the $r > 1$ case from this. Given a variety \mathcal{M} , one obtains $\text{Sym}^r(\mathcal{M})$ as a quotient of \mathcal{M}^r by a group action of the symmetric group S_r . From this, it follows that $H^*(\text{Sym}^r(M), \mathbb{Q}) = H^*(M^r, \mathbb{Q})^{S_r}$, the S_r -equivariant cohomology ring of M^r . In [81] it is shown that

$$P_p H^*(\text{Sym}^r(M), \mathbb{Q}) = (P_p H^*(M^r, \mathbb{Q}))^{S_r}$$

Furthermore, naturality of the weight filtration implies that

$$W_{2k} H^*(\text{Sym}^r(M), \mathbb{Q}) = W_{2k+1} H^*(\text{Sym}^r(M), \mathbb{Q}) = (W_{2k} H^*(M^r, \mathbb{Q}))^{S_r}$$

This is how the $P = W$ conjecture is proved for $r > 1$ in [15], for any abelian variety. It also implies the result that we wanted to verify.

Corollary 7.2.8. *Let \mathcal{X} be an abelian variety and let $\mathcal{M}_H(r, 0)$ and $\mathcal{M}_B(r, 0)$ be the Higgs moduli space and the Betti moduli space, respectively. Then*

$$h_{\mathfrak{P}}^{p,q}(\mathcal{M}_H(r, 0)) = h^{p,q}(\mathcal{M}_B(r, 0))$$

Now, one can use the results from [82] to obtain the perverse Hodge numbers of $\mathcal{M}_H(r, 0)$. In the citation, a formula is given for the Deligne-Hodge polynomials of symmetric products of certain algebraic groups, such as $\text{Sym}^r((\mathbb{C}^\times)^{2n})$. Let $\sigma \in S_r$ and let M_σ denote the corresponding permutation matrix. Then the result states that

$$\mathfrak{W}_{\text{Sym}^r((\mathbb{C}^\times)^{2n})}(t, u, v, \cdot) = \frac{1}{r!} \sum_{\sigma \in S_r} \det(I + tuvM_\sigma)^{2n}$$

And indeed, we see that $r = 1$ just yields the familiar formula $(1 + tuv)^{2n}$.

7.3 Enumerative Geometry and $P = W$

We saw that the moduli spaces that are involved in the $P = W$ conjecture can be obtained by infinite dimensional symplectic reduction, akin to the moduli space of holomorphic structures on a complex vector bundle on a Kähler manifold. Aside from this analogy, and the appearance of a special Lagrangian fibration, it turns out that the $P = W$ phenomenon is closely related to enumerative geometry, and can be used to obtain (local) Donaldson-Thomas invariants and (local) Gromov-Witten invariants. We briefly explain how this comes about.

Let \mathcal{X} be a Riemann surface. It turns out (see [83]) that the perverse Hodge numbers $h_{\mathfrak{P}}^{p,q}(\mathcal{M}_H(r, d))$ of the associated Higgs moduli space are the Gopakumar-Vafa invariants of the local Calabi-Yau threefold $\text{Tot}(K_{\mathcal{X}} \oplus \mathcal{L})$, where $\deg(\mathcal{L}) = 2 - 2g$ and Tot denotes the total space of the vector bundle. In the case of an elliptic curve, which we treated above, this is just $\mathcal{X} \times \mathbb{C}^2$. The Gopakumar-Vafa invariants are certain invariants of a Calabi-Yau threefold which are motivated by M -theory, and they count curves in the Calabi-Yau

7.3. ENUMERATIVE GEOMETRY AND $P = W$

threefold together with line bundles on them, in some appropriate sense. There are also invariants called Pandharipande-Thomas invariants (introduced in [84]) which count curves with points on them (scheme theoretically) in a Calabi-Yau threefold. It has been shown that the generating functions of these invariants coincide with the (reduced) generating functions of the Donaldson-Thomas invariants, i.e.

$$Z_{\beta}^{\text{PT}}(q) = Z_{\beta}^{\text{red}}(q)$$

This is established in [85, 86]. So essentially Donaldson-Thomas theory and Pandharipande-Thomas theory are equivalent ways of counting curves in a Calabi-Yau threefold.

Gopakumar-Vafa invariants are a refined invariant in the sense that they count curves, as well as line bundles on them. In the case of the local Calabi-Yau threefold associated to a curve, these are the spectral curves of the Hitchin system, which lie in the total space of $K_{\mathcal{X}} \oplus \mathcal{O}_{\mathcal{X}}$, with line bundles on them. The corresponding moduli space is the Higgs moduli space $\mathcal{M}_H(r, d)$ which is why $h_{\mathfrak{X}}^{p,q}(\mathcal{M}_H(r, d))$ arise as Gopakumar-Vafa invariants. In [87], there is a conjectured correspondence between Pandharipande-Thomas and (conjecturally defined) Gopakumar-Vafa invariants. In other words, the perverse Hodge numbers $h_{\mathfrak{X}}^{p,q}(\mathcal{M}_H(r, d))$ determine the Pandharipande-Thomas invariants, and in turn, the Donaldson-Thomas invariants of the local Calabi-Yau threefold $\text{Tot}(K_{\mathcal{X}} \oplus \mathcal{L})$. So in fact, the $P = W$ conjecture tells us that all of these enumerative invariants of the local Calabi-Yau threefold can be obtained from the Hodge numbers of Deligne's weight filtration on the cohomology of $\mathcal{M}_B(r, d)$. Quite a remarkable result, because the perverse weight filtration is (in general) not at all well-understood, where the weight filtration on an affine variety is well-understood. For more details, we refer to [88, 69].

With this in mind, it is interesting to consider whether the $P = W$ phenomenon appears in more generality. For instance, suppose we have a Calabi-Yau threefold which is fibred by $K3$ surfaces, and looks like $\mathcal{S} \times \mathbb{C}$ away from the singular fibres. Here, \mathcal{S} is a $K3$ surface. Then one can also consider the Gopakumar-Vafa invariants of the local Calabi-Yau threefold $\mathcal{S} \times \mathbb{C}$. To compute these, one needs to consider the moduli space of curves with line bundles on them for a given $K3$ surface. Mathematically, this moduli space is constructed as the moduli space of 1-dimensional stable sheaves \mathcal{E} on \mathcal{S} such that $\text{Supp}(\mathcal{E}) = \beta \in H_2(\mathcal{S}, \mathbb{Z})$ and $\chi(\mathcal{E}) = 1$, where β is the homology class of the curve in $\mathcal{S} \subset X$. Denote this moduli space by \mathcal{M}_{β} . Then there is a morphism $\pi : \mathcal{M}_{\beta} \rightarrow \text{Chow}_{\beta}(\mathcal{S})$, the Chow variety which parameterises algebraic cycles of dimension 1 and homology class β in \mathcal{S} . Naturally, π sends the sheaf \mathcal{E} to its support. It turns out (see [89]) that for a $K3$ surface, this map is induced by a morphism

$$\mathcal{S}^{[n]} \rightarrow \mathbb{C}\mathbb{P}^n$$

with $n = \frac{1}{2}\beta^2 + 1$ (recall that the $K3$ lattice is even, so β^2 is always even) and $\mathcal{S}^{[n]}$ the Hilbert scheme of n -points on \mathcal{S} . To extract the Gopakumar-Vafa invariants, one needs to compute the perverse Hodge numbers associated to this morphism π . In general, this is not a simple task, so it would be much easier if the $P = W$ phenomenon could be used for $\mathcal{S}^{[n]}$. The main result of [89] is that this can indeed be done.

7.3. ENUMERATIVE GEOMETRY AND $P = W$

Theorem 7.3.1. [89] *For any projective irreducible holomorphic symplectic variety \mathcal{M} , equipped with a holomorphic Lagrangian fibration $\pi : \mathcal{M} \rightarrow \mathcal{B}$, it holds that*

$$h_{\mathfrak{F}}^{p,q}(\mathcal{M}) = h^{p,q}(\mathcal{M})$$

In this way, the various enumerative invariants (the local Gromov-Witten, Pandharipande-Thomas and Donaldson-Thomas invariants) of $\mathcal{S} \times \mathbb{C}$ can be deduced through the $P = W$ phenomenon and the Göttsche formula [90], which states

$$\sum_{n=1}^{\infty} \mathfrak{W}_{\mathcal{S}^{[n]}}(1, u, v) t^n = \prod_{k=1}^{\infty} \prod_{p,q=0}^2 (1 + (-1)^{p+q+1} u^{p+k+1} v^{q+k+1} t^k)^{(-1)^{p+q+1} h^{p,q}(\mathcal{S})}$$

This holds for any projective complex surface. Note that the formula on the right hand side depends only on the Hodge numbers of the surface \mathcal{S} , so the Hodge numbers of $\mathcal{S}^{[n]}$ can be determined explicitly. In the case of a $K3$ surface, one finds

$$\begin{aligned} \mathfrak{W}_{\text{tot}} &= \sum_n \mathfrak{W}_{\mathcal{S}^{[n]}} t^n = \\ & \prod_{k=1}^{\infty} \left((1 - u^{k+1} v^{k+1} t^k) (1 - u^{k+3} t^k) (1 - u^{k+2} v^{k+2} t^k)^{20} (1 - v^{k+3} t^k) (1 - u^{k+3} v^{k+3} t^k) \right)^{-1} \end{aligned}$$

Running a program, one can then obtain $h^{p,q}(\mathcal{S}^{[n]}) = h_{\mathfrak{F}}^{p,q}(\mathcal{S}^{[n]})$ as

$$\frac{1}{n!p!q!} \left(\frac{\partial^{n+p+q} \mathfrak{W}_{\text{tot}}}{\partial t^n \partial u^p \partial v^q} \right) (0, 0, 0)$$

Chapter 8

The Thomas-Yau Conjecture

We are now almost ready to discuss the Thomas-Yau conjecture, although we need one more piece of the puzzle. This piece comes from the realm of string theory, and is related to the notion of stability and BPS (Bogomolnyi-Prasad-Sommerfield) branes. These are certain distinguished branes in the Calabi-Yau target space of the supersymmetric non-linear open string sigma model which preserve some supersymmetry, and so they are also called supersymmetric branes. We will give a very brief overview of these BPS branes, and then start to piece together the puzzle.

8.1 BPS Branes

We refer the reader to [55] and [5] for extensive discussions on the topic that we are about to discuss. As before, we will assume in the following that we are at the large volume limit. This means that the volume of the Kähler form tends to infinity. In this case, we may view A -branes as objects in the Fukaya category, and B -branes as coherent sheaves.

We recall that the homological mirror symmetry conjecture pertains to the category of A -branes and B -branes, which are the D -branes of the A -model and B -model, respectively. In turn, these models were topological twistings of some underlying string theory, the "untwisted" theory. Recall that a D -brane in the untwisted theory with Calabi-Yau target space (\mathcal{X}, g) is (simplistically) viewed as a submanifold $Y \subset X$ together with a vector bundle $E \rightarrow Y$ and a connection ∇ on E . Some branes from the twisted theory arise as branes in the untwisted theories, namely those which correspond to so-called BPS branes. If an open string ends on two BPS branes, this means that its equations of motion are invariant under half of the supersymmetry, which is the maximal amount of supersymmetry that can be preserved for open strings, and are therefore of great interest to string theorists, e.g. [91]. Supersymmetry demands that the submanifolds corresponding to these branes are volume minimising with respect to the metric on the Calabi-Yau manifold (\mathcal{X}, g) , and the existence of covariantly constant spinors on the submanifolds. Mathematically, this translates into the condition that the submanifolds need to be calibrated (see 2.4.2 where we discussed calibrations on Calabi-Yau manifolds).

On a Calabi-Yau manifold (\mathcal{X}, g) which is the target space of our theory, a calibrated submanifold can either be a complex submanifold $\mathcal{Y} \subseteq \mathcal{X}$ with respect to ω^k , or a special Lagrangian submanifold $L \subseteq X$ with respect to $\text{Re} \exp(i\theta)\Omega$. These arise as the submanifolds for BPS branes in type IIA, respectively type IIB string theory. The terminology

8.1. BPS BRANES

is very confusing for us, since BPS branes for type IIA string theory resemble B -branes while BPS branes for type IIB string theory resemble A -branes. This is why we will continue to avoid mentioning type IIA/IIB string theory.

There are further restrictions which are imposed on the triple (Y, E, ∇) , for it to correspond to BPS a brane in the untwisted theory. In each case, this additional condition comes in the form of a partial differential equation that must be satisfied. The first type of BPS branes consist of

1. A complex submanifold $\mathcal{Y} \subseteq \mathcal{X}$.
2. A connection ∇ on $E \rightarrow Y$ which satisfies the HYM equation.

We call these B -type BPS branes. The second type of BPS branes consist of

1. A special Lagrangian submanifold $L \subseteq X$, i.e. $\text{Im } \exp(i\theta)\Omega|_L = 0$ and $\omega|_L = 0$
2. A Hermitian line bundle $E \rightarrow L$ with unitary connection ∇ and $F_\nabla = 0$.

These are called A -type BPS branes. Evidently, every special Lagrangian can be graded, and defines an object in the derived Fukaya category. Similarly, a B -type BPS brane defines an object in $\text{Coh}(\mathcal{X})$. But in both cases, some additional criteria are satisfied.

Recall that we discussed the SYZ picture of mirror symmetry in 4.4.2. As mentioned there, the SYZ approach to mirror symmetry, particularly in the semi-flat limit (when the metric is flat along the torus fibres), provides a nice testing ground for certain ideas. For instance, in [67], the authors show that the smooth Fourier-Mukai transform (also called T -duality transform) of a special Lagrangian fibration, takes A -type BPS branes on \mathcal{X} to B -type BPS branes on \mathcal{X}^\vee , whenever the Lagrangian submanifold underlying the A -brane can be expressed as a (multi-)section of the fibration. Whilst not conclusive by any means, it is evidence in favour of homological mirror symmetry. But the condition for a B -brane to be BPS actually depends on the the Kähler form, since the Hermitian-Yang-Mills equation does. Conversely, the special Lagrangian criterion for A -type branes requires the presence of the holomorphic volume form. So the citation is evidence in favour of more than homological mirror symmetry - it is evidence in favour of mirror symmetry at the level of the untwisted string theory, which is what string theorists believe to be true (sometimes referred to as quantum mirror symmetry):

Conjecture 8.1.1. *The set of A -type BPS branes on \mathcal{X} is isomorphic to the set of B -type BPS branes on \mathcal{X}^\vee .*

This is a consequence of the isomorphism (whatever this means is not mathematically rigorous) between type IIB string theory on \mathcal{X} , and type IIA string theory on \mathcal{X}^\vee . By the Kobayashi-Hitchin correspondence, B -type BPS branes are precisely complex submanifolds $\mathcal{Y} \subseteq \mathcal{X}$ together with a μ -stable holomorphic vector bundle, which give μ -stable coherent sheaves on \mathcal{X} via pushforward. So in fact, we will view the set of B -type BPS branes as the set of μ -stable coherent sheaves on \mathcal{X} . Thus, B -type BPS branes are intricately tied to some notion of stability, at least at the large volume limit. Is the

8.2. THE MIRROR IMAGE

same true of A -type BPS branes? This is the question that the Thomas-Yau conjecture attempts to answer, as we discuss next.

A more rigorous notion of stability was introduced in the string theory literature under the name Π -stability, see [6, 92], which applies away from the large volume limit as well. This was adapted to the framework of triangulated categories by Bridgeland, see [93]. The Thomas-Yau conjecture was reinterpreted in this context by Joyce, see [12]. We will discuss this later on.

Remark 8.1.2. Note that the Donaldson-Thomas invariants of (\mathcal{X}, g) can be interpreted as the count of BPS branes for the large volume limit of type IIA string theory compactified on (\mathcal{X}, g) .

8.2 The Mirror Image

The Thomas-Yau conjecture has its origins in a paper written by Thomas [1], where he looks for an analogue of the notion of μ -stability for A -branes. Mirror symmetry is used as a guiding principle, and he presents a table which provides the dictionary between the A -side and the B -side of the story. The entries that we have not yet discussed will be explained shortly. We fix some basepoint ∇_0 in the affine space of connections on a vector bundle E , and we denote by $\bar{\partial}_0$ its $(0, 1)$ part.

Mirror Symmetry Dictionary	
$(\mathcal{X}, \omega, \Omega)$	$(\mathcal{X}^\vee, \omega^\vee, \Omega^\vee)$
$\Omega \in H^{3,0}(\mathcal{X})$	$\omega^\vee \in H^{1,1}(\mathcal{X}^\vee)$
Connections on E with charge $\eta = \text{ch}(E)\sqrt{\text{td}(\bar{X})} \in H^{\text{ev}}(X, \mathbb{Q})$	Submanifolds L in homology class $[L] \in H_3(X^\vee)$ with connection on $L \times \mathbb{C}$
$S_{\mathbb{C}}(\bar{\partial} = \bar{\partial}_0 + A) = \int_X \text{tr}(A \wedge \bar{\partial}_0 A + \frac{2}{3} A^3) \wedge \Omega$	$f_{\mathbb{C}}(A, L) = \int_{L_0}^L (F + \omega^\vee)^2$
Critical points: holomorphic bundles	Critical points: Lagrangians with flat line bundles
Invariant: $\text{DT}_\eta(\mathcal{X})$	Invariant: "counting" special Lagrangian submanifolds
Symmetry group: bundle automorphisms	Symmetry group: Hamiltonian deformations
$\omega \in H^{1,1}(\mathcal{X})$	$\Omega^\vee \in H^{3,0}(\mathcal{X}^\vee)$
Momentum map: $A \mapsto F_A \wedge \omega^2$	Momentum map: $(A, L) \mapsto \text{Im } \Omega^\vee _L$
Slope: $\mu(E) = \frac{1}{\text{rank}(E)} \int_X c_1(E) \wedge \omega^2$	Slope: $\mu(L) = \frac{1}{\text{Vol}(L)} \int_L \text{Im } \Omega^\vee$

Note that, aside from the functionals $S_{\mathbb{C}}$ and $f_{\mathbb{C}}$ and the associated invariants, the story actually applies to any Calabi-Yau manifold. Next we will explain this table in more detail, and how it leads to the conjecture by Thomas in [1].

We begin on the left hand side of the table, which concerns holomorphic vector bundles on a Calabi-Yau threefold \mathcal{X} . Suppose we fix a smooth vector bundle with the topo-

8.2. THE MIRROR IMAGE

logical data of the charge vector $\text{ch}(E)\sqrt{\text{td}(X)}$.¹ Endow it with a Hermitian metric, and consider the space of unitary connections on E . Equivalently, consider the space of $\bar{\partial}$ -operators on E compatible with the metric. We can define a functional on this infinite dimensional space, the holomorphic Chern-Simons functional $S_{\mathbb{C}}(A)$. Its critical points correspond to connections whose curvature satisfies $F_{\nabla}^{0,2} = 0$, i.e. to holomorphic structures on E . From this, we learn that the moduli space of holomorphic vector bundles of a given topological type has virtual dimension zero, and we may define the Donaldson-Thomas invariants $\text{DT}_{\eta}(\mathcal{X})$. This may also be viewed as the process of symplectic reduction, by taking the bundle automorphisms \mathcal{G} to act on the space of unitary connections (with curvature of type $(1,1)$), denoted $\mathcal{A}^{(1,1)}$. We can give a momentum map for this symplectic group action using the Kähler form ω , namely $A \mapsto F \wedge \omega^2$. The Kobayashi-Hitchin correspondence tells us that this is equivalent to the formal GIT quotient $\mathcal{A}^{(1,1)}/\mathcal{G}_{\mathbb{C}}$, where the notion of stability is defined by the slope $\mu(E)$. That is, the moduli space of stable holomorphic vector bundles coincides with the symplectic quotient $\mathcal{A}^{(1,1)} // \mathcal{G}$, and the orbit spaces of the latter may be represented by a unique HYM connection, whereas points in the GIT quotient correspond to orbits of (semi-)stable points in $\mathcal{A}^{(1,1)}$. In other words, solutions to some PDE imply a certain well-behavedness on the algebraic side. This is a retelling of the story we discussed in 6.3.

So can a similar theory be developed for the A -type branes? Answering this question is the goal of [1]. To start with, we should find some infinite dimensional space which plays the role of \mathcal{A} . The general principles behind this are outlined in [94]. Denote by \mathcal{M} the space whose elements consist of pairs (L, A) where L is a submanifold $L \subseteq X$ in a fixed homology class $[L] \in H_3(X, \mathbb{Z})$ and $\nabla = d + A$ is a unitary connection on $L \times \mathbb{C}$. We are interested in the case where L is a Lagrangian submanifold. In this case, the complex structure on \mathcal{X} allows us to identify normal vectors with tangent vectors, i.e. $TL \cong NL$. Infinitesimal deformations of L can be viewed as normal vector fields on L , and using the complex structure and the metric, these can be represented by elements of $\Omega^1(L)$. On the other hand, infinitesimal deformations of the unitary connection are given by 1-forms $\Omega^1(L, \mathfrak{u}(1)) \cong i\Omega^1(L)$, since this is the tangent space to the space of connections. As such, we see that the tangent space of \mathcal{M} at a Lagrangian submanifold may be identified with $\Omega^1(L) \oplus i\Omega^1(L) = \Omega^1(L) \otimes \mathbb{C}$, so that $T_{(L,A)}\mathcal{M}$ acquires a natural complex structure, at least when L is Lagrangian. Since \mathcal{M} is not an affine space, we cannot guarantee that this almost complex structure is integrable by some elementary argument. Let us proceed under the assumption that it is, for the purpose of developing an analogy with the B -type branes.

The next step in doing so is to define a functional $f_{\mathbb{C}}$ on \mathcal{M} whose critical point yield A -branes, i.e. Lagrangian submanifolds with flat line bundles. In the same way that we fixed a reference connection ∇_0 for the holomorphic Chern-Simons functional, we now

¹Recall that the Chern character is an isomorphism $K(X) \otimes_{\mathbb{Z}} \mathbb{Q} \rightarrow H^{\text{ev}}(X, \mathbb{Q})$, where $K(X)$ denotes the topological K -theory group which consists of isomorphism classes of topological vector bundles. So up to torsion, a vector bundle is determined by its Chern class, and $\sqrt{\text{Td}(X)}$ is just a change of basis.

8.2. THE MIRROR IMAGE

fix $L_0 \in \mathcal{M}$ to represent the homology class of a given $L \in \mathcal{M}$. Then we can define

$$f_{\mathbb{C}}(L, A) := \int_{L_0}^L (F + \omega)^2$$

where the integration is over a 4-cycle whose boundary is $L - L_0$. Such a cycle is guaranteed to exist since $[L] = [L_0]$. The curvature form F is that of a connection on a bundle on the 4-cycle, which restricts to A on L , and to A_0 on L_0 . Therefore, $f_{\mathbb{C}}(L, A)$ is only well-defined up to addition of some period. Evidently, critical points correspond to pairs (L, A) such that $\omega|_L = 0$ and $F = 0$, i.e. Lagrangian submanifolds together with flat connections. Like the Chern-Simons functional, the functional $f_{\mathbb{C}}$ is holomorphic with respect to the complex structure on \mathcal{M} . We denote its space of critical points by \mathcal{M}^0 . That is,

$$\mathcal{M}^0 = \{(L, A) \mid [L] = [L_0] \in H_3(X, \mathbb{Z}), \omega|_L = 0, F_A = 0\}$$

Next in the table is some enumerative invariant, which should be mirror to the Donaldson-Thomas invariant. In other words, it should count A -type BPS branes. A first attempt at defining this invariant was made by Joyce in [95], in which the proposal is to count special Lagrangian homology spheres, together with gauge equivalence classes of flat line bundles. This is because special Lagrangian homology spheres are isolated, so they have a 0-dimensional moduli space. Note that the publication of this paper (in 1999) predates much of the development of D -branes in string theory, so this was not thought of as an A -side invariant, although Joyce speculates that his invariants are related to Gromov-Witten invariants. Either way, the attempt was inconclusive.

So the "right" definition of the A -side invariant has not yet been given. Perhaps the method of symplectic reduction can give us a hint as to what this invariant should be, through a similar infinite dimensional Kempf-Ness picture as the Kobayashi-Hitchin correspondence. First, we can try to find the symmetry group based on string theory arguments. On the B -side, two holomorphic vector bundles are isomorphic as branes if they are related by a unitary gauge transformation, which is why we obtained \mathcal{G} as a symmetry group. On the A -side, two branes are isomorphic if they are related by a Hamiltonian deformation, since the Floer cohomology is invariant under Hamiltonian deformations in both arguments. Can these act by symplectomorphisms with respect to an appropriately defined symplectic structure on \mathcal{M}^0 ? To answer this, we first need to understand what $T_L \mathcal{M}^0 \subseteq \Omega^1(L, \mathbb{C})$ looks like.

We recall that we obtained $\Omega^1(L, \mathbb{C})$ as $\Omega^1(L, \mathbb{R}) \oplus \Omega^1(L, \mathfrak{u}(1))$. These were infinitesimal deformations of the submanifold L and the connection $\nabla = d + A$, respectively. Now, we want L to remain Lagrangian under the infinitesimal deformation, and we want ∇ to remain flat. For the former, we require symplectic vector fields², i.e. vector fields v such that $\mathcal{L}_v \omega = 0$. Equivalently, this means $d \circ \iota_v \omega = 0$. As a result, the subspace in $\Omega^1(L, \mathbb{R})$ which corresponds to Lagrangian deformations is $\ker(d : \Omega^1(L, \mathbb{R}) \rightarrow \Omega^2(L, \mathbb{R}))$, which we

²We are still considering Calabi-Yau manifolds in the strict sense. Thus, $H^1(X, \mathbb{R}) = 0$, which means that every symplectic vector field is in fact Hamiltonian.

8.2. THE MIRROR IMAGE

write as $Z^1(L, \mathbb{R})$. Similarly, suppose that $\nabla = d + A \mapsto \nabla_t = \nabla + ta$ for $a \in \Omega^1(L, \mathfrak{u}(1))$. Then the curvature of ∇_t is given by

$$F_A + t(da + [A, a]) + \frac{t^2}{2}[a, a] = F_A + tda$$

since $U(1)$ is abelian. Differentiating, we find that for ∇_t to remain flat, we need $da = 0$, i.e. $a \in Z^1(L, \mathfrak{u}(1))$. We conclude that $T_L \mathcal{M}^0 = Z^1(L, \mathbb{C})$. Now we can attempt to define a symplectic form on \mathcal{M}^0 . By analogy with the B -side, we would like to use the 3-form $\Omega|_L \in \Omega^3(L, \mathbb{C})$ to define this symplectic form. We use this form to define a metric on \mathcal{M}^0 , which combines with the complex structure to yield a symplectic form. Fix a homology class for L and normalise Ω so that $\int_L \text{Im } \Omega = 0$. For the metric, we can then take

$$\langle \cdot, \cdot \rangle : T_L \mathcal{M}^0 \times T_L \mathcal{M}^0 \rightarrow \mathbb{R} \quad ((a_1, a_2), (b_1, b_2)) \mapsto \int_L a_1 \wedge (\iota_{\tilde{b}_1} \text{Im } \Omega|_L) + \int_L a_2 \wedge (\iota_{\tilde{b}_2} \text{Im } \Omega|_L)$$

where $\tilde{b} \in \Gamma(L, TL)$ is obtained from $b \in Z^1(L, \mathbb{R})$ from the isomorphism $T^*X|_L \rightarrow TX|_L$ provided by ω , and the isomorphism $NL \cong TL$ provided by J and the fact that L is Lagrangian. As written, it is not so clear that this should define a metric. However, if we consider $\Omega|_L = \exp(i\theta)d\text{Vol}_g$ then the expression may be written as

$$((a_1, a_2), (b_1, b_2)) \mapsto \int_L \cos(\theta) a_1 \wedge \star b_1 + \int_L \cos(\theta) a_2 \wedge \star b_2$$

We are already restricting our attention to graded Lagrangians, but for this to define an inner product, we cannot allow $\theta = \pm\pi/2$. As such, we restrict our attention to almost calibrated Lagrangians, which means that $\theta(x) \in (-\pi/2, \pi/2)$ for all $x \in L$. We redefine \mathcal{M}^0 to consist of pairs (L, A) where L is an almost calibrated Lagrangian submanifold, together with a flat connection.

Proposition 8.2.1. *The map $\langle \cdot, \cdot \rangle : T_L \mathcal{M}^0 \times T_L \mathcal{M}^0 \rightarrow \mathbb{R}$ defines a Hermitian metric on \mathcal{M}^0 .*

Proof. Let $(L, A) \in \mathcal{M}^0$. Since L is almost calibrated, the function $\cos(\theta) \in C^\infty(L)$ is nowhere vanishing. As such, $\int_L \cos(\theta) a \wedge \star b$ is just a deformed version of the Hodge inner product on $\Omega^1(L, \mathbb{R})$. Then $\langle \cdot, \cdot \rangle$ is simply the induced inner product on the direct sum

$$T_A \mathcal{M}^0 = Z^1(L, \mathbb{C}) \cong Z^1(L, \mathbb{R}) \oplus Z^1(L, \mathbb{R})$$

Next, we must show that this metric is compatible with the complex structure. Writing the tangent space as $T_L \mathcal{M}^0 = \Omega^1(L, \mathbb{R}) \oplus \Omega^1(L, \mathbb{R})$, the complex structure acts as the matrix

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} : \Omega^1(L, \mathbb{R}) \oplus \Omega^1(L, \mathbb{R}) \rightarrow \Omega^1(L, \mathbb{R}) \oplus \Omega^1(L, \mathbb{R})$$

As such, we find that

$$\begin{aligned} \langle J(a_1, a_2), J(b_1, b_2) \rangle &= \langle (-a_2, a_1), (-b_2, a_2) \rangle = \int_L \cos(\theta) (-a_2) \wedge \star (-b_2) + \int_L \cos(\theta) a_1 \wedge \star b_1 = \\ &= \int_L \cos(\theta) a_1 \wedge \star b_2 + \int_L \cos(\theta) a_2 \wedge \star b_2 = \langle (a_1, a_2), (b_1, b_2) \rangle \end{aligned}$$

We conclude that this is indeed a Hermitian metric. \square

8.2. THE MIRROR IMAGE

Since we now have a Hermitian manifold $(\mathcal{M}^0, J, \langle \cdot, \cdot \rangle)$, we also have a symplectic structure by taking $\vartheta_L(a, b) = \langle a, Jb \rangle$, assuming that the Hermitian metric is Kähler.³

Definition 8.2.2. Let (X, ω) be a symplectic manifold. A Hamiltonian deformation of X is a diffeomorphism $\varphi : X \rightarrow X$ which is generated by a Hamiltonian vector field. Let $L \subseteq X$ be a Lagrangian submanifold. Then a Hamiltonian deformation of L is a deformation of L which is pulled back from a Hamiltonian deformation $\varphi : X \rightarrow X$.

For strict Calabi-Yau manifolds, one has $H^1(X, \mathbb{R}) = 0$ so every symplectic vector field on X is Hamiltonian. The group of Hamiltonian deformations of L acts on \mathcal{M}^0 by $\varphi : (L, \nabla) \mapsto (\varphi(L), (\varphi^{-1})^* \nabla)$. It combines with the group of gauge transformations into a group action of the semidirect product $\mathcal{H} := \text{Ham}(L, \omega) \ltimes \mathcal{G}$.

Proposition 8.2.3. *The group \mathcal{H} acts on $(\mathcal{M}^0, \vartheta)$ by symplectomorphisms.*

Proof. This follows straightforwardly from the invariance of ϑ under the gauge group part of the action, and the diffeomorphism invariance of the integral. That is,

$$\int_L \eta = \int_{\varphi(L)} (\varphi^{-1})^* \eta$$

Together with the naturality of the Hodge star operator, this yields the result. \square

Given that we have a group acting by symplectomorphisms, we would now like to find a momentum map. First, then, we must identify the dual of the Lie algebra of \mathcal{H} . In fact, since the gauge transformations do not affect the symplectic form, we need only know the dual of the Lie algebra of the group of Hamiltonian deformations of L . The Lie algebra of this group consists of closed 1-forms, since every Hamiltonian deformation is symplectic when $H^1(X, \mathbb{R}) = 0$. Thus, we may identify this vector space with $V := C^\infty(L)/\mathbb{R}$, since every closed 1-form is exact.

Proposition 8.2.4. *Let $\Omega_0^3(L)$ denote those 3-forms η such that $\int_L \eta = 0$. Then we identify $\Omega_0^3(L)$ with V^* as vector spaces.*

Note that we do not require this to be a morphism of Lie algebras. We are content with being able to identify elements in V^* with some other vector space (or the converse).

Proof. We already have $\Omega^0(L) \times \Omega^3(L) \rightarrow \mathbb{R}$ which is given by integration. We used this to make the identification $\Omega^3(L) \cong C^\infty(L)^*$. We may normalise this isomorphism so that $d\text{Vol}$ is mapped to 1. In this case, taking the quotient $C^\infty(L)/\mathbb{R}$ amounts to taking $\Omega^3(L)/(\mathbb{R} \cdot d\text{Vol})$. We identify the latter with $\Omega_0^3(L)$, by taking the unique representative of each equivalence class which integrates to 0. \square

³If we instead model the space \mathcal{M} on maps of some fixed compact manifold into X , then one can certainly construct a Kähler metric. However, not all homologous Lagrangian submanifolds arise in this way.

8.2. THE MIRROR IMAGE

Thus, a momentum map for the action of \mathcal{H} on \mathcal{M}^0 will be a map $\mathcal{M}^0 \rightarrow \Omega_0^3(L)$. This notation is somewhat imprecise since the submanifold L varies as we move within \mathcal{M}^0 . We would really like to identify it with $C^\infty(L_0)/\mathbb{R}$ for some fixed L_0 in the homology class we have chosen. However, this causes some issues which are also discussed in [1]. We will ignore these rather subtle issues and continue to outline the main ideas.

Just as we used the Kähler form to construct a momentum map for the B -side story, here we want to use the holomorphic volume form. Recall that we normalised Ω so that $\int_L \text{Im } \Omega = 0$.

Proposition 8.2.5. *A momentum map for the action of \mathcal{H} on \mathcal{M}^0 is given by $L \mapsto \text{Im } \Omega|_L$.*

Proof. The moment map μ pairs with the Lie algebra $C^\infty(L)/\mathbb{R}$ via $\langle \mu, h \rangle = \int_L h \text{Im } \Omega|_L$, for $h \in C^\infty(L)/\mathbb{R}$. We will take a tangent vector to \mathcal{M}^0 to consist of a single closed 1-form a , even though the tangent space is $Z^1(L) \oplus Z^1(L)$. This is because one of these summands represents infinitesimal deformations of the connection, which will not be affected. So we pick a 1-form $a \in Z^1(L)$ representing an infinitesimal deformation of L . That is, it corresponds to a normal vector field $v \in \Gamma(L, NL)$ via the Kähler form. The derivative of $\langle \mu, h \rangle$ in the direction of v is just

$$\int_L \mathcal{L}_v(h \text{Im } \Omega|_L)$$

where \mathcal{L} is the Lie derivative. We extend h to an open neighbourhood $L \subset U \subseteq X$ such that h is constant along the flow of v . This is possible since v is normal to L . As a result, we get

$$\int_L \mathcal{L}_v(h \text{Im } \Omega|_L) = \int_L 0 \cdot \text{Im } \Omega|_L + \int_L h \mathcal{L}_v \text{Im } \Omega|_L = \int_L h(d \circ \iota_v + \iota_v \circ d) \text{Im } \Omega|_L = \int_L h(d \circ \iota_v \text{Im } \Omega|_L)$$

using Cartan's formula and the fact that Ω is closed. Integrating by parts, this is just

$$- \int_L dh \wedge \iota_v \text{Im } \Omega|_L = \vartheta(dh, a) = (\iota_{\rho(h)} \vartheta)(a)$$

□

So now, we can once again carry out formal symplectic reduction. Can we complexify the group \mathcal{H} to get a group action of $\mathcal{H}_\mathbb{C}$ on \mathcal{M}^0 ? It is clear that the complexified Lie algebra should be $C^\infty(L, \mathbb{C})/\mathbb{C}$. However, there appears to be no natural way to complexify the group action, and so the analogy with a formal GIT quotient or formal Kempf-Ness theorem seems to break down here. Nevertheless, we can perform the formal symplectic reduction and find that $\mathcal{M}^0 // \mathcal{H}$ is the moduli space of special Lagrangian submanifolds together with flat $U(1)$ -connections, since $\text{Im } \Omega|_L = 0$ is precisely the special Lagrangian condition for the set \mathcal{M}^0 (whose elements are already Lagrangian).

At the infinitesimal level, the analogy is also very tantalising. Mirror symmetry relates the Ext^\bullet -groups to the Lagrangian-Floer cohomology groups. Presently, $\text{Ext}^1(\mathcal{E}, \mathcal{E}) \cong$

8.3. WALL-CROSSING

$H^1(\mathcal{X}, \text{End}(\mathcal{E}))$ and $HF^1(L, L) \cong H^1(L, \mathbb{C})$ (modulo the technicality of taking coefficients in the Novikov field). These are precisely the tangent spaces of the respectively moduli spaces, which is another hint that there may be truth behind the analogy presented above.

The B -side story had a very nice feature: the Kobayashi-Hitchin correspondence. Unfortunately, the lack of a formal GIT quotient picture for the A -side story does not allow for a direct analogy with this. However, an initial version of the Thomas-Yau conjecture may now be informally understood as follows (meaning: this is not how Thomas formulated the conjecture).

Conjecture 8.2.6 (Thomas). *There exists a condition on \mathcal{M}^0 , say a stability condition, such that (L, A) is stable if and only if (L, A) has a unique special Lagrangian representative in its orbit under \mathcal{H} .*

The stability condition in the above conjecture should play the same role as the stability of holomorphic vector bundles, which ensures that a given connection ∇ on E has a unique HYM representative in its gauge orbit. In fact, Thomas proposed a stability condition on \mathcal{M}^0 , which leads to the actual conjecture as it is presented in [1]. Before we get there, we first discuss the evidence for the conjecture presented in loc. cit., which will also introduce the necessary definitions that are needed to understand Thomas's initial definition of stability for A -branes.

8.3 Wall-Crossing

We return to the B -side picture of holomorphic bundles on a Calabi-Yau threefold (\mathcal{X}, g) . From both physical considerations (BPS states) and mathematical considerations (enumerative invariants), it is interesting to consider the moduli space of semi-stable coherent sheaves on \mathcal{X} of some fixed Chern character. Looking at the smooth manifold X , there are many choices of complex structure (provided X is not rigid), and these change which sheaves are the coherent ones. Furthermore, we can let the Kähler form vary, and this would change which coherent sheaves are stable, since stability is defined using the Kähler class. In his thesis, Thomas proves that the invariant which bears his name is in fact invariant under the former.

Theorem 8.3.1. [14] *The Donaldson-Thomas invariant $DT_\eta(\mathcal{X})$ is invariant under complex deformations of \mathcal{X} .*

This result justifies the moniker "invariant". However, the Kähler deformations do not have this convenient property. Consider some fixed coherent sheaf \mathcal{E} over \mathcal{X} . It may happen that we choose a 1-parameter family of Kähler forms $\omega_t : \mathbb{R} \rightarrow \Omega^2(X)$ such that \mathcal{E} is stable if $t > 0$, semi-stable if $t = 0$ and unstable if $t < 0$.

8.3. WALL-CROSSING

Example 30. [1] Suppose we have two coherent sheaves \mathcal{E}_i with $\text{Ext}^1(\mathcal{E}_2, \mathcal{E}_1) \cong \mathbb{C}$. Then there exists a unique non-trivial extension (up to scaling), since these are classified by $\text{Ext}^1(\mathcal{E}_2, \mathcal{E}_1)$. So define \mathcal{E} by

$$0 \rightarrow \mathcal{E}_1 \rightarrow \mathcal{E} \rightarrow \mathcal{E}_2 \rightarrow 0$$

Suppose furthermore that we can choose a 1-parameter family ω_t of Kähler forms with the property that $\Delta_t := \mu_t(\mathcal{E}_2) - \mu_t(\mathcal{E}_1) : \mathbb{R} \rightarrow \mathbb{Q}$ is such that the sign of Δ_t agrees with the sign of t for all $t \in (-\varepsilon, \varepsilon)$.⁴ Furthermore, suppose that the \mathcal{E}_i are stable w.r.t. ω_t for $t \in (-\varepsilon, \varepsilon)$. Then \mathcal{E} will be stable w.r.t. ω_t for sufficiently small $t > 0$, but unstable for $t < 0$. To see this, we fix $\mu_t(\mathcal{E}_2) := \mu$, and let $\mu_t(\mathcal{E}_1) = \mu - t$, for convenience. Because of the injection $\mathcal{E}_1 \hookrightarrow \mathcal{E}$, it is then immediate that \mathcal{E} is no longer stable for $t \leq 0$, because $\mu_{t \leq 0}(\mathcal{E}_1) \geq \mu_{t \leq 0}(\mathcal{E})$. However, let $t > 0$ be sufficiently small. Stability of \mathcal{E}_2 implies there are no subsheaves $\mathcal{F} \subset \mathcal{E}_2$ with $\mu_t(\mathcal{F}) > \mu - t$. Suppose that $\mathcal{F} \subset \mathcal{E}$ is a stable destabilising subsheaf, meaning $\mu_t(\mathcal{F}) \geq \mu_t(\mathcal{E})$, but \mathcal{F} itself is stable. Then the composition $\mathcal{F} \hookrightarrow \mathcal{E} \rightarrow \mathcal{E}_2$ cannot be an injection, unless it is an isomorphism. Since the extension above does not split by virtue of the fact that the extension is defined by a non-trivial element of $\text{Ext}^1(\mathcal{E}_2, \mathcal{E}_1)$, it cannot be an isomorphism. We conclude that $\mathcal{F} \cap \mathcal{E}_1 \neq 0$. The quotient does define an injection $\mathcal{F}/(\mathcal{F} \cap \mathcal{E}_1) \hookrightarrow \mathcal{E}_2$. We get $\mu_t(\mathcal{F}/(\mathcal{F} \cap \mathcal{E}_1)) > \mu_t(\mathcal{F}) > \mu - t$, following from the stability of \mathcal{F} . This is a contradiction, since \mathcal{E}_2 is stable. Therefore, \mathcal{E} must be stable.

The example above demonstrates the so-called wall-crossing phenomenon. By varying the stability condition, stable objects may simply disappear when we cross certain walls in the space of stability conditions. This can be made more precise, as we will see later, using the notion of Bridgeland stability. The wall-crossing phenomenon is an active area of research within the theory of Donaldson-Thomas invariants, see e.g. [96] or [97].⁵

This is perhaps the first case in which the B -side picture is more difficult to describe than the A -side picture. Before we outline how, we recall that the phase of a graded, oriented Lagrangian submanifold $L \subset X$ is defined as the function $\theta : L \rightarrow \mathbb{R}$ such that

$$\Omega|_L = \exp(i\theta) d\text{Vol}_g|_L$$

We also define the average phase $\phi(L) \in \mathbb{R}$ as

$$\int_L \Omega|_L = A \exp(i\phi(L))$$

⁵It should be noted that these articles investigate the wall-crossing phenomenon for stability conditions in the sense of Bridgeland. The stability of sheaves does not define a Bridgeland stability condition on Calabi-Yau 3-folds. Nevertheless, there is a close relation.

8.3. WALL-CROSSING

This is defined up to shifts $\phi(L) \mapsto \phi(L) + 2\pi n$ for integers n . The slope of L , which will be denoted $\mu(L)$ by analogy with the slope of a coherent sheaf, is defined as

$$\mu(L) = \frac{\int_L \operatorname{Im} \Omega}{\int_L \operatorname{Re} \Omega} = \tan(\phi(L))$$

The slope and the phase capture similar information, but the slope is invariant under a change of orientation of L . This may seem like somewhat of a deficiency if we wish to develop a close analogy to holomorphic vector bundles. However, since we are considering vector bundles on the B -side, rather than objects in the derived category, we have thrown away the grading of the complexes in the derived category. As such, it is no wonder that we lost some information by naively considering bundles only.

The starting point for wall-crossings on the A -side is the notion of a Lagrangian connected sum. We consider two Lagrangian submanifolds $L_0, L_1 \subset X$ such that $L_0 \cap L_1 = \{x_0\}$. In this case, one can use the Lagrangian connected sum $L_0 \natural L_1$. The notation is meant to reflect the fact that this construction is asymmetric in its arguments: $L_0 \natural L_1 \not\approx L_1 \natural L_0$, where \approx denotes being Hamiltonian isotopic, as we will now see. Consider an embedding of the unit ball $\iota: \mathbb{C}^n \supset B^{2n} \hookrightarrow X$ such that the following hold:


1. $\iota(0) = x_0$
2. $\iota^{-1}(L_0) = \mathbb{R}^n \cap B^{2n} \subset \mathbb{C}^n$
3. $\iota^{-1}(L_1) = i\mathbb{R}^n \cap B^{2n} \subset \mathbb{C}^n$
4. $\iota^* \omega = \alpha \omega^{\text{std}}$ for some $\alpha > 0$

where ω^{std} is the standard symplectic form on \mathbb{C}^n . This seems very restrictive, but this is purely a statement about the differential geometric aspects of these Lagrangian submanifolds so there are no holomorphicity requirements on the embedding $B^{2n} \rightarrow X$. As such, a variant on the Weinstein tubular neighbourhood theorem will guarantee the existence of such an embedding. Next, we take an embedding $\gamma: \mathbb{R} \rightarrow \mathbb{C}$ which satisfies $\gamma(t) = t$ for $t \leq -1/2$, $\gamma(t) = it$ for $t \geq 1/2$ and $\gamma(\mathbb{R}) \cap -\gamma(\mathbb{R}) = \emptyset$. One may show that $H = \cup_t \gamma(t) S^{n-1} \subset \mathbb{C}^n$ is a Lagrangian submanifold. By construction, $H \cap (\mathbb{C}^n \setminus B^{2n}) = (\mathbb{R}^n \cup i\mathbb{R}^n) \setminus B^{2n}$. The plan is now clear: we define


$$L_0 \natural L_1 = (L_0 \setminus \iota(B^{2n})) \cup (L_1 \setminus \iota(B^{2n})) \cup \iota(H \cap B^{2n})$$

This is a submanifold of X , and we call it the Lagrangian connected sum of L_0 and L_1 . This procedure may be illustrated as follows:

8.3. WALL-CROSSING



A different γ or ι will lead to a Hamiltonian isotopic submanifold, so for our purposes, this is a good definition. We can already see that this connected sum is asymmetric. The two options correspond to the two distinct ways of resolving the singularity of $\mathbb{R} \cup i\mathbb{R} \subset \mathbb{C}$ at the origin. For example, if we picture the above scenario as happening on the torus, then the following two diagrams are not Hamiltonian isotopic:



As abstract smooth manifolds, $L_0 \natural L_1$ is obviously diffeomorphic to $L_1 \natural L_0$ because they are both the familiar connected sum $L_0 \# L_1$, which is symmetric (up to diffeomorphisms). However, in the ambient space X , the two may not be deformed into one another by a Hamiltonian isotopy. This is the crucial difference.

Now we know how to construct the Lagrangian connected sum. But for our purposes, Lagrangians also come with a grading. Hence, we need to perform a graded Lagrangian connected sum. This is not always possible. In fact, rather remarkably, it is possible to give $L_0 \natural L_1$ a grading which is compatible with the gradings on L_0 and L_1 if and only if the intersection point $\{x_0\} = L_0 \cap L_1$ has absolute Maslov index 1.

Proposition 8.3.2. [98] *Suppose the intersection point x_0 has absolute Maslov index 1. Then there exists a grading on $L_0 \natural L_1$ which agrees with that of L_0 and L_1 on the respective intersections $(L_0 \natural L_1) \cap L_0$ and $(L_0 \natural L_1) \cap L_1$.*

8.3. WALL-CROSSING

The converse is likewise true. Observe that this forces the Floer cohomology $HF^\bullet(L_0, L_1)$ to be concentrated purely in degree 1, i.e. $HF^\bullet(L_0, L_1) = HF^1(L_0, L_1)$. It is suggested in the references above that the graded connected sum should be the mirror analogue of extensions of bundles, which are controlled by $\text{Ext}^1(\mathcal{E}, \mathcal{F})$. Once again, this seems very tantalising, since mirror symmetry relates these two cohomology groups.

In the same way that a bundle extension can cease to be stable, so too is there the possibility that a Lagrangian connected sum (which we will always assume to be compatibly graded from now on) will cease to be stable as we vary the complex structure on (X, ω) . This time, we can construct quite explicit examples of this wall-crossing phenomenon occurring, due to Joyce [95], predicated on the so-called Lawlor neck. As per the construction of the Lagrangian connected sum, we would like to find some special Lagrangian submanifolds in \mathbb{C}^n . To start with, we look for special Lagrangian planes $\mathbb{R}^n \subset \mathbb{C}^n$. We write such an embedding by specifying its angles, say

$$\Pi^\phi = \{(\exp(i\phi_1)x_1, \dots, \exp(i\phi_n)x_n) \mid x_i \in \mathbb{R}\}$$

With respect to the standard holomorphic volume form $\Omega = dz^1 \wedge \dots \wedge dz^n$, the phase is given by

$$\phi(\Pi^\phi) = \arg \int_{\Pi^\phi} \Omega = \sum_k \phi_k$$

Naturally, we fix $0 \leq \phi_1 < 2\pi$ and $0 \leq \phi_k < \pi$ for $k > 1$. Taking $\phi_1 \mapsto \phi_1 + \pi$ reverses the orientation, and we will denote by $\Pi^{(\pi, 0)}$ the standard embedding $\mathbb{R}^n \subset \mathbb{C}^n$ with its orientation reversed. Now, suppose we have a transverse intersection $\Pi^{(\pi, 0)} \cup \Pi^\alpha$. In particular, this forces $\phi_k > 0$ for all k . We assume that $\phi(\Pi^{(\pi, 0)}) = \phi(\Pi^\phi)$, which can hold only if $n > 1$. In this case, the union of these Lagrangian planes may be viewed as the limit of a family of special Lagrangian submanifolds of \mathbb{C}^n . Smooth members of this family are called Lawlor necks, and they can be described explicitly through equations. Namely, define a rational function

$$P(x) = \frac{\prod^n (1 + a_k x^2) - 1}{x^2}$$

and a real number A . The latter is denoted as such because it will be the A appearing in $\int_L \Omega = A \exp(i\phi(L))$. Then the a_j in the above equation are uniquely determined by the following constraints:

$$\phi_k = a_k \int_{-\infty}^{\infty} \frac{dx}{(1 + a_k x^2) \sqrt{P(x)}} \quad A = \frac{\text{Vol}(S^n)}{\sqrt{a_1 \dots a_n}}$$

The reason we assumed that the phases of $\Pi^{(\pi, 0)}$ and Π^ϕ agree, is that these equations have a solution if and only if $\sum \phi_k = \pi$. One may then define functions $\gamma_k : \mathbb{R} \rightarrow \mathbb{C}$ by

$$\gamma_k(y) = \exp\left(i a_k \int_{-\infty}^y \frac{dx}{(1 + a_k x^2) \sqrt{P(x)}}\right) \sqrt{\frac{1}{a_k} + y^2}$$

8.3. WALL-CROSSING

This leads to the Lawlor neck, which is defined by the real constant A and the angles $\{\phi_k\}$. We write

$$L^{A,\phi} = \{(\gamma_1(y)x_1, \dots, \gamma_n(y)x_n) \mid y, x_1, \dots, x_n \in \mathbb{R}, \sum x_k^2 = 1\}$$

Such a manifold is diffeomorphic to $S^{n-1} \times \mathbb{R}$. The x_k are the spherical coordinates, and the y is the \mathbb{R} -coordinate.

Proposition 8.3.3 ([99]). *$L^{A,\phi}$ is a special Lagrangian submanifold of $(\mathbb{C}^n, \omega^{std}, \Omega^{std})$.*

It is quite easy to picture what this family looks like, as A varies, for the case $n = 2$. We get a cylinder which is squeezed as we approach the origin, and as A goes to 0, this shape approaches the union of two cones, whose tips meet at the origin. Note, however, that we should really imagine this as happening in $\mathbb{C}^2 = \mathbb{R}^4$. Viewed in this 4-dimensional world, we asymptotically approach $\Pi^{(\pi,0)}$ as $y \rightarrow -\infty$, and Π^ϕ as $y \rightarrow \infty$. As promised, then, the union of the two linear subspaces may be viewed as the limit of this family.

Next, the idea is to use the above construction globally. That is, suppose (\mathcal{X}, ω) is a Calabi-Yau manifold. Suppose we have two BPS A -branes L_0, L_1 on X , i.e. special Lagrangians with flat line bundles, which intersect transversally in a point $x_0 \in X$ with absolute Maslov index 1. After using some unitary transformation if necessary, we may assume that $T_{x_0}L_0 = \Pi^0 \subset T_{x_0}X \cong \mathbb{C}^n$ and $T_{x_0}L_1 = \Pi^\phi$ for some $\phi \in [0, 2\pi) \times [0, \pi) \times \dots \times [0, \pi)$. Suppose we let the complex structure vary slightly over $(-\varepsilon, \varepsilon)$, so that the phase difference remains sufficiently small. We denote the fibre over $t \in (-\varepsilon, \varepsilon)$ by X_t .

Theorem 8.3.4 (Joyce). *For all $t \in (-\varepsilon, \varepsilon)$, there exists a special Lagrangian submanifold $L_0 \natural L_1 \subset X_t$ which is close⁶ to $L_0 \cup L_1$ if and only if $\phi_t(L_1) \leq \phi_t(L_0)$.*

The Lagrangian connected sum in the theorem above is constructed using the Lawlor necks, and gluing them into the union $L_0 \cup L_1$ as the Lagrangian surgery procedure dictates. Since the intersection point had index 1, the special Lagrangian $L_0 \natural L_1$ furthermore has a grading which is compatible with those of L_0 and L_1 respectively.

Proposition 8.3.5. [1] *In the above theorem, we may view the family of special Lagrangians as a family of Lagrangian submanifolds in a fixed symplectic manifold. The special Lagrangians are special representatives for this family as the complex structure varies.*

In the subsequent paper by Thomas and Yau, the following is proved.

Theorem 8.3.6. [2] *Let $L \subset X$ be an A -brane. Then there exists at most one (smooth) special Lagrangian in the Hamiltonian deformation class of L .*

This theorem is some remarkable evidence for a potential Kobayashi-Hitchin-like correspondence, on the side of A -branes.

In summary: let $\mathcal{M}_c(X)$ denote the complex moduli space of X . Locally, the holomorphic volume form Ω gives a coordinate in $\mathcal{M}_c(X)$ via the period integrals. Suppose

⁶Formulating this statement precisely requires the theory of currents, from geometric measure theory.

8.4. THE THOMAS-YAU CONJECTURE

$U \subset \mathcal{M}_c(X)$ is such that the holomorphic volume form provides local coordinates. Let L_0 and L_1 be special Lagrangian submanifolds of X . Define

$$U^\pm = \{\Omega \in U \mid \pm(\phi_\Omega(L_0) - \phi_\Omega(L_1) + \pi) < 0\}$$

Then for U^- , there exists a special Lagrangian submanifold in the homology class $[L_0] + [L_1]$. On the wall where $\phi_1(L_0) + \pi = \phi(L_1)$, this special Lagrangian degenerates to a singular union of special Lagrangian, which are of the same phase (since we take the phase to be modulo π). For U^+ , it disappears. This is the wall-crossing that happens on the mirror side. By analogy, notice that on U^- we have $\mu(L_0) < \mu(L_1)$. On the wall, we have $\mu(L_0) = \mu(L_1)$, while $\mu(L_0) > \mu(L_1)$ on U^+ . Naturally, the objects still exist as Lagrangian submanifolds (or even as A -branes), but they are no longer special (or BPS branes). In the same way, the bundle extension [30](#) continues to define a coherent sheaf, even when it is no longer a (semi-)stable sheaf, i.e. when it no longer admits a HYM connection. This HYM connection on the bundle side is the unique representative of its gauge orbit, in the same way that the special Lagrangians are unique representatives of orbits of Lagrangian submanifolds under the group of Hamiltonian deformations.

8.4 The Thomas-Yau Conjecture

The wall-crossing phenomenon motivates the following definition, found in [\[1\]](#). Compare to what happens in [30](#).

Definition 8.4.1. Let L_0 and L_1 be A -branes such that $L_0 \looparrowright L_1$ can be defined.⁷ Let L be another A -brane. Then L is said to be destabilised by (L_0, L_1) if $L \approx L_0 \looparrowright L_1$ and $\phi(L_0) \geq \phi(L_1)$. If no such pair exists, then L is said to be stable.

If the A -branes are mirror to some coherent sheaves, then the above is a very close analogue for the stability of sheaves. Note, however, that we are using the phase $\phi(L_i)$ rather than the slope $\mu(L_i)$. They are closely related, but the phase can also detect changes in the orientation of L_i .

Given the above definition, the Thomas-Yau conjecture as it was originally formulated is then an analogue of the Kobayashi-Hitchin correspondence, on the A -side of mirror symmetry.

Conjecture 8.4.2. *[Thomas [1]] Let $L \subset X$ be an A -brane. Then there exists a special Lagrangian submanifold in the Hamiltonian deformation class of L if and only if L is stable.*

In the cited source, Thomas demonstrates that this conjecture holds on elliptic curves. Of course, an elliptic curve \mathbb{C}/Λ is not a Calabi-Yau threefold. However, besides the motivation from Donaldson-Thomas theory, most of the story can be carried over to Calabi-Yau manifolds of any dimension, in which case the elliptic curve becomes the easiest

⁷We have not discussed the case in which there are multiple intersections, and we will not do so explicitly.

8.4. THE THOMAS-YAU CONJECTURE

example. In the follow-up paper [2], Thomas and Yau attempt to relate Thomas's conjecture to the Lagrangian mean curvature flow.

Proposition 8.4.3. [2] *Let (L, θ) be a graded Lagrangian submanifold of a Calabi-Yau manifold. Then $\vec{H} = J\nabla\theta$.*

The mean curvature flow of a submanifold Y of a Riemannian manifold (X, g) is defined to be the 1-parameter family of submanifolds Y_t satisfying

$$\left. \frac{d}{dt} \right|_{t=t_0} Y_t = \vec{H}_{Y_{t_0}} \quad Y_0 = Y$$

A relatively messy proof reveals that in fact, on a Calabi-Yau manifold, the Lmcf exists, in the sense that the evolution of a Lagrangian submanifold under this flow remains Lagrangian. See [100]. Aside from the fact that the grading appears, it is not so clear how this relates to the conjecture by Thomas. The answer is given in loc. cit., by the following result, in which μ denotes the momentum map $L \mapsto \text{Im } \Omega|_L$ that we saw previously.

Proposition 8.4.4. [2] *The gradient flow of $-|\mu|^2$ with respect to an appropriate metric on $\Omega_0^3(L)$ is the Lmcf. This metric is constructed from the volume form on L induced by g .*

As such, it follows that $\frac{d}{dt}\theta = -\Delta\theta$ under the Lmcf, and the range of θ must decrease in time by the maximum principle. Of course, the smallest possible range is when θ is constant, i.e. when L is special Lagrangian. With this in mind, the conjecture is rephrased (in fact sharpened) into the following statement:

Conjecture 8.4.5 (The Thomas-Yau Conjecture [2]). *Suppose that $L \subset X$ is a stable Lagrangian submanifold. Then the Lmcf of L exists for all time, and it converges to the unique special Lagrangian in the Hamiltonian isotopy class of L conjectured in 8.4.2.*

The authors then work out an example for which the conjecture holds, namely a family of affine quadrics. This is a non-compact manifold, and it is not quite Calabi-Yau. However, it is sufficiently close to being Ricci flat that the Lmcf has similar behaviour, and so the fact that the conjecture holds for this example is a good indication that the underlying philosophy is correct, even if the conjecture as it is stated does not turn out to be true.

We should note that according to e.g. [12], the Thomas-Yau conjecture as stated above is very unlikely to be true, but mostly for technical reasons, namely the long time existence of solutions to the mean curvature flow. He then outlines how to modify the conjecture within the framework of Bridgeland stability, to incorporate the singularities which a Lagrangian submanifold may develop under the mean curvature flow. See also [101] for a recent re-interpretation of the Thomas-Yau conjecture, which is less ambitious than the proposal by Joyce (which we will discuss later).

8.5 Example: The Elliptic Curve

The simplicity of elliptic curves allows us to write everything down explicitly, so let us do that. We remind the reader that we discussed the Fukaya category of the elliptic curve in 4.4. The upshot was as follows: every object in the Fukaya category can be represented by a geodesic, together with a flat unitary line bundle on it. This data is specified completely by three numbers, so we denote the corresponding object by $L_{(y_0, k, \beta)}$. Here, y_0 is the y -intercept of the geodesic, lifted to \mathbb{C} , and k is its slope (although we should really consider its lift to a grading α). The flat connection ∇ is defined by $\nabla = d - 2\pi i \beta dx$.

Firstly, identifying special Lagrangian submanifolds is easy. They are geodesics. The slope $\mu(L_\gamma) = k$ is going to have a nice geometric interpretation in this example. Namely, we consider the pullback:

$$\gamma^* \Omega = A(dt + ikdt)$$

Write $A(1 + ik) = r \exp(-i\theta)$. Then we can see that $\exp(i\theta)\Omega|_{L_\gamma} = r dt$, implying that L_γ is indeed special Lagrangian of phase θ . As such, the slope is give by $\mu(L_\gamma) = \tan(\theta) = k$, which is precisely the slope of the line $\gamma: \mathbb{R} \rightarrow \mathbb{C}$.


Recall also from 4.4 that the graded Lagrangian submanifolds are those which are homotopically non-trivial. The intuition is clear, then. The Thomas-Yau conjecture suggests that the Lmcf for homotopically non-trivial Lagrangians should converge to straight lines with rational slope. This is in fact a known result within geometric measure theory, see [102], which predates the inception of mirror symmetry.

Theorem 8.5.1. *The mean curvature flow of a homotopically non-trivial curve on a Riemannian 2-manifold converges to a geodesic.*


In this case, the Lmcf is known as a curve-shortening flow. This tells us that every graded Lagrangian submanifold will converge to a line with rational slope, which is a special Lagrangian submanifold. Hence, the Thomas-Yau conjecture additionally asserts that every graded Lagrangian submanifold in \mathbb{C}/Λ is stable. This is also the case, as we now explain.

Suppose that (L, θ) is an unstable Lagrangian. By definition, this means that $L \approx L_0 \natural L_1$ for two graded Lagrangians (L_i, θ_i) , with $\phi(L_0) \geq \phi(L_1)$. Using a Hamiltonian isotopy, we may as well assume that the L_i are special Lagrangian. For concreteness, take $\phi(L_0) = 3\pi/4$ and $\phi(L_1) = \pi/4$. Then it is not even possible to form the graded connected sum, because the grading would become discontinuous. Instead, we have to first reverse the orientation on, say, L_0 so that we can form the graded connected sum $L_0[1] \natural L_1$, but now we have $\phi(L_0) = -\pi/4 < \pi/4 = \phi(L_1)$ and so this does not destabilise (L, θ) anymore. In conclusion, it is impossible for a Lagrangian submanifold of \mathbb{C}/Λ to be unstable. To see this pictorially, we use the following diagram:

8.5. EXAMPLE: THE ELLIPTIC CURVE




This is the torus viewed as a square, with the usual identifications. The directions of the arrows indicate the phase of the special Lagrangian submanifolds. Notice, then, the impossibility of the diagram on the right hand side. The arrows which result from the connected sum point in opposite direction, and so as we approach the point at which the intersection was resolved, we necessarily get a discontinuity in the phase, arising from a mismatch between the orientations. Instead, we need to take the shifted special Lagrangian $L_0[1]$, reversing the orientation on L_0 , which results in the following diagram:



This time, the phases align in the right way, as can be seen by following the arrows, so that we can in fact form the connected sum $L_0[1] \natural L_1$ without issues. The same argument applies whenever $\phi(L_0) \geq \phi(L_1)$, and so it is not possible for a Lagrangian submanifold of \mathbb{C}/Λ to be unstable. This proves the Thomas-Yau conjecture for the 2-torus, as was already noted by Thomas in his original paper.

Bearing in mind mirror symmetry for the elliptic curve, we could say that the diagrams we drew above really do represent sheaves, and their extensions. For example, the special Lagrangian with $y_0 = 0$ and $k = 1$ corresponds to $\mathcal{O}(p_0)$. The special Lagrangian with $y_0 = 0$ and $k = 0$ corresponds to $\mathcal{O}_{\mathcal{X}} = \mathcal{O}$, the structure sheaf. Moreover, we can concretely relate the Lagrangian connected sum to bundle extensions. To do this, we calculate $\text{Ext}^1(\mathcal{O}(p_0), \mathcal{O})$. By Serre duality, this is $\text{Ext}^0(\mathcal{O}, \mathcal{O}(p_0))^* = \text{Hom}(\mathcal{O}, \mathcal{O}(p_0)) = H^0(\mathcal{X}, \mathcal{O}(p_0))^* \cong \mathbb{C}$, since $\mathcal{O}(p_0)$ has degree 1. Therefore, there exists a unique non-trivial extension \mathcal{E} . Pictorially, this extension corresponds to the Lagrangian connected sum (notice that the resulting slope of $L_0 \natural L_1$ is not integer, which is why it corresponds to a higher rank bundle):

8.6. EXAMPLE: COHOMOGENEITY ONE LAGRANGIAN SUBMANIFOLDS IN T^{2n}



Notice that $\mu(L_{(y_0, k, \beta)}) = k = \deg(\mathcal{L}(y_0, k, \beta)) = \mu(\mathcal{L}(y_0, k, \beta))$. For the given example, we furthermore have $\deg(\mathcal{E}) = 1$, so $\mu(\mathcal{E}) = 1/2 = \mu(L_0 \rightsquigarrow L_1)$. This is a general feature: the slopes are invariant under mirror symmetry. This is non-trivial, because a priori, homological symmetry only knows about symplectic data on the A -side, not the holomorphic volume form which defines the notion of stability. Likewise, the B -side does not know about the Kähler structure which defines the notion of stability for sheaves. This is another indication that there really is more to the story than just homological mirror symmetry.

8.6 Example: Cohomogeneity one Lagrangian submanifolds in T^{2n}

If we want to do any sort of explicit calculations, then we need to know an explicit expression for the Calabi-Yau metric. The only compact Calabi-Yau manifolds for which such metrics are known, are flat tori. Thus, we either need to work with non-compact Calabi-Yau manifolds, or the flat torus $(\mathbb{C}^n/\Lambda, g)$ where Λ is the lattice generated by the standard (real) basis for \mathbb{C}^n . In this section, we will prove a Thomas-Yau-type result for Lagrangian submanifolds with a large symmetry group. These are known as cohomogeneity one Lagrangian submanifolds in the literature. Imposing these restrictions on the symmetry group greatly simplifies the problem to the case of the Thomas-Yau conjecture on a 2-torus.

So we take $(\mathbb{C}^n/\Lambda, g)$, where the metric is the induced flat metric which we write in coordinates $(x^1, \dots, x^n, y^1, \dots, y^n)$ as

$$g = \sum_i dx^i \otimes dx^i + \sum_i dy^i \otimes dy^i$$

Naturally we may view \mathbb{C}^n/Λ as $T^{2n} = S^1 \times \dots \times S^1$. We will consider an action of T^{n-1} on T^{2n} such that the total space becomes a principal T^{n-1} -bundle $\pi: T^{2n} \rightarrow T^{n+1}$. Then we will restrict our attention to submanifolds $L \subset T^{2n}$ which are invariant under this T^{n-1} action. We will denote $k = n - 1$. If L is a T^k -invariant submanifold in T^{2n} , then $\pi(L) \subset T^{k+2}$ must be a curve. We will assume that L is closed, so that $\pi(L)$ is a closed curve. Conversely, every closed curve $\gamma: I \rightarrow T^{k+2}$ determines a T^k -invariant submanifold in

8.6. EXAMPLE: COHOMOGENEITY ONE LAGRANGIAN SUBMANIFOLDS IN T^{2n}

T^{2n} by $L_\gamma := \pi^{-1}(\gamma)$, where we also denote by γ the image of γ . Thus, we have a bijective correspondence between closed curves in T^{k+2} , and closed T^k -invariant submanifolds in T^{2n} . We denote the submanifolds associated to a closed curve γ by L_γ , so $\pi(L_\gamma) = \gamma$. Clearly, each L_γ is a $(k+1)$ -torus as an abstract smooth manifold. We will denote the coordinates on the base of the principal T^k -bundle by $\{p^i\} \subset \{x^1, \dots, x^n, y^1, \dots, y^n\}$, and the coordinates on the fibre by $\{q^i\}$, i.e. $\{q^i\} = \{x^1, \dots, x^n, y^1, \dots, y^n\} \setminus \{p^i\}$.

Proposition 8.6.1. *The mean curvature flow of L_γ coincides with the preimage of the mean curvature flow of γ .*

Proof. We assume that γ is parameterised by arc-length, and we denote by $\dot{\gamma}$ the derivative with respect to arc-length. We also fix a global orthonormal frame $\{e_i\}$ for T^n , namely

$$\left\{ \frac{\partial}{\partial p^1}, \dots, \frac{\partial}{\partial p^{k+2}}, \frac{\partial}{\partial q^1}, \dots, \frac{\partial}{\partial q^k} \right\} = \{e_i\}$$

Then we compute

$$\nabla_{\dot{\gamma}} \dot{\gamma} = \nabla_{\dot{\gamma}^i e_i} (\dot{\gamma}^j e_j) = \ddot{\gamma} + \dot{\gamma}^i \dot{\gamma}^j \nabla_{e_i} e_j = \ddot{\gamma}$$

since $\nabla_{e_i} e_j = \Gamma_{ij}^k e_k$, and the connection coefficients vanish for the flat metric. In fact, $\dot{\gamma}^i = 0$ whenever $i > k+2$, because the curve lies in the base space T^{k+2} .

Now, the mean curvature of L_γ is the normal projection of

$$\left(\sum_{j=k+3}^{2n} \nabla_{e_j} e_j \right) + \nabla_{\dot{\gamma}} \dot{\gamma}$$

Immediately, the sum drops out because $\nabla_{e_j} e_j = 0$ for the flat metric. Since $\dot{\gamma}^j = 0$ whenever $j > k+2$, and $\langle \dot{\gamma}, \ddot{\gamma} \rangle = 0$ the normal projection becomes

$$\ddot{\gamma} - \langle \dot{\gamma}, \ddot{\gamma} \rangle \dot{\gamma} - \sum_{j>k+2} \langle e_j, \ddot{\gamma} \rangle e_j = \ddot{\gamma}$$

where we used that $\{\dot{\gamma}\} \cup \{e_j\}_{j>k+2}$ is a global frame for L_γ . It follows that the mean curvature of L_γ at a point in $\pi^{-1}(\gamma(t))$ is just $\ddot{\gamma}(t)$. It is clear, then, that the mean curvature flow will preserve T^k -invariance, so the time evolution of L_γ along the mean curvature flow will be of the form L_{γ_s} for some 1-parameter family of curves γ_s . Such a L_{γ_s} is a solution to the mean curvature flow if and only if

$$\frac{\partial L_{\gamma_s}}{\partial s} = \ddot{\gamma}_s$$

Thus, we see that the mean curvature flow of L_γ is L_{γ_s} where γ_s is the curve shortening flow (i.e. 1-dimensional mean curvature flow) of the curve γ in T^{k+2} . \square

The Thomas-Yau conjecture for T^k -invariant Lagrangians may then be reduced to studying the curve shortening flow. Thus, we are very close to the Thomas-Yau conjecture on T^2 . The difference is that the curve shortening flow for space curves, rather than planar curves, is not as easy to study. Hence, we would like to show the following.

8.6. EXAMPLE: COHOMOGENEITY ONE LAGRANGIAN SUBMANIFOLDS IN T^{2n}

Proposition 8.6.2. *Let $\gamma : I \rightarrow T^{k+2}$ be a closed curve. Then L_γ is a Lagrangian submanifold of T^n only if γ is contained in a 2-torus.⁸*

Proof. Suppose that we have a pair (x^i, y^i) of coordinates such that one of the following holds:

1. Both x^i and y^i are coordinates for the fibre.
2. One of the coordinates is a coordinate for the fibre, and γ has non-zero derivative along the other coordinate.

In either of these cases, L_γ cannot be a Lagrangian submanifold. Indeed, since the module $\Omega^2(L_\gamma)$ is free:

1. We may complete $dx^i \wedge dy^i$ to a basis for $\Omega^2(L_\gamma)$. Hence, $\omega|_{L_\gamma} \neq 0$.
2. We may complete $dt \wedge dx^i$ and $dt \wedge dy^i$ to a basis for $\Omega^2(L_\gamma)$. Then $\omega|_{L_\gamma}$ expressed as a linear combination of the basis elements contains the terms

$$\pi_{y_i}(\dot{\gamma})dx^i \wedge dt + \pi_{x_i}(\dot{\gamma})dt \wedge dy^i$$

implying that $\omega|_{L_\gamma} \neq 0$. We denote by π_{y_i} and π_{x_i} the projection onto the respective components of the vector $\dot{\gamma}$.

As such, neither of these two scenarios can occur if L_γ is Lagrangian. But this means we need to make k choices out of $k+1$ pairs $(x^1, y^1), \dots, (x^n, y^n)$, which will be the coordinates of the fibres. Once a pair has been chosen, we choose one of the two coordinates to be along the fibre. The derivative of γ in the direction of the other coordinate is necessarily zero, by the above observation. We are left with a single pair, say (x^j, y^j) . Since we had to take $\pi_{x_i}(\dot{\gamma}) = \pi_{y_i}(\dot{\gamma}) = 0$ for $i \neq j$, we find

$$\omega|_{L_\gamma} = \pi_{x_j}(\dot{\gamma})\pi_{y_j}(\dot{\gamma})dt \wedge dt = 0$$

Therefore, γ is unconstrained in the (x^j, y^j) -plane, but it cannot move in any other direction. Thus, γ is contained in a 2-torus. \square

This reduces the study of the mean curvature flow of T^k -invariant Lagrangians, to the mean curvature flow of curves on T^2 , which is well-understood. It remains to identify which of the T^k -invariant Lagrangians can be graded. Since the mean curvature can also be given as $H = J\tilde{d}\theta$ it follows that $\dot{\theta} = \kappa$, where κ is the curvature of γ in the (x^j, y^j) -plane.

Proposition 8.6.3. *The T^k -invariant Lagrangian L_γ can be graded if and only if $\int_\gamma \kappa = 0$.*

⁸By which we mean: the image of a real affine plane $V \subset \mathbb{C}^n$ under the projection to T^{2n} , which is diffeomorphic to T^2 .

8.6. EXAMPLE: COHOMOGENEITY ONE LAGRANGIAN SUBMANIFOLDS IN T^{2n}

Proof. Let $[\frac{d\theta}{2\pi}]$ denote the cohomology class of $\frac{d\theta}{2\pi}$ in $H^1(L_\gamma, \mathbb{Z})$. Since L_γ is topologically a $(k+1)$ -torus, a basis for its first homology is given by the cycles $\alpha, \beta_1, \dots, \beta_k \in H_1(L_\gamma, \mathbb{Z})$, where we choose $\alpha = [\gamma]$ and β_j the class of the circle which is defined by the fibre coordinate q^j on T^k . We see that $\int_{\beta_j} d\theta = 0$ since θ is constant along the q^j -coordinate by T^k -invariance. Furthermore,

$$\frac{1}{2\pi} \int_\alpha d\theta = \frac{1}{2\pi} \int_\gamma d\theta = \frac{1}{2\pi} \int_\gamma \kappa$$

Therefore, $[\frac{d\theta}{2\pi}] = 0$ if and only if $\int_\gamma \kappa = 0$, which proves the result. \square

Corollary 8.6.4. *The T^k -invariant Lagrangian L_γ can be graded if and only if γ is homotopically non-trivial in T^{k+2} .*

Proof. If γ is a homotopically trivial closed loop in T^{k+2} , then $\int_\gamma \kappa = \pm 2\pi$. Conversely, if γ is homotopically non-trivial, then $\int_\gamma \kappa = 0$. \square

Putting it all together:

Theorem 8.6.5. *Let $\pi : T^{2n} \rightarrow T^{k+2}$ be the trivial principal T^k -bundle over T^{k+2} . Let L be a graded, closed, T^k -invariant Lagrangian submanifold of T^{2n} . Then $L = \pi^{-1}(\gamma) := L_\gamma$ for some homotopically non-trivial closed curve $\gamma : I \rightarrow T^{k+2}$, such that γ is contained in a 2-torus in T^{k+2} . Furthermore, L_γ converges under the mean curvature flow to L_{γ_0} , where γ_0 is the geodesic (i.e. straight line) in the flat 2-torus to which γ converges under the curve shortening flow.*

Proof. The only statements which we have not yet discussed is that L_{γ_0} is in fact a special Lagrangian submanifold. Consider $\theta \in [0, 2\pi)$ such that

$$\left(\text{Im} \exp(-i\theta)\Omega|_{L_{\gamma_0}} \right)_{\gamma_0(0)} = 0$$

Since the derivative of γ_0 along x^j and y^j is constant, this implies that the same number θ suffices along all of γ . Consequently, T^k -invariance of the holomorphic volume form $\Omega = dz^1 \wedge \dots \wedge dz^n$ implies that the same constant suffices on all of L_{γ_0} , so that $\text{Im} \exp(-i\theta)\Omega|_{L_{\gamma_0}} = 0$, whence L_{γ_0} is a special Lagrangian submanifold. \square

A natural question arises: which homology classes may be represented by such T^k -invariant Lagrangians? The answer is as follows. In 8.6.2, we outlined the possible choices that one can make for the fibre coordinates, which leaves the curve γ constrained to some plane with coordinates (x^j, y^j) . We can then choose for γ to wind either of these two directions as well. As a result, we may obtain 2^{k+1} out of $\binom{2k+2}{k+1} = \dim_{\mathbb{R}} H_n(T^{2n}, \mathbb{R})$ generators as T^k -invariant Lagrangian submanifolds.⁹ In other words, the theorem becomes increasingly less relevant as n increases. For example, for $k = 1, 2, 3$ we have:

⁹Note: not all of these are T^k -invariant with respect to the same T^k -action. We get 2 generators for each possible choice of symplectic T^k -action on T^{2n} .

8.6. EXAMPLE: COHOMOGENEITY ONE LAGRANGIAN SUBMANIFOLDS IN T^{2n}

1. 4 out of 6 generators
2. 8 out of 20 generators
3. 16 out of 70 generators

The trend is clear. More generally, there are

$$\frac{(2n)!}{(n!)^2} - 2^n$$

generators for the homology of T^{2n} which do not admit a T^k -invariant Lagrangian representative. Nevertheless, we have shown the following result, since every T^k -invariant Lagrangian submanifold converges to a special Lagrangian submanifold in T^{2n} .

Corollary 8.6.6. *The Thomas-Yau conjecture holds for T^k -invariant Lagrangians in the trivial T^k -bundle over T^{k+2} , with respect to T^k -invariant Hamiltonian isotopies.*

Proof. When we restrict our attention to T^k -invariant Hamiltonian isotopies, the problem once again reduces to the Thomas-Yau conjecture on the 2-torus, which we already know to be true. \square

In essence, the entirety of this argument boils down to the fact that the Thomas-Yau conjecture is known to hold on the 2-torus. Suppose that we know that the Thomas-Yau conjecture holds on some Calabi-Yau manifold \mathcal{Y} . Then we are naturally led to the following.

Theorem 8.6.7. *If the Thomas-Yau conjecture holds on \mathcal{Y} , then it holds for T^k -invariant Lagrangians in $\mathcal{X} = \mathcal{Y} \times \mathbb{C}^k / \Lambda$, where T^k acts symplectically on \mathbb{C}^k / Λ .*

We denote by $\pi_1 : \mathcal{X} \rightarrow \mathcal{Y}$ and $\pi_2 : \mathcal{X} \rightarrow \mathbb{C}^k / \Lambda$ the natural projection maps.

Proof. Suppose that $N \subset Y$ is a Lagrangian submanifold. Then $L = N \times T^k \subset X$, the T^k -invariant lift of N , is a Lagrangian submanifold, since $T_x L = T_{\pi_1(x)} L \oplus T_{\pi_2(x)} T^k$, and X is equipped with the symplectic structure of the product $Y \times T^{2k}$. Similarly, the mean curvature vector of L in X is the T^k -invariant lift of $\vec{H}_N \in \Gamma(Y, TN^\perp)$ to a section of $\Gamma(L, TL^\perp) \subset \Gamma(L, TX|_L)$. This is because the product metric on X is flat along the fibres of π_2 . Therefore, the mean curvature flow of L is the T^k -invariant lift of the mean curvature flow of N . Now, by assumption, the Thomas-Yau conjecture holds on Y , so supposing that N is stable, let N_t be a family of Lagrangians converging to $N_\infty \subset Y$ under the mean curvature flow, with N_∞ special Lagrangian. Let L_t be the T^k -invariant lift of N_t . Then L_t is also a family of Lagrangian submanifolds. We claim that L_∞ is special Lagrangian. Indeed, the holomorphic volume form on \mathcal{X} is $\pi_1^* \Omega_1 \wedge \pi_2^* \Omega_2$, where Ω_j are the holomorphic volume forms on \mathcal{Y} and \mathbb{C}^k / Λ respectively. Then it is easy to verify that $e^{-i\theta} \Omega|_L = 0$, where $\theta = \theta_1 + \theta_2$ with θ_1 the Lagrangian angle of $\pi_1(L_\infty)$ in \mathcal{Y} , and θ_2 the Lagrangian angle of $\pi_2(L_\infty)$ in \mathbb{C}^k / Λ . It follows that L is stable w.r.t. T^k -invariant Hamiltonian isotopies

8.7. EXAMPLE: LAGRANGIAN FIBRATIONS OF $\mathbb{C}\mathbb{P}^n$

if and only if L converges to a T^k -invariant special Lagrangian submanifold under the mean curvature flow. Hence, the T^k -invariant Thomas-Yau conjecture holds on \mathcal{X} . \square

For example, in [9] an S^1 -invariant version of the Thomas-Yau conjecture is shown to hold on hyper-Kähler manifolds which arise through the Gibbons-Hawking ansatz. Whilst these are non-compact manifolds, it has been shown that the Ooguri-Vafa metric (which belongs to this class of metrics) is a very close approximation to the local model of a Ricci flat Kähler metric on an elliptic $K3$ surface (see [103]). We get the following corollary:

Corollary 8.6.8. *Let \mathcal{X} be an elliptic $K3$ surface with the Ooguri-Vafa metric. Then the Thomas-Yau conjecture holds for T^2 -invariant Lagrangians in $\mathcal{X} \times \mathbb{C}/\Lambda$, where we have an S^1 -action on \mathbb{C}/Λ and a (local) S^1 -action on \mathcal{X} , as per the Gibbons-Hawking ansatz.*

8.7 Example: Lagrangian fibrations of $\mathbb{C}\mathbb{P}^n$

As noted, restricting our attention to Calabi-Yau manifolds poses a major problem: we do not know of any Calabi-Yau metrics on compact Calabi-Yau manifolds. As such, what we have instead opted to do here, is to investigate the behaviour of the Lagrangian fibration defined by the standard T^n -action on $\mathbb{C}\mathbb{P}^n$, under the mean curvature flow. These prototypical Kähler-Einstein manifolds are Fano varieties, and so there exists a version of mirror symmetry for them. In fact, one may define the notion of a special Lagrangian submanifold of a Fano variety, after making a choice of anti-canonical divisor D . Indeed, since the anti-canonical bundle of a Fano variety is, by definition, positive, there exists a section $s \in H^0(\mathcal{X}, \mathcal{O}(D))$ which vanishes precisely along the divisor D . Hence, there is a dual section $\Omega \in \Omega_{\mathcal{X}}^n(\mathcal{X} \setminus D)$, which is nowhere vanishing on $\mathcal{X} \setminus D$, which makes $\mathcal{X} \setminus D$ a non-compact Calabi-Yau manifold.

Definition 8.7.1. Let (\mathcal{X}, D) be a Fano variety with an anti-canonical divisor D . A special Lagrangian submanifold of (\mathcal{X}, D) is a special Lagrangian submanifold L of $\mathcal{X} \setminus D$.

We can also construct a category of A -branes and a category of B -branes for topologically twisted string theories compactified on Fano varieties. The category of A -branes will again contain Lagrangian submanifolds, possibly with additional restrictions imposed. However, we must note that special Lagrangian submanifolds in a Fano variety do not enjoy the same properties as special Lagrangian submanifolds of a Calabi-Yau variety. Indeed, when \mathcal{X} is Calabi-Yau, the closed n -form $\text{Re } \Omega$ is a calibration, which means that special Lagrangian submanifolds are volume minimising. When \mathcal{X} is Fano, the closed n -form $\text{Re } \Omega \in \Omega^1(\mathcal{X} \setminus D)$ is not a calibration with respect to the pullback metric, and so these special Lagrangian submanifolds are not volume minimising. For more on mirror symmetry for Fano varieties (and their mirror Landau-Ginzburg models), we refer to [104, 10].

We will consider the Fano variety $\mathbb{C}\mathbb{P}^n$ with its Fubini-Study metric. This is a toric

8.7. EXAMPLE: LAGRANGIAN FIBRATIONS OF $\mathbb{C}\mathbb{P}^n$

symplectic manifold, meaning it carries a Hamiltonian T^n -action $\mathbb{C}\mathbb{P}^n \times T^n \rightarrow \mathbb{C}\mathbb{P}^n$ defined by

$$([z_0 : \cdots : z_n], (e^{i\theta_1}, \dots, e^{i\theta_n})) \mapsto [z_0 : z_1 e^{i\theta_1} : \cdots : z_n e^{i\theta_n}]$$

The corresponding momentum map will be denoted $\mu : \mathbb{C}\mathbb{P}^n \rightarrow (\mathfrak{u}(1)^*)^n \cong \mathbb{R}^n$ and is given by

$$\mu : [z_0 : \cdots : z_n] \mapsto \frac{1}{|z|^2} (|z_1|^2, \dots, |z_n|^2)$$

By the Delzant theorem, the image of μ is a convex polytope, called the Delzant polytope. For our chosen momentum map, this is the standard polytope corresponding to the unit basis vectors and the origin. We denote this polytope by Δ . The boundary divisor $D \subset \mathbb{C}\mathbb{P}^n$ is defined as $D = \mu^{-1}(\partial\Delta)$, which is the set

$$D = \{[z_0 : \cdots : z_n] \in \mathbb{C}\mathbb{P}^n \mid z_0 \dots z_n = 0\}$$

This is an anti-canonical divisor, and a holomorphic volume form on $\mathbb{C}\mathbb{P}^n \setminus D$ is given by $\Omega = d \log z_1 \wedge \cdots \wedge d \log z_n$, in local coordinates $z_0 \neq 0$.

Proposition 8.7.2. *The Lagrangian fibres of the momentum map $\mu : \mathbb{C}\mathbb{P}^n \rightarrow \Delta$ are special Lagrangian in $\mathbb{C}\mathbb{P}^n \setminus D$.*

Proof. The fibres of μ are the orbits of the T^n -action. In $\mathbb{C}\mathbb{P}^n \setminus D \cong (\mathbb{C}^\times)^n$, we can take coordinates (z_1, \dots, z_n) and write the action as

$$((z_1, \dots, z_n), (e^{i\theta_1}, \dots, e^{i\theta_n})) \mapsto (z_1 e^{i\theta_1}, \dots, z_n e^{i\theta_n})$$

Therefore, the orbits are products of circles $L_R = S_{R_1}^1 \times \cdots \times S_{R_n}^1$, where $R = (|z_1|, \dots, |z_n|) \in \mathbb{R}^n$. Then we can compute the phase of Ω on L_R , as follows. Evidently, $\Omega = \frac{1}{z_1 \dots z_n} dz_1 \wedge \cdots \wedge dz_n$. Take radial coordinates $z_k = r_k e^{i\theta_k}$. Then $dz_k = e^{i\theta_k} dr_k + i r_k e^{i\theta_k} d\theta_k$. Pulling back to L_R gives $dz_k|_{L_R} = i r_k e^{i\theta_k} d\theta_k$. Therefore,

$$\Omega|_{L_R} = \frac{i^n r_1 \dots r_n e^{i(\theta_1 + \cdots + \theta_n)}}{r_1 \dots r_n e^{i(\theta_1 + \cdots + \theta_n)}} d\theta_1 \wedge \cdots \wedge d\theta_n = i^n d\theta_1 \wedge \cdots \wedge d\theta_n$$

It follows that L_R is special Lagrangian of phase $n\pi/2$ w.r.t. Ω . □

As previously noted, the reader should not be under the misapprehension that these fibres of the momentum map are volume minimising. In fact, let us consider the case $\mathbb{C}\mathbb{P}^1 \cong S^2$ first, to see what is going on. In this case, the Delzant polytope is the unit interval $\Delta = [0, 1]$. We have $\mu^{-1}((0, 1)) \cong \mathbb{C}^\times$, and the individual fibres are the circles centered at the origin. We can identify these with the fibres of the height function $\mathbb{R}^3 \supset S^2 \rightarrow \mathbb{R}$ sending $(x_1, x_2, x_3) \mapsto x_3$. In particular, the fibre of $1/2 \in (0, 1)$ is the equator of S^2 , which is volume maximising. It is also clear what the behaviour of these fibres will be under the mean curvature flow. If $x_3 < 0$, it will flow towards the south pole. If $x_3 > 0$ it will flow towards the north pole. The equator, i.e. $x_3 = 0$, remains stationary under the flow. As we will see, this behaviour generalises in the obvious way to higher dimensions.

8.7. EXAMPLE: LAGRANGIAN FIBRATIONS OF $\mathbb{C}\mathbb{P}^n$

The idea is as follows. The manifold $\mathbb{C}\mathbb{P}^n$ is obtained as the Kähler reduction of \mathbb{C}^{n+1} with its standard S^1 -action. Therefore, any compact Lagrangian $L \subset \mathbb{C}\mathbb{P}^n$ may be lifted to a compact Lagrangian $\tilde{L} \subset \mathbb{C}^{n+1}$. Since \mathbb{C}^{n+1} is Calabi-Yau, we may compute the mean curvature vector of \tilde{L} from its Lagrangian angle. This is a section of $\tilde{L} \times \mathbb{C}^{n+1} \subset T\mathbb{C}^{n+1}$. We may then project to TS^{2n+1} using a normal vector field, and this is essentially the the mean curvature vector of L , since the resulting projection will be S^1 -invariant.

Given $R \in \mathbb{R}_{>0}^n$, we will consider

$$L_R = \{[1 : R_1 e^{i\theta_1} : \dots : R_n e^{i\theta_n}] \in \mathbb{C}\mathbb{P}^n \mid \theta \in [0, 2\pi)^n\}$$

After reparameterising, this lifts to a Lagrangian submanifold $\tilde{L}_R \subset S^{2n+1}$ given by

$$\tilde{L}_R = \left\{ \frac{1}{\sqrt{1+R^2}} (e^{i\theta_0}, R_1 e^{i\theta_1}, \dots, R_n e^{i\theta_n}) \mid \theta \in [0, 2\pi)^{n+1} \right\}$$

So we want to know the Lagrangian angle of Lagrangian submanifolds L which are products of circles of some given radius. It will be convenient to work in radial coordinates $(r_0, \theta_0, \dots, r_n, \theta_n)$, i.e. $z_k = r_k e^{i\theta_k}$. The standard holomorphic volume form on \mathbb{C}^n then restricts to L as

$$\Omega|_L = i^{n+1} r_0 \dots r_n e^{i(\theta_0 + \dots + \theta_n)} d\theta_0 \wedge \dots \wedge d\theta_n$$

On the other hand, using the obvious orthonormal frame for the normal bundle NL w.r.t. the Euclidean metric, we see that the induced volume form of the metric on L is

$$\text{vol}_L = r_0 \dots r_n d\theta_0 \wedge \dots \wedge d\theta_n$$

As such, the Lagrangian angle of L is given by $\theta_0 + \dots + \theta_n + (n+1)\pi/2$, so the mean curvature 1-form is $H = d\theta_0 + \dots + d\theta_n$. The symplectic form in radial coordinates is given by

$$\omega = r_0 dr_0 \wedge d\theta_0 + \dots + r_n dr_n \wedge d\theta_n$$

The mean curvature vector \vec{H} is related to the mean curvature 1-form H by $H = -\iota_{\vec{H}}\omega$, so it is clear that

$$\vec{H} = -\frac{1}{r_0} \partial_{r_0} - \dots - \frac{1}{r_n} \partial_{r_n}$$

A unit normal vector field to S^{2n+1} is given by $\vec{n} = r_0 \partial_{r_0} + \dots + r_n \partial_{r_n}$. The Euclidean metric is

$$dr_0^2 + \dots + dr_n^2 + r_0^2 d\theta_0^2 + \dots + r_n^2 d\theta_n^2$$

so we have $\langle \partial_{r_j}, \partial_{r_k} \rangle = \delta_{jk}$. Hence, we find $\langle \vec{H}, \vec{n} \rangle \vec{n} = -(n+1)\vec{n}$, implying that

$$\vec{H} - \langle \vec{H}, \vec{n} \rangle \vec{n} = \left((n+1)r_0 - \frac{1}{r_0} \right) \partial_{r_0} + \dots + \left((n+1)r_n - \frac{1}{r_n} \right) \partial_{r_n}$$

which is the mean curvature vector of $L \subset S^{2n+1}$. Since this depends only on the r_k coordinates, it induces a vector field on the Delzant polytope Δ by $d\mu(\pi_* \vec{H})$, where

8.7. EXAMPLE: LAGRANGIAN FIBRATIONS OF $\mathbb{C}\mathbb{P}^n$

we denote by $\mu : \mathbb{C}\mathbb{P}^n \rightarrow \Delta$ the momentum map and $\pi : S^{2n+1} \rightarrow \mathbb{C}\mathbb{P}^n$ is the projection $(z_0, \dots, z_n) \mapsto [z_0 : \dots : z_n]$. Equivalently, this is $d\tilde{\mu}(\vec{H})$, where $\tilde{\mu} : S^{2n+1} \rightarrow \mathbb{R}^n$ is the obvious lift of μ to \mathbb{C}^{n+1} , defined by

$$\tilde{\mu}(z_0, \dots, z_n) = \frac{1}{|z|^2} (|z_1|^2, \dots, |z_n|^2) = \frac{1}{r^2} (r_1^2, \dots, r_n^2)$$

The resulting vector field on Δ , whose coordinates are $x_k = r_k^2$, is then

$$d\tilde{\mu}(\vec{H}) = (2(n+1)x_1 - 2)\partial_1 + \dots + (2(n+1)x_n - 2)\partial_n$$

Clearly, the flow has a stationary point at the barycentre of the polytope, whose fibre corresponds to the Clifford torus in $\mathbb{C}\mathbb{P}^n$. All the other points flow radially outwards, see the cases of $\mathbb{C}\mathbb{P}^2$ and $\mathbb{C}\mathbb{P}^1 \times \mathbb{C}\mathbb{P}^1$ in the images below. The red outline denotes the Delzant polytope, which is $[0, 1] \times [0, 1]$ for $\mathbb{C}\mathbb{P}^1 \times \mathbb{C}\mathbb{P}^1$.




Figure 8.1: The flow on the Delzant polytope of $\mathbb{C}\mathbb{P}^2$ induced by the mean curvature flow.




Figure 8.2: The flow on the Delzant polytope of $\mathbb{C}\mathbb{P}^1 \times \mathbb{C}\mathbb{P}^1$ induced by the mean curvature flow.

8.7. EXAMPLE: LAGRANGIAN FIBRATIONS OF $\mathbb{C}\mathbb{P}^n$

In [11], a different method is used to show a much stronger result. The authors show the following. The main ingredient that goes into this, which we have not touched upon, is a potential function for the Kähler form/metric, denoted here by u .

Theorem 8.7.3 ([11]). *Let (X, ω, g, μ) be a toric Kähler manifold, and give the Delzant polytope Δ the pushforward metric g_Δ (which is possible because of T^n -invariance). The mean curvature of the Lagrangian orbits can be written as $H = -\nabla \log V$, where $V = \det(\text{Hess}(u))^{-1/2}$, with $u : \Delta \rightarrow \mathbb{R}$ a potential for g_Δ . Given $x \in \text{int}(\Delta)$, one has $\mu(L_t) = x_t$ and $\{x_t\}$ is a solution of the negative g_Δ -gradient flow of $\log V$, starting at x .*

This naturally leads one to wonder: can the argument be adapted to say something about the mean curvature flow of the fibres of an almost toric Calabi-Yau manifold (i.e. T^4 or a K3 surface)?

Definition 8.7.4. An almost toric 4-manifold is a Lagrangian fibration of a symplectic 4-manifold such that every critical point of the fibration is an elliptic or a focus-focus singularity.

Neither T^4 nor a K3 are themselves toric manifolds, but they are almost toric Kähler manifolds, see [105]. They are the only Calabi-Yau almost toric manifolds. In their classification, the authors show that any Lagrangian fibration which renders T^4 an almost toric manifold turns T^4 into a principal T^2 -bundle, via the angle coordinates of the Liouville-Arnold theorem. As such, they are cohomogeneity one Lagrangian submanifolds for an S^1 -action on T^4 , which means we have the following corollary to 8.6.5, telling us that the answer to the above question is negative for T^4 (and thus likely also for K3 surfaces).

Corollary 8.7.5. *Let T^4 be given the structure of an almost toric manifold. Then the mean curvature flow does not preserve the fibration, unless it is a special Lagrangian fibration.*

Indeed, this is the expected result in light of the Thomas-Yau conjecture; if we have a Lagrangian fibration, then each of the fibres should converge to a special Lagrangian submanifold. There is no reason for this special Lagrangian to be a fibre in the fibration, and so the mean curvature flow will not preserve the fibration. It is interesting that the toric case differs so much from the almost toric case, in this regard.

Notice that the theorem from [11] applies to toric Kähler manifolds in complete generality, the corresponding T^n -invariant Kähler metric need not even be Kähler-Einstein. For example the first Hirzebruch surface \mathcal{H}_1 is a blow-up of $\mathbb{C}\mathbb{P}^2$ in a point. It has been shown in [106] that such a complex surface does not admit a Kähler-Einstein metric. This is very closely related to a notion of stability, just as for the Hermitian-Yang-Mills equation. Using the results from the paper, we can calculate the mean curvature flow induced on the polytope, purely from the combinatorial data of the polytope. We will do this for the Hirzebruch surfaces \mathcal{H}_n .

8.7. EXAMPLE: LAGRANGIAN FIBRATIONS OF $\mathbb{C}\mathbb{P}^n$

Example 31. The polytope Δ_n of \mathcal{H}_n lies in \mathbb{R}^2 , with vertices given by

$$\{(0, 0), (n + 1, 0), (0, 1), (1, 1)\}$$

The first step is to express this in the form

$$\Delta_n = \{x = (x_1, x_2) \in \mathbb{R}^2 \mid \langle x, v_i \rangle \geq \lambda_i \quad i \in \{0, 1, 2, 3\}\}$$

where the v_i are primitive inward pointing normal vectors. Presently, we have $v_0 = (0, 1)$, $v_1 = (-1, -n)$, $v_2 = (0, -1)$ and $v_3 = (1, 0)$. Since $x_1 \geq 0$, $x_2 \geq 0$, $x_2 \leq 1$ and $x_1 + nx_2 \leq n + 1$, this results in

$$\Delta_n = \{x \in \mathbb{R}^2 \mid x_2 \geq 0, -x_1 - nx_2 \geq -n - 1, -x_2 \geq -1, x_1 \geq 0\}$$

Next, one defines $\ell_i(x) = \langle x, v_i \rangle - \lambda_i$, which presently yields $\ell_0(x) = x_2$, $\ell_1(x) = n + 1 - x_1 - nx_2$, $\ell_2(x) = 1 - x_2$ and $\ell_3(x) = x_1$ and the potential function $u(x) = \sum_i \ell_i(x) \log(\ell_i(x)) - \ell_i(x)$. We obtain a metric on Δ_1 by declaring the metric tensor to be given by the matrix $\text{Hess}(u)$, which is presently given by

$$g_\Delta = \begin{pmatrix} \frac{1}{x_1} + \frac{1}{n+1-x_1-nx_2} & \frac{n}{n+1-x_1-nx_2} \\ \frac{n}{n+1-x_1-nx_2} & \frac{1}{x_2} + \frac{1}{1-x_2} + \frac{n^2}{n+1-x_1-nx_2} \end{pmatrix}$$

Next, one defines

$$V(x) = \det(\text{Hess}(u))^{-1/2} = \left(\frac{x_1 x_2 (x_1 + nx_2 - n - 1)}{n + 1 - nx_2 - n^2 x_2} \right)^{\frac{1}{2}}$$

From this, one can explicitly calculate $-\nabla \log(V)$, as we will do for $n = 1$, because the expression for arbitrary n becomes rather lengthy. For $n = 1$, we get:

$$d \log(V) = \left(\frac{2x_1 + x_2 - 2}{x_1(x_1 + x_2 - 2)}, \frac{x_1(x_2^2 - 4x_2 + 2) + x_2^4 - 8x_2^2 + 12x_2 - 4}{x_2(x_2 - 1)(x_2^2 - 2)(x_1 + x_2 - 2)} \right)$$

The gradient is then obtained simply by multiplying by g_Δ^{-1} (or the isomorphism $T\Delta_1 \rightarrow T^*\Delta_1$, more invariantly). Plotting the result yields an outward radial flow, just as for $\mathbb{C}\mathbb{P}^2$. Use e.g. this mathematica code:

```
n = 2;
u[x, y] = y*Log[y] + (n + 1 - x - n*y)*Log[n + 1 - x - n*y] - n - 2 +
n*y + (1 - y)*Log[1 - y] + x*Log[x];
hess[x, y] = D[u[x, y], {{x, y}, 2}];
ginv[x, y] = Inverse[hess[x, y]];
V[x, y] = 1/Det[hess[x, y]];
diff = Grad[-Log[V[x, y]], {x, y}];
ginvsimp = Simplify[ginv[x, y]];
diffsimp = Simplify[diff];
VectorPlot[ginvsimp . diffsimp, {x, 0, n + 1}, {y, 0, 1}]
```

8.7. EXAMPLE: LAGRANGIAN FIBRATIONS OF $\mathbb{C}\mathbb{P}^n$

The resulting vector plots very closely resemble the behaviour that is exhibited by $\mathbb{C}\mathbb{P}^2$, as can be seen in the images below. Note that $n = 0$ simply corresponds to $\mathbb{C}\mathbb{P}^1 \times \mathbb{C}\mathbb{P}^1$, which we discussed previously.




Figure 8.3: $\mathcal{H}_0 = \mathbb{C}\mathbb{P}^1 \times \mathbb{C}\mathbb{P}^1$




Figure 8.4: \mathcal{H}_1




Figure 8.5: \mathcal{H}_2

Chapter 9

The Thomas-Yau-Joyce Conjecture

The way in which Thomas and Yau phrased their conjecture was somewhat of a departure from the framework of homological mirror symmetry. In fact, this already started in Thomas's original paper, which was motivated by Donaldson-Thomas theory. The Donaldson-Thomas invariants are extracted from moduli spaces of sheaves - the structure of $D^b(\mathcal{X})$ does not make an appearance. Of course, this is unimportant if one's sole purpose is to extract invariants. However, if one wants to relate the Donaldson-Thomas invariants to string theory, then the category $\text{Coh}(\mathcal{X})$ is the category of B -branes at the large volume limit. As such, the Donaldson-Thomas invariants, which count certain BPS states, should have a more general interpretation when we move away from the large volume limit. There should be a notion of "stability" in the category $D^b(\mathcal{X})$ which tells us what the BPS states are, for the corresponding point in the Kähler moduli space. This notion of stability is conjectured to be Bridgeland stability, and its development was in part inspired by the Thomas-Yau conjecture. We will explain what a Bridgeland stability condition is, and how it inspired Joyce's reformulation of the Thomas-Yau conjecture. After doing so, we will conclude by outlining the ideas of Joyce and Song, and independently, Kontsevich and Soibelman, on how to obtain invariants from a Bridgeland stability condition on any CY3-category. These are categories which have the same properties as $D^b(\mathcal{X})$ for a Calabi-Yau threefold \mathcal{X} , and in this specific case, the invariants should be the Donaldson-Thomas invariants of \mathcal{X} (or a suitable refinement). We will explain how this fits into the framework of the Thomas-Yau conjecture, and its underlying philosophy.

9.1 Bridgeland Stability

We considered the notion of μ -stability for vector bundles, in relation to the Hermitian-Yang-Mills equation. Furthermore, we considered (Gieseker) stability to form the moduli space of sheaves. In particular, μ -stability should be the limiting notion of stability on $D^b(\mathcal{X})$ at the large volume limit that we assumed throughout our discussion of homological mirror symmetry. More generally, we have Bridgeland stability, first defined in [93], inspired by the notion of Π -stability which dates back to [6]. Π -stability is the string theoretic notion of stability which was put into a more mathematical framework by Bridgeland. First, we must define a slicing of a triangulated category.

Definition 9.1.1. Let \mathcal{A} be a triangulated category. A slicing of \mathcal{A} is a collection of full, additive subcategories $\{\mathcal{P}(\phi)\}$ for $\phi \in \mathbb{R}$ such that

9.1. BRIDGELAND STABILITY

1. $\mathcal{P}(\phi)[1] = \mathcal{P}(\phi + 1)$
2. If $\phi_1 > \phi_2$, then $\text{Hom}_{\mathcal{A}}(A_1, A_2) = 0$ whenever $A_i \in \mathcal{P}(\phi_i)$
3. Every object $E \in \mathcal{A}$ admits a Harder-Narasimhan filtration, i.e. a sequence of exact triangles

$$\begin{array}{ccccccc}
 0 = E_0 & \longrightarrow & E_1 & \longrightarrow & \dots & \longrightarrow & E_{n-1} & \longrightarrow & E_n = E \\
 & & \swarrow & & & & \swarrow & & \swarrow \\
 & & A_1 & & & & \dots & & A_n
 \end{array}$$

such that $A_i \in \mathcal{P}(\phi_i)$ for some sequence $\phi_1 > \dots > \phi_n$.

Intuitively, this axiomatises the Harder-Narasimhan filtrations which are exhibited by sheaves. The objects of $\mathcal{P}(\phi)$ are the semi-stable objects of phase ϕ , and the first axiom is the behaviour of the phase under the shift functor in the derived category (resp. derived Fukaya category) of a Calabi-Yau threefold.

Definition 9.1.2. Let $(\mathcal{A}, \mathcal{P})$ be a triangulated category together with a slicing. Then a central charge is a group homomorphism¹ $Z : K(\mathcal{A}) \rightarrow \mathbb{C}$ such that $Z(E) = m(E) \exp(i\pi\phi)$ with $m(E) \in \mathbb{R}_{>0}$ whenever $0 \neq E \in \mathcal{P}(\phi)$.

This definition essentially says that the phase ϕ should really be the angle w.r.t. the positive real axis of the complex number $Z(E)$, and the terminology "central charge" comes from physical considerations.

Definition 9.1.3. A Bridgeland stability condition on a triangulated category is a slicing of \mathcal{A} together with a homomorphism $\text{cl} : K(\mathcal{A}) \rightarrow \Gamma$ for some free abelian group Γ , and a homomorphism $Z : \Gamma \rightarrow \mathbb{C}$ such that $Z \circ \text{cl}$ is a central charge for $(\mathcal{A}, \mathcal{P})$. The lattice Γ is called the charge lattice.

By abuse of notation, we also simply denote $Z \circ \text{cl} = Z$.

Remark 9.1.4. We have deviated slightly from the definition of a Bridgeland stability condition as it is normally given, in that the homomorphism $\text{cl} : K(\mathcal{A}) \rightarrow \Gamma$ usually does not make an appearance. However, we will only be interested in cases where such a homomorphism already naturally exists, and so we include it as part of the data.

The point of the homomorphism cl , short for class, is that the objects in \mathcal{A} vary in continuous families. Given $E \in K(\mathcal{A})$, we should think of $\text{cl}(E)$ as being some topological invariant such as the Chern character of a holomorphic vector bundle, or the homology class of a Lagrangian submanifold.

We would like to give some explicit examples of stability conditions. Unfortunately, the notion of μ -stability does not yield a Bridgeland stability condition on $D^b(\mathcal{X})$, when-

¹We denote by $K(\mathcal{A})$ the numerical Grothendieck group, which is a quotient of the usual Grothendieck group in order to make it a countable set.

9.1. BRIDGELAND STABILITY

ever $\dim \mathcal{X} > 1$. Rather, μ -stability is thought to be the large volume limit of some appropriate stability condition. For Calabi-Yau threefolds, there are in general no known Bridgeland stability conditions on $D^b(\mathcal{X})$, see [107]. For abelian threefolds, there has been some progress, namely [108]. Similarly, there are no known Bridgeland stability conditions on the derived Fukaya category of a Calabi-Yau threefold. However, this is where Joyce's interpretation of the Thomas-Yau conjecture will come into play. Before explaining this, let us first present the Bridgeland stability condition provided by μ -stability when $\dim \mathcal{X} = 1$.

We refer to D.4.4 for the definition of a t -structure, and its heart. Suppose that \mathcal{A} is given a t -structure, and let \mathcal{A}^\heartsuit be its heart. We will define a stability function on \mathcal{A}^\heartsuit to be a group homomorphism $Z : K(\mathcal{A}^\heartsuit) \rightarrow \mathbb{C}$, such that $Z(E) = m(E) \exp(i\pi\phi(E))$ with $m(E) \in \mathbb{R}_{>0}$ whenever $E \neq 0$, and $0 < \phi(E) \leq 1$. Once again $\phi(E)$ is called the phase of the object $E \in \mathcal{A}^\heartsuit$. An object $E \in \mathcal{A}^\heartsuit$ is said to be semi-stable (w.r.t. Z) if every subobject $0 \neq F \subset E$ satisfies $\phi(F) \leq \phi(E)$. The Harder-Narasimhan property of Z then amounts to the condition that every object should admit a filtration whose sub-quotients are semi-stable, of increasing phase, similar to the above definition. If these properties are satisfied, the group homomorphism Z is said to be a stability function on \mathcal{A}^\heartsuit .

Theorem 9.1.5. [5] *A Bridgeland stability condition on \mathcal{A} is equivalent to a bounded t -structure together with a stability function on its heart.*

The proof is almost immediate from the definitions. With this in mind, we present an instructive example.

Example 32. Let \mathcal{X} be a smooth algebraic curve (i.e. $\dim \mathcal{X} = 1$). Define a monoid homomorphism $Z : \text{Coh}(\mathcal{X}) \rightarrow \mathbb{C}$ by $Z(E) = -\deg(E) + i\text{rank}(E)$. By additivity of rank and degree on curves, this is indeed a monoid homomorphism, and so can be extended to give a group homomorphism $Z : K(\text{Coh}(\mathcal{X})) \rightarrow \mathbb{C}$. Then Z is a stability function for the standard t -structure on $D^b(\mathcal{X})$ (see right below D.4.4 for this standard t -structure). It is easy to see that the slope of the line going through 0 and $Z(E)$ is precisely the slope of E , i.e. $\mu(E)$. We have omitted the charge lattice here, but it is simply the image of the Chern character in $H^0(X, \mathbb{Q}) \oplus H^2(X, \mathbb{Q})$. Note that we have to take $-\deg(E)$ to make the Harder-Narasimhan property hold, i.e. the Harder-Narasimhan filtration from B.2.

For $\dim \mathcal{X} > 1$, this no longer works because of sheaves with support on submanifolds of codimension greater than 1. For K3 surfaces, there are known stability conditions. As stated, for Calabi-Yau threefolds, there are not, except in the special case of abelian threefolds. Now we can state the conjecture by Joyce, which is the content of [12].

Conjecture 9.1.6 (Joyce). *Let $(\mathcal{X}, \omega, \Omega, g)$ be a Calabi-Yau manifold with derived Fukaya category $D^b \text{Fuk}(X, \omega)$. Then there exists a natural Bridgeland stability condition on $D^b \text{Fuk}(X, \omega)$, with the following properties.*

9.1. BRIDGELAND STABILITY

1. *The charge lattice is $H_3(X, \mathbb{Z})$, and the class map is $(L, E, \nabla) \mapsto [L] \in H_3(X, \mathbb{Z})$.*
2. *The central charge is $Z : [L] \mapsto \int_L \Omega$.*
3. *If (L, E, ∇) lies in the isomorphism class of a special Lagrangian submanifold of phase ϕ , then $(L, E, \nabla) \in \mathcal{P}(\phi)$.*
4. *Every isomorphism class of objects in $\mathcal{P}(\phi)$ contains a unique special Lagrangian representative of phase ϕ , possibly immersed.*

In fact, Joyce is more cautious: he indicates that the final point in his conjecture is probably false, as stated, and this is because, unlike Thomas-Yau, Joyce does not expect the mean curvature flow of stable Lagrangians to exist for all time without acquiring singularities, so there will be some technical issues that need to be addressed. Nevertheless, the bigger picture is clear: this is the natural extension of the Thomas-Yau conjecture to the realm of the derived Fukaya category, together with the notion of Bridgeland stability (which was conceived after the conjecture by Thomas and Yau).

In his paper, Joyce outlines the following programme for proving his conjecture, improving the Thomas-Yau conjecture. It is clear that $Z : (L, E, \nabla) \mapsto \int_L \Omega$ is a central charge, once the main obstacle has been overcome, which is to prove that defining $\mathcal{P}(\phi)$ as in the conjecture indeed yields a slicing of $D^b\text{Fuk}(X, \omega)$. The difficult part is to prove the Harder-Narasimhan property, and the proposal by Joyce is the following. Construct a family $\{(L^t, E^t, \nabla^t)\}_{t \in [0, \infty)}$ such that

1. $(L^0, E^0, \nabla^0) = (L, E, \nabla)$.
2. There is a series $0 < T_1 < \dots$ of singular times such that, for $t \in [0, \infty) \setminus \{T_1, \dots\}$, there is an isomorphism $(L^t, E^t, \nabla^t) \cong (L, E, \nabla)$ in $D^b\text{Fuk}(X, \omega)$.
3. For $t \in [0, \infty) \setminus \{T_1, \dots\}$, the family $\{L_t\}$ satisfies the Lagrangian mean curvature flow.
4. At the singular times T_i , the (possibly singular) Lagrangian submanifolds have to undergo surgery so that the Lagrangian mean curvature flow may be continued.
5. $\lim_{t \rightarrow \infty} L^t = L_1 \cup \dots \cup L_n$ with each L_i a (possibly singular) special Lagrangian with phase ϕ_i , ordered so that $\phi_1 > \dots > \phi_n$, yielding a Harder-Narasimhan filtration for L .

Joyce notes how this is similar to Perelman's proof of the Poincaré conjecture, and speculates that the (complex) dimension 3 case will be of similar difficulty, which higher dimensions being more difficult yet. In complex dimension 2, the Thomas-Yau-Joyce conjecture was verified in [109] for circle invariant Lagrangians in hyper-Kähler manifolds of the Gibbons-Hawking ansatz type.

Remark 9.1.7. A very observant reader may have noticed that, when we stated the homological mirror symmetry conjecture 0.0.1, we made no mention of the so-called Karoubi envelope of the derived Fukaya category. Usually, this is included to make the derived Fukaya category idempotent closed, as the derived category of an algebraic variety is

naturally. Otherwise, it would not be possible for there be an equivalence of categories. Our reason for omitting this, is that the Thomas-Yau-Joyce conjecture implies (if true) that the derived Fukaya category is already idempotent closed, and so we do not need to go through this extra construction. Moreover, the Karoubi envelope would enlarge the derived Fukaya category by including objects for which the central charge cannot be defined.

9.2 Refined Donaldson-Thomas Invariants

Denote by $\mathbb{C}[d]$ the chain complex which is trivial in all degrees except d , where it is \mathbb{C} .

Definition 9.2.1. A CY3 category consists of an A_∞ -category \mathcal{A} , and for each pair of objects $A, B \in \mathcal{A}$, a morphism of chain complexes

$$\langle \cdot, \cdot \rangle : \text{Hom}(A, B) \otimes \text{Hom}(B, A) \rightarrow \mathbb{C}[3]$$

which is

1. Non-degenerate for each pair of objects.
2. Cyclically invariant: $\langle m_{k-1}(a_0 \otimes \cdots \otimes a_{k-2}), a_{k-1} \rangle = (-1)^* \langle m_{k-1}(a_1 \otimes \cdots \otimes a_{k-1}), a_0 \rangle$

The map $\langle \cdot, \cdot \rangle$ is meant to encode Serre duality in a more categorical framework, and will be called the trace map.

Example 33. The most important examples of CY3 categories for us will be:

1. The category $D^b(\mathcal{X})$ for \mathcal{X} a Calabi-Yau threefold, where the trace map is given by Serre duality.
2. The category $D^b\text{Fuk}(X, \omega)$ for \mathcal{X} a Calabi-Yau threefold. The construction of the relevant trace map is the content of e.g. [110].

In a series of papers, Kontsevich-Soibelman [13] (and, independently, Joyce-Song [111]²) outline how to extract enumerative invariants from a Bridgeland stability condition on a CY3 category. We will sketch the procedure below, but explaining it in full detail would require far too much time, so we refer the reader to the cited papers and references therein.

First, suppose that we are given a free abelian group $\Gamma \cong \mathbb{Z}^n$ together with a skew-symmetric bilinear form $\langle \cdot, \cdot \rangle : \Gamma \times \Gamma \rightarrow \mathbb{Z}$. From this, one may construct a Γ -graded Lie algebra over \mathbb{Q} , defined as $\mathfrak{g}_\Gamma := \bigoplus_{\gamma \in \Gamma} \mathbb{Q} \cdot e_\gamma$ with Lie bracket

$$[e_\gamma, e_\mu] = (-1)^{\langle \gamma, \mu \rangle} \langle \gamma, \mu \rangle e_{\gamma+\mu}$$

²The machinery of Joyce-Song only applies to abelian categories, not to triangulated categories. However, their methods are more concrete and less conjectural.

9.2. REFINED DONALDSON-THOMAS INVARIANTS

We can also define a commutative, associative multiplication on the vector space \mathfrak{g}_Γ by setting

$$e_\gamma \cdot e_\mu = (-1)^{\langle \gamma, \mu \rangle} e_{\gamma+\mu}$$

Denote this ring, which is a \mathbb{Q} -algebra, by S_Γ . We set $\mathbb{T}_\Gamma := \text{Spec} S_\Gamma$. This scheme has the natural structure of a Poisson manifold, since its algebra of functions is $S_\Gamma = (\mathfrak{g}_\Gamma, \cdot)$, so it has a Poisson structure given by the Lie bracket on \mathfrak{g}_Γ .

Next, we will need the notion of a stability condition on the graded Lie algebra \mathfrak{g}_Γ . This is a pair (Z, a) where $Z : \Gamma \rightarrow \mathbb{C}$ is a group homomorphism, and $a = \{a(\gamma)\}_{\gamma \in \Gamma} \subset \mathfrak{g}_\Gamma$ is a collection with $a(\gamma) \in \mathbb{Q} \cdot e_\gamma := \mathfrak{g}_\gamma$, such that, for a given norm $\|\cdot\|$ on $\Gamma \otimes_{\mathbb{Z}} \mathbb{R}$, there exists $C > 0$ with the property that $\|\gamma\| \leq C|Z(\gamma)|$ for all $\gamma \in \text{Supp}(a)$.

Suppose we are given a stability condition (Z, a) on \mathfrak{g}_Γ as above. Then there is a function $\Omega : \Gamma \setminus \{0\} \rightarrow \mathbb{Q}$ such that we can uniquely write

$$a(\gamma) = - \sum_{\substack{n \geq 1 \\ \frac{\gamma}{n} \in \Gamma \setminus \{0\}}} \frac{\Omega(\frac{\gamma}{n})}{n^2} e_\gamma$$

We then have an equality

$$\exp\left(\sum_{n \geq 1} a(n\gamma)\right) = \exp\left(-\sum_{n \geq 1} \Omega(n\gamma) \sum_{k \geq 1} \frac{e_{kn\gamma}}{k^2}\right)$$

as elements of the Lie group³ $\exp(\mathfrak{g}_\Gamma)$. The numbers $\Omega(\gamma)$ for $\gamma \in \Gamma \setminus \{0\}$ will be the refined Donaldson-Thomas invariants. So it remains for us to explain how to associate a stability condition (Z, a) to \mathfrak{g}_Γ for some lattice Γ with a bilinear form, given a Bridgeland stability condition on the CY3 categories above.

First, we should specify the lattices and the bilinear forms.

1. In the case of $D^b(\mathcal{X})$, the lattice Γ is the image of $\text{ch} : K(\text{Coh}(\mathcal{X})) \rightarrow H^{\text{ev}}(X, \mathbb{Q})$. The bilinear form is given by the Euler form⁴, and the class map is the Chern character.
2. In the case of $D^b\text{Fuk}(X, \omega)$, the lattice Γ is $H^3(X, \mathbb{Z})$. The bilinear form is given by the polarised Hodge structure on $H^3(X, \mathbb{Z})$, and the class map takes an object to the (Poincaré dual of the) homology class of the Lagrangian submanifold.

From this data, we get the graded Lie algebra \mathfrak{g}_Γ from above, as well as the algebraic Poisson manifold $(\mathbb{T}_\Gamma, \{\cdot, \cdot\})$. The next step is to specify the stability condition (Z, a) on \mathfrak{g}_Γ . To do this, we assume that the respective categories carry a Bridgeland stability whose charge lattice is the one stated above. In this case, the homomorphism $Z : \Gamma \rightarrow \mathbb{C}$

³We are working over \mathbb{Q} here, so this should be called a pro-nilpotent Lie group, but we do not wish to digress into this.

⁴The Euler form is defined on $K(\text{Coh}(\mathcal{X}))$ by defining it on objects $E, F \in \text{Coh}(\mathcal{X})$ as $\chi(E, F) = \sum (-1)^k \dim_{\mathbb{C}} \text{Ext}^k(E, F)$.

9.2. REFINED DONALDSON-THOMAS INVARIANTS

is given by Bridgeland stability condition, so it remains to explain how the collection $a = \{a(\gamma)\}_{\gamma \in \Gamma}$ is obtained. This is the part which is rather involved, and in their paper, Kontsevich-Soibelman use the theory of motives to achieve this. We do not wish to digress into this, and so we will take the following to be a black box:

Theorem 9.2.2 (Kontsevich-Soibelman). *Let \mathcal{A} be a CY3 category with a Bridgeland stability condition and an orientation. Then for every strict sector⁵ in \mathbb{C} , there exist elements $A_V^{\text{mot}} \in \mathcal{R}_V$, where \mathcal{R}_V is an associative ring called the quantum torus associated with V . The collection*

$$\{A_V^{\text{mot}} \mid V \subset \mathbb{C} \text{ a strict subsector}\}$$

are called the motivic Donaldson-Thomas invariants of $(\mathcal{A}, \mathcal{P}, Z)$. Furthermore, let $\mathcal{R}_{\Gamma, q}$ be the $D_q = \mathbb{Z}[q^{1/2}, q^{-1/2}, ((q^n - 1)^{-1})_{n \geq 1}]$ -algebra generated by $\{\widehat{e}_\gamma\}_{\gamma \in \Gamma}$, subject to the relations

$$\widehat{e}_\gamma \widehat{e}_\mu = q^{\langle \gamma, \mu \rangle / 2} \widehat{e}_{\gamma + \mu} \quad \widehat{e}_0 = 1$$

For every strict subsector $V \subset \mathbb{C}$, there is a subalgebra $\mathcal{R}_{V, q} \subset \mathcal{R}_{\Gamma, q}$, and a homomorphism $\mathcal{R}_V \rightarrow \mathcal{R}_{V, q}$.

Denote the image of A_V^{mot} under the homomorphism $\mathcal{R}_V \rightarrow \mathcal{R}_{V, q}$ by $A_{V, q}$. Then each $A_{V, q}$ is a series in \widehat{e}_γ whose coefficients are rational functions in $q^{1/2}$, possibly with poles as $q^n = 1$ for some $n \geq 1$. The numerical Donaldson-Thomas invariants should be obtained from the motivic ones by taking a semi-classical limit, which means taking $q^{1/2} \rightarrow -1$. The "integer quantum torus", i.e. the ring $\mathcal{R} = \bigoplus_{\gamma \in \Gamma \cap C_0(V)} \mathbb{Z}[q^{\pm 1/2}] \widehat{e}_\gamma$, has a semi-classical limit which is a Poisson algebra with basis $\{e_\gamma\}_{\gamma \in \Gamma \cap C_0(V)}$ and relations $e_\gamma e_\mu = (-1)^{\langle \gamma, \mu \rangle} e_{\gamma + \mu}$. The Poisson bracket is $\{e_\gamma, e_\mu\} = (-1)^{\langle \gamma, \mu \rangle} \langle \gamma, \mu \rangle e_{\gamma + \mu}$.

Conjecture 9.2.3 (Kontsevich-Soibelman). *The automorphism $Ad(A_{V, q})$ of \mathcal{R} preserves the subring*

$$\prod_{\gamma \in \Gamma \cap C_0(V)} \mathbb{Z}[q^{\pm 1/2}] \cdot \widehat{e}_\gamma$$

This automorphism should then yield a formal symplectic automorphism (i.e. expressed as a not necessarily convergent power series) A_V of $\mathbb{T}_\Gamma = \text{Spec } S_\Gamma$ in the semi-classical limit. There is a bijection of sets between \mathfrak{g}_Γ and $\exp(\mathfrak{g}_\Gamma)$, and this yields elements $a(\gamma) \in \mathfrak{g}_\gamma$ such that $\exp(a) = A_V$, which yields a stability condition, and hence defines the function $\Omega : \Gamma \setminus \{0\} \rightarrow \mathbb{Q}$. The (a priori rational, conjecturally integral for generic Z) number $\Omega(\gamma)$ is the Donaldson-Thomas invariant which counts semi-stable objects of class γ . More explicitly, there exists a decomposition

$$A_V = \overrightarrow{\prod}_{Z(\gamma) \in V} T_\gamma^{\Omega(\gamma)}$$

where the arrow over the product indicates that the ordering of the product is done in clockwise order w.r.t. the rays in the plane $Z(\Gamma) \subset \mathbb{C}$, and T_γ denotes the exponentiated

⁵This means it is a cone which does not contain a straight line. Also, the origin is removed.

infinitesimal Poisson automorphism

$$T_\gamma = \exp(-\{\sum_{k \geq 1} \frac{e_\gamma^k}{k^2}, \cdot\})$$

which acts as $T_\gamma e_\mu = (1 - e_\gamma)^{\langle \gamma, \mu \rangle} e_\mu$.

Definition 9.2.4. The rational number $\Omega(\gamma)$ is called the refined Donaldson-Thomas invariant counting semi-stable objects of class γ in \mathcal{A} w.r.t. the given Bridgeland stability condition.

Even without presenting the main ingredients that go into this definition, the construction of the refined Donaldson-Thomas invariants requires some ingenious combinatorics and algebra. The essential point that we wished to convey with our brief presentation of this construction, is the importance of stability conditions.

9.3 The A - and B -model Donaldson-Thomas Invariants

Finally, then, this brings us full circle. We started our discussion of the Thomas-Yau conjecture in 8.2 by presenting the table in which Thomas outlines the mirror dual objects on a Calabi-Yau threefold, and we explained how this gives rise to the Thomas-Yau conjecture. On either side of mirror symmetry, when looked at on a Calabi-Yau threefold, we should obtain some invariant which counts BPS branes. In full generality, this invariant should not just be the classical Donaldson-Thomas invariants, which are extracted from semi-stable objects in $\text{Coh}(\mathcal{X})$, but the refined Donaldson-Thomas invariants extracted from semi-stable objects in $D^b(\mathcal{X})$ with some Bridgeland stability condition. Granting the Thomas-Yau-Joyce conjecture, we make the following definition.

Definition 9.3.1. Let (\mathcal{X}, ω) be a Calabi-Yau threefold. We define the A -model Donaldson-Thomas invariants DT_A of \mathcal{X} as the refined Donaldson-Thomas invariants of the CY3 category $D^b\text{Fuk}(X, \omega)$, with Bridgeland stability condition given by the Thomas-Yau-Joyce conjecture.

This is somewhat peculiar, for the following reason: the classical Donaldson-Thomas invariants, extracted from the moduli space of sheaves, has no known symplectic analogue. When we move to the derived category (that is: away from the large volume limit), we need a Bridgeland stability condition to define semi-stable objects. Since we knew what the stable objects were on the B -side at the large volume limit, while we did not know this on the A -side, it stands to reason that it should be easier to find Bridgeland stability conditions on $D^b(\mathcal{X})$ than on $D^b\text{Fuk}(X, \omega)$. As it turns out, this is not the case: according the Thomas-Yau-Joyce conjecture, there is a very natural Bridgeland stability condition on the latter category. This does not appear to be the case on $D^b(\mathcal{X})$, since there is still no known Bridgeland stability condition for most Calabi-Yau threefolds. This is related to the fact that the B -side central charge requires instanton correction

9.3. THE A - AND B -MODEL DONALDSON-THOMAS INVARIANTS

terms, see [5]. However, we can make the following definition, if we also grant the homological mirror symmetry conjecture:

Definition 9.3.2. Let (\mathcal{X}, ω) and $(\mathcal{X}^\vee, \omega^\vee)$ be a mirror pair of Calabi-Yau threefolds. Then we define the B -model Donaldson-Thomas invariants DT_B of \mathcal{X} as the A -model Donaldson-Thomas invariants of \mathcal{X}^\vee .

So what goes wrong if we try to make an ansatz for a stability condition on $D^b(\mathcal{X})$? As noted, the charge lattice is the image of the Chern character in $H^{\text{ev}}(X, \mathbb{Q})$, and so there is an obvious choice for a central charge. Namely, we take the standard t -structure on $D^b(\mathcal{X})$ whose heart is $\text{Coh}(\mathcal{X})$. Then we are looking for a stability function $Z : K(\text{Coh}(\mathcal{X})) \rightarrow \mathbb{C}$, which we may try to define by

$$Z(E) = \int_X \exp(i\omega) \text{ch}(E) \sqrt{\text{Td}(X)}$$

We are assuming that the B -field is turned off. Then this is the natural way to pair the charge vector $\text{ch}(E) \sqrt{\text{Td}(X)}$ with the Kähler class to produce a complex number. However, this does not actually produce a stability function. There is a much stronger result, found in [112].

Proposition 9.3.3. [112] *Let \mathcal{X} be a smooth projective variety of dimension ≥ 2 . Then there is no stability function on $K(\text{Coh}(\mathcal{X}))$.*

We sketch the proof, in order to understand what goes wrong when $\dim \mathcal{X} > 1$.

Proof. Since $\dim \mathcal{X} > 1$, there exists a smooth subvariety $\iota : \mathcal{Y} \hookrightarrow \mathcal{X}$ with $\dim \mathcal{Y} = 2$. Suppose that we have a stability functions on $K(\text{Coh}(\mathcal{X}))$. Then $Z \circ \iota_*$ is a stability function on $K(\text{Coh}(\mathcal{Y}))$, and so it suffices to consider the case $\dim \mathcal{X} = 2$. Take a curve $\mathcal{C} \subset \mathcal{X}$ and a divisor D on \mathcal{C} . The proof then uses the fact that we can take D to be of any degree, to show that $Z(\mathcal{O}_x) = 0$ for any $x \in \mathcal{X}$, which contradicts the assumption that Z is a stability function. \square

To conclude, then, the presence of subvarieties of codimension greater than 1 results in the fact that we cannot choose a Bridgeland stability structure on $D^b(\mathcal{X})$ whose t -structure has $\text{Coh}(\mathcal{X})$ as its heart. This is the reason why there is no obvious choice of Bridgeland stability condition on $D^b(\mathcal{X})$, since the most natural way to work with $D^b(\mathcal{X})$ is to work with the heart of the standard t -structure, i.e. $\text{Coh}(\mathcal{X})$. This does not mean that $D^b(\mathcal{X})$ does not admit a Bridgeland stability condition, it simply means that we need to consider a different t -structure. Using the technique known as tilting, one can create new t -structures from old ones, and this method is used extensively in attempts to construct Bridgeland stability conditions on $D^b(\mathcal{X})$, but it is still a work in progress.

On the other hand, it should be noted that the Hermitian-Yang-Mills equation is the large volume limit of the so-called deformed Hermitian-Yang-Mills (dHYM) equation.

9.3. THE A - AND B -MODEL DONALDSON-THOMAS INVARIANTS

Given a triple $(\mathcal{Y} \subseteq \mathcal{X}, E, \nabla)$ with $\dim_{\mathbb{C}} \mathcal{Y} = k$, $E \rightarrow Y$ a line bundle, and $F = F_{\nabla}$, the deformed Hermitian-Yang-Mills equations are

$$\begin{cases} F^{2,0} = 0 \\ \operatorname{Im} e^{i\theta} (\omega + F)^k = 0 \end{cases}$$

To leading order, this is the HYM equation (e.g. substitute $\omega \mapsto c\omega$ with $c \in \mathbb{R}$ and look at the leading order in c), and when moving away from the large volume limit, the correction terms are no longer negligible. The dHYM equations are the equations of motions for B -type BPS branes, as is shown in [67] - at least for line bundles. So when looking for Bridgeland stability conditions on $D^b(\mathcal{X})$, there are some string theoretic ideas which should light the path as well. Indeed, the following conjecture is stated in [113].

Conjecture 9.3.4. *A line bundle $E \rightarrow Y$ admits a metric whose Chern connection is a solution to the dHYM equation if and only if \mathcal{E} is stable in $D^b(\mathcal{X})$ with respect to an appropriate Bridgeland stability condition.*

As we can see, this picture is noticeably less complete than the Thomas-Yau-Joyce conjecture, as it has only been formulated for line bundles thusfar. In conclusion: away from the large volume limit, the B -side is somehow more mysterious than the A -side, at least from this perspective, in spite of the fact that the A -side in the large volume limit required many more technicalities to be addressed (some of which we omitted entirely).

Some final remarks: the reader should not be under the misapprehension that our definition of the B -model Donaldson-Thomas invariants is practical in any way. The original problem was to find a Bridgeland stability condition on $D^b(\mathcal{X})$, and then to understand the category $D^b(\mathcal{X})$ (e.g. finding sheaves which split-generate the category) so that one can calculate the invariants. Using our definition (which rests on two open conjectures), one needs to find a mirror manifold \mathcal{X}^{\vee} and then understand the category $D^b\text{Fuk}(X, \omega)$. The derived Fukaya category is typically much more difficult to understand than the derived category of sheaves, and finding a mirror manifold explicitly is certainly no easy matter.

Finally, we would like to note that in [114], the authors give an explicit mapping

$$\Phi : \operatorname{im} (\operatorname{ch} : K(\mathcal{X}) \rightarrow H^{\text{ev}}(X, \mathbb{Q})) \rightarrow H_3(X^{\vee}, \mathbb{Q})$$

which tells us the homology class of the Lagrangian submanifold that a given coherent sheaf gets mapped to under homological mirror symmetry. Thus, modulo issues of torsion in $H_3(X^{\vee}, \mathbb{Z})$, given $\gamma \in H^{\text{ev}}(X, \mathbb{Q})$, one can argue that $\text{DT}_A(\gamma) := \text{DT}_B(\Phi(\gamma))$.

Chapter 10

Conclusion

Conclusion

To conclude, let us briefly outline what was discussed in this thesis. The overarching story is homological mirror symmetry, and how it can be used to motivate new ideas in symplectic and algebraic geometry. We started by explaining homological mirror symmetry in a way which is (hopefully) accessible to a relatively broad audience. Essentially no understanding of physics is required, and we kept the technicalities which arise in the mathematics to a minimum. We then discussed mirror symmetry for $K3$ surfaces, and homological mirror symmetry for the elliptic curve (albeit not in full generality), as well as aspects of homological mirror symmetry for complex tori of any dimension.

We discussed two ideas which can be motivated by homological mirror symmetry: the Thomas-Yau conjecture, and the $P = W$ conjecture (the latter has been shown to hold in many cases). The elliptic curve, being the simplest Calabi-Yau manifold, also proved to be helpful in understanding both of these conjectures, and we discussed their generalisations to complex tori of higher dimensions.

For the $P = W$ conjecture, we discussed the recent publication [15] which generalises the $P = W$ conjecture to moduli spaces associated to abelian varieties, i.e. a specific class of complex tori, of any dimension.

For the Thomas-Yau conjecture, we discussed how the conjecture holds for certain cohomogeneity one Lagrangian submanifolds. This also led us to conclude that a certain invariant version of the Thomas-Yau conjecture 8.6.7 which can be stated for Calabi-Yau manifolds which are a product of a Calabi-Yau manifold on which the Thomas-Yau conjecture is known to hold, and a complex torus.

Some interesting follow-up questions which remain open:

- A precise formulation of what happens in the semi-flat case, when we have a special Lagrangian torus fibration of a Calabi-Yau threefold for which the torus fibres are flat. In this case, one can also define a T^k -action locally, which may have fixed points in the singular fibres. We can still lift submanifolds in the base space under this action, and investigate their mean curvature flow. The lifted submanifolds will not always be tori anymore.
- A version of 8.6.7 where the total space is instead fibred by complex tori, which is quite different from the semi-flat picture, where the fibres are Lagrangian. This

should allow one to inductively construct more examples of Calabi-Yau manifolds on which the Thomas-Yau conjecture holds.

- A general understanding of mean curvature flow of higher dimensional submanifolds is lacking, which makes it difficult to study examples of the Thomas-Yau conjecture explicitly. However, given an understanding of the mean curvature flow of surfaces in \mathbb{R}^4 , for instance, one could prove the Thomas-Yau conjecture on T^4 and subsequently prove an invariant version of the Thomas-Yau conjecture in higher dimensional Calabi-Yau manifolds, similar to what we did. This would already give quite a large class of examples of Lagrangian submanifolds in Calabi-Yau threefolds for which the Thomas-Yau conjecture holds.
- An understanding of the singularities that develop and how to resolve them. We only treated examples in which such singularities are avoided, because the curve shortening flow in the plane does not result singularities. For higher dimensional submanifolds, one has to develop techniques to resolve singularities.

We also looked at the mean curvature flow of Lagrangian fibrations of $\mathbb{C}\mathbb{P}^n$. The result would lead one to suspect that something similar can be said for toric Kähler-Einstein manifolds, and we noted that [11] proves that this is indeed the case. It also leads one to wonder about almost toric Calabi-Yau twofolds, which are T^4 and the $K3$ surface. We answered the question of what happens in the case of T^4 , but an almost toric fibration of a $K3$ surface is not yet understood. This is probably the next simplest example that can be explored, and the close approximations for Ricci flat Kähler metrics on elliptic $K3$ surfaces such as the quartic may allow one to do this.

We commented on how the Thomas-Yau-Joyce conjecture, if true, may be combined with homological mirror symmetry to define a Bridgeland stability condition on $D^b(\mathcal{X})$. Since the existence of such stability conditions is, in general, an open problem, this may be worth investigating, as there is evidence in favour of the Thomas-Yau-Joyce conjecture. On the other hand, there is also progress being made towards proving the existence of Bridgeland stability conditions on $D^b(\mathcal{X})$, which is likely nearer to completion than proving the Thomas-Yau-Joyce conjecture.

Finally, we commented on how the Thomas-Yau-Joyce conjecture may be used to define the Donaldson-Thomas invariants of a Calabi-Yau threefold in full generality, by using the machinery from Kontsevich-Soibelman. This answers a question (which we were certainly not the first to address) raised by Thomas in his paper which motivated the Thomas-Yau conjecture - finding an analogue of Donaldson-Thomas invariants for the A -model. It seems poetic that a refinement of his conjecture will, if true, lead to the correct definition of these invariants.

Final Remarks

In closing, I would like to thank both of my supervisors, dr. N. Martynchuk and prof. M. Kool, for agreeing to supervise this master's thesis. I had decided that I wanted to write my thesis about mirror symmetry well before I started looking for supervisors, and so the task became to find two people people who were willing to supervise the project that I envisioned, rather than having my supervisor(s) suggest a topic to me. This proved somewhat challenging, so I am very thankful to both of them - for agreeing to supervise a master's thesis about a topic that, at least partially, falls outside of their area of expertise, as well as for many helpful suggestions along the way.

Homological mirror symmetry, as well as the classical (closed string) version of mirror symmetry are both very rich subjects. It was at times difficult to select the right topics and examples to include, that would not stray too far from the main narrative, whilst also presenting the beauty of the subject. It is my hope that the right balance was struck, and that readers of this text come to appreciate that beauty. May it serve as a motivation to explore all the other facets of (homological) mirror symmetry, of which there are many.

Part III

Appendices

Appendix A

Appendix: Differential Geometry

Here, we will recall some elementary notions from differential geometry.

A.1 Vector Bundles and Connections

We denote by \mathbb{K} either \mathbb{R} or \mathbb{C} .

Definition A.1.1. A vector bundle on a smooth manifold X is a smooth submersion $\pi : E \rightarrow X$ such that:

1. Every $x \in X$ has an open neighbourhood U and a diffeomorphism $\Phi : \pi^{-1}(U) \rightarrow U \times \mathbb{K}^r$.
2. The restriction $\Phi|_x : \pi^{-1}(\{x\}) \rightarrow \{x\} \times \mathbb{K}^r$ is a linear isomorphism.
3. $\pi_U \circ \Phi = \pi$, where $\pi_U : U \times \mathbb{K}^r \rightarrow U$ is the projection.

We denote the fibre of E at x by E_x . A morphism of vector bundles is a smooth map $\varphi : E_1 \rightarrow E_2$ such that $\pi_1 = \pi_2 \circ \varphi$ and $\varphi_x : E_{1,x} \rightarrow E_{2,x}$ is a linear transformation for all $x \in X$. If U_α and U_β both trivialise E , then there is a function $\varphi_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow \text{GL}(r, \mathbb{K})$, which is defined by

$$\Phi_\beta \circ \Phi_\alpha^{-1}(x, e) = (x, \varphi_{\alpha\beta}(x)e)$$

This function $\varphi_{\alpha\beta}$ is called the transition function, and it satisfies the cocycle condition on triple intersections $U_{\alpha\beta\gamma}$

$$\varphi_{\alpha\beta} \circ \varphi_{\beta\gamma} = \varphi_{\alpha\gamma}$$

A vector bundle can be equivalently defined by an open cover $\{U_\alpha\}$ of X , together with functions $\varphi_{\alpha\beta}$ satisfying the cocycle condition on triple intersections, as well as $\varphi_{\alpha\alpha} = 1$. In practice, this is often the most convenient way to define a vector bundle.

Example 34. The tangent bundle is defined by taken an open cover of X by coordinate neighbourhoods, and defining $\varphi_{\alpha\beta} := D(\psi_{\alpha\beta})$ where $\psi_{\alpha\beta} : \psi_\alpha(U_\alpha) \rightarrow \mathbb{R}^n$ is the transition function associated to the chart maps ψ_α and ψ_β . The fibre $T_x X$ can be identified with $\text{Der}_{\mathbb{R}}(C_x^\infty)$, the \mathbb{R} -linear derivations of the germ at x of the sheaf of smooth functions on X (see the corresponding appendix B).

A.1. VECTOR BUNDLES AND CONNECTIONS

The most important examples of vector bundles are the trivial bundle $X \times \mathbb{K}$ and the tangent bundle TX . We can obtain many other vector bundles from these, by extending the linear algebraic operations to vector bundles. In particular, suppose that $F \subset E$, E_1 and E_2 are vector bundles over X . Then we can form the following vector bundles over X :

1. The direct sum $E_1 \oplus E_2$
2. The tensor product $E_1 \otimes E_2$
3. The exterior powers $\wedge^k E$
4. The symmetric powers $\text{Sym}^k(E)$
5. The dual E^*
6. The homomorphism bundle $\text{Hom}(E_1, E_2)$
7. The quotient bundle E/F

Furthermore, if $f : X \rightarrow Y$ is a smooth map and E is a vector bundle on Y , then we may define the pullback bundle $f^*E \rightarrow X$. This bundle is characterised by a pullback diagram:

$$\begin{array}{ccc} f^*E & \longrightarrow & E \\ \downarrow & & \downarrow \pi \\ X & \xrightarrow{f} & Y \end{array}$$

As such, $f^*E = \{(x, e) \in X \times E \mid f(x) = \pi(e)\}$.

Importantly, a morphism $\varphi : E_1 \rightarrow E_2$ does not have a kernel, image, or cokernel which is a vector bundle, unless φ has constant rank.

Definition A.1.2. A section of a vector bundle $E \rightarrow X$ is a smooth map $s : X \rightarrow E$ such that $\pi \circ s = \text{id}_X$.

This means that for each $x \in X$, we get a vector $s(x) \in E_x$. For example, a morphism of vector bundles is just a section of $\text{Hom}(E_1, E_2)$. Crucially, every smooth vector bundle admits non-zero sections. This is easily proven using cutoff functions. As such, the analogous argument fails for holomorphic bundles on complex manifolds, and indeed, many holomorphic bundles do not have non-trivial global sections. Sections of $X \times \mathbb{R}$ are just smooth functions, while sections of TX are vector fields, and sections of $\wedge^k T^*X$ are differential k -forms.

We denote the space of sections of E over $U \subseteq X$ by $\Gamma(U, E)$. This is naturally an infinite dimensional \mathbb{K} -vector space. In fact, it is a module over $\Gamma(U, U \times \mathbb{R})$ in the obvious way, by setting $(f \cdot s)(x) = f(x) \cdot s(x)$.

Exceptions to the above notation $\Gamma(-, E)$ are the vector bundles $X \times \mathbb{R}$ and $\wedge^k T^*X$. We denote their sections by $C^\infty(U)$ and $\Omega^k(U)$, respectively.

A.1. VECTOR BUNDLES AND CONNECTIONS

The trivial vector bundle $X \times \mathbb{K}$ has a canonical way to differentiate sections. Sections of this bundle are just smooth functions on X , and we use the exterior derivative to differentiate these functions. More generally, we can do this for $X \times \mathbb{K}^r$. However, for arbitrary vector bundles, there is no canonical way to differentiate sections. Nevertheless, we can still develop exterior calculus for vector bundles. This is the role of a connection on a vector bundle. First, we note that the exterior derivative is characterised by the following properties:

1. $d(f \cdot g) = g \cdot df + f \cdot dg$ for $f, g \in C^\infty(X)$.
2. $d(a \cdot f) = a \cdot df$ for $a \in \mathbb{R}$, $f \in C^\infty(X)$.

These are the natural properties that we would like a derivative to satisfy: linearity, and the Leibniz rule. We denote by $\Omega^k(X, E) := \Gamma(X, \wedge^k T^*X \otimes E)$.

Definition A.1.3. A connection ∇ on a vector bundle E is a linear map $\nabla : \Gamma(X, E) \rightarrow \Omega^1(X, E)$ such that for all $f \in C^\infty(X)$ and $s \in \Gamma(X, E)$,

$$\nabla(f \cdot s) = df \otimes s + f \cdot \nabla s$$

This should be interpreted as follows. Given $\nabla s \in \Omega^1(X, E)$, we can take a vector field $v \in \Gamma(X, TX)$. Then we should interpret $\nabla_v s$ as being the derivative of s along v , with respect to the connection ∇ . The expression $\nabla_v s$ is defined as follows, in local coordinates. We have $\nabla s \in \Omega^1(X, E)$, so locally $\nabla s|_U = \sum_i \eta_i \otimes e_i$ for some $\eta_i \in \Gamma(U, T^*U)$ and $e_i \in \Gamma(U, E)$. We define $\nabla_v s|_U$ as $\sum_i \eta_i(v) \cdot e_i \in \Gamma(U, E)$.

Using a partition of unity, it is easy to show that any vector bundle admits a connection. Furthermore, if $\varphi \in \Omega^1(X, \text{End}(E))$, then $\nabla + \varphi$ defines another connection on E . Locally, we have $E|_U \cong U \times \mathbb{K}^r$, which has the canonical connection d . Hence, any connection ∇ on E may be restricted to $E|_U$, and can then be written as $\nabla = d + A$ for some $A \in \Omega^1(U, \text{End}(E))$. This endomorphism-valued 1-form is called the local connection 1-form. If we have two trivialisations over U_α and U_β with $U_\alpha \cap U_\beta$ non-empty, then these 1-forms are related by the transformation rule

$$A_\beta = \varphi_{\alpha\beta}^{-1} \circ A_\alpha \circ \varphi_{\alpha\beta} + \varphi_{\alpha\beta}^{-1} d\varphi_{\alpha\beta}$$

where $\varphi_{\alpha\beta} : U_{\alpha\beta} \rightarrow \text{GL}(r, \mathbb{K})$ is the transition function.

Like the exterior derivative, a connection can be extended to the exterior covariant derivative on higher degree forms. Recall that the exterior derivative on k -forms is defined as

$$\begin{aligned} d\eta(v_0, \dots, v_k) &:= \sum_i (-1)^i v_i \eta(v_0, \dots, \hat{v}_i, \dots, v_k) + \\ &\sum_{i < j} (-1)^{i+j} \eta([v_i, v_j], v_0, \dots, \hat{v}_i, \dots, \hat{v}_j, \dots, v_k) \end{aligned}$$

A.1. VECTOR BUNDLES AND CONNECTIONS

Definition A.1.4. Let (E, ∇) be a vector bundle with a connection. Then the exterior covariant derivative

$$d_{\nabla} : \Omega^k(X, E) \rightarrow \Omega^{k+1}(X, E)$$

is defined by the same formula as the exterior derivative, for $\eta \in \Omega^k(X, E)$:

$$d_{\nabla}\eta(v_0, \dots, v_k) := \sum_i (-1)^i \nabla_{v_i} \eta(v_0, \dots, \widehat{v}_i, \dots, v_k) + \sum_{i < j} (-1)^{i+j} \eta([v_i, v_j], v_0, \dots, \widehat{v}_i, \dots, \widehat{v}_j, \dots, v_k)$$

We had to use the connection ∇ to prescribe how the section $\eta(v_0, \dots, \widehat{v}_i, \dots, v_k)$ is differentiated along v_i , and this is the only difference.

Observe that $\Omega^*(X, E)$ is a module over $\Omega^*(X)$ in the obvious way. Then d_{∇} can be characterised uniquely by the following axioms.

Theorem A.1.5. *The exterior covariant derivative d_{∇} for a vector bundle with connection (E, ∇) is the unique linear operator $d_{\nabla} : \Omega^*(X, E) \rightarrow \Omega^*(X, E)$ which restricts to ∇ on $\Omega^0(X, E)$, and satisfies*

$$d_{\nabla}(\alpha \wedge \eta) = d\alpha \wedge \eta + (-1)^{|\alpha|} \alpha \wedge d_{\nabla}\eta$$

for $\alpha \in \Omega^k(X)$ and $\eta \in \Omega^l(X, E)$.

Suppose that (E, ∇) , (E_1, ∇^1) and (E_2, ∇^2) are vector bundles with connections on X , and $f : Y \rightarrow X$ is a smooth map. Then we get

1. The dual connection ∇^* on E^* by setting $d\langle s, \sigma \rangle = \langle \nabla s, \sigma \rangle + \langle s, \nabla^* \sigma \rangle$ for $s \in \Gamma(X, E)$, $\sigma \in \Gamma(X, E^*)$ and $\langle \cdot, \cdot \rangle$ the natural pairing.
2. The direct sum connection $\nabla^1 \oplus \nabla^2$ on $E_1 \oplus E_2$ by setting $(\nabla^1 \oplus \nabla^2)(s_1, s_2) = (\nabla^1 s_1, \nabla^2 s_2)$ for $s_i \in \Gamma(X, E_i)$.
3. The tensor product connection $\nabla^1 \otimes \nabla^2$ on $E_1 \otimes E_2$ by setting $(\nabla^1 \otimes \nabla^2)(s_1 \otimes s_2) = (\nabla^1 s_1) \otimes s_2 + s_1 \otimes (\nabla^2 s_2)$ for $s_i \in \Gamma(X, E_i)$ and extending by linearity.
4. The pullback connection $f^* \nabla$ on $f^* E$ by setting $(f^* \nabla)(f^* s) = f^* \nabla s$ and extending by linearity.¹

In particular, this yields a connection on $\text{Hom}(E_1, E_2)$ and $\text{End}(E)$.

A very convenient, and in fact defining property of the exterior derivative $d : \Omega^k(X) \rightarrow \Omega^{k+1}(X)$ is that $d^2 = 0$. This allows us to define de Rham cohomology, a very important invariant of the manifold. The same is not true for d_{∇} .

Theorem A.1.6. *There exists a section $F_{\nabla} \in \Omega^2(X, \text{End}(E))$ such that*

$$d_{\nabla}^2 \eta = F_{\nabla} \wedge \eta$$

¹The module $\Gamma(Y, f^* E)$ is generated over $C^\infty(Y)$ by the sections $f^* s$.

A.2. CHERN CLASSES

This section F_∇ is called the curvature 2-form of ∇ . If $F_\nabla = 0$, then ∇ is called a flat connection, in which case it does define a chain complex. If our connection is obtained as the dual, direct sum, tensor product, or pullback of some existing connection(s), then the curvature forms are related in the following way:

1. $F_{\nabla^*} = -F_\nabla^T$
2. $F_{\nabla^1 \oplus \nabla^2} = F_{\nabla^1} \oplus F_{\nabla^2}$
3. $F_{\nabla^1 \otimes \nabla^2} = F_{\nabla^1} \otimes \text{id}_2 + \text{id}_1 \otimes F_{\nabla^2}$
4. $F_{f^*\nabla} = f^*F_\nabla$

If a local trivialisation over U is given, we get a local connection 1-form A , discussed above. This gives us an explicit local formula for the curvature 2-form.

Theorem A.1.7 (Cartan's Structural Equation). *Let A be the local connection 1-form of a connection ∇ on E . Then the curvature 2-form may be locally written as*

$$F_\nabla = dA + A \wedge A = dA + \frac{1}{2}[A, A]$$

where the wedge product of endomorphism valued forms is defined using the composition of endomorphisms, respectively the Lie bracket on endomorphisms.

Using this theorem, one may easily deduce the transformation rule of the curvature 2-form. In fact, we already know its transformation rule, since it is a 2-form. Nevertheless, it is given by

$$F_\beta = \varphi_{\alpha\beta}^{-1} \circ F_\alpha \circ \varphi_{\alpha\beta}$$

Theorem A.1.8 (Bianchi Identity). *The curvature 2-form satisfies $d_\nabla F_\nabla = 0$, where ∇ also denotes the connection on $\text{End}(E)$ induced by the connection on E .*

On a local trivialisation with connection 1-form A , the exterior covariant derivative on $\text{End}(E)$ is given by $d_\nabla \Phi = d\Phi + A \wedge \Phi$ for $\Phi \in \Omega^k(X, \text{End}(E))$. As such, the Bianchi identity can also be written as $dF = F \wedge A$. Inductively, one can show that $d(F^k) = F^k \wedge A - A \wedge F^k$.

A.2 Chern Classes

We now present the differential geometric approach to Chern classes. We will not cover the Euler or Pontrjagin classes, since we will not be using them. As such, we will be considering complex vector bundles $E \rightarrow X$ of rank r . Suppose that $P \in \text{Inv}(\mathfrak{gl}(r, \mathbb{C}))$, meaning that P is an element of $\text{Sym}^\bullet(\mathfrak{gl}(r, \mathbb{C}))$ which is invariant under the adjoint representation of $\text{GL}(r, \mathbb{C})$.

Theorem A.2.1. *Let $F_\nabla \in \Omega^2(X, \text{End}(E))$ be the curvature 2-form of the connection ∇ . Then*

A.2. CHERN CLASSES

1. The differential form $P(F_\nabla)$ on X is closed.
2. The cohomology class $[P(F_\nabla)] \in H^{ev}(X, \mathbb{C})$ is independent of the choice of connection.
3. The map $\Xi : \text{Inv}(\mathfrak{gl}(r, \mathbb{C})) \rightarrow H^{ev}(X, \mathbb{C})$ defined by $\Xi(P) = [P(F_\nabla)]$ is an algebra homomorphism.

Proof. We give a sketch of the proof. To start with, we note that the invariance of P implies that the local expressions indeed glue to give a differential form $P(F_\nabla)$ on X . Next, one shows that $\text{Inv}(\mathfrak{gl}(r, \mathbb{C}))$ is generated by the trace polynomials, i.e. the polynomials $A \mapsto \text{tr}(A^k)$ for $A \in \mathfrak{gl}(r, \mathbb{C})$. Then we only need to prove the relevant statements for the trace polynomials. Then:

1. Follows from the fact that $d\text{tr}(F_\nabla^k) = \text{tr}(d(F_\nabla^k))$ and applying the Bianchi identity.
2. Is the most involved part, and requires the fact that the space of all connection is an affine vector space modelled on $\Omega^1(X, \text{End}(E))$. Then we can take the line segment between two connections, and using this, one may prove that $\text{tr}(F_{\nabla_0}^k) - \text{tr}(F_{\nabla_1}^k)$ is exact. This is done via the so-called transgression formula

$$\int_0^1 \frac{d}{dt} \text{tr}(F_{\nabla_t}^k) dt = d \int_0^1 k \text{tr}(F_{\nabla_t}^{k-1} \wedge \frac{d}{dt} A_t) dt = \text{tr}(F_{\nabla_0}^k) - \text{tr}(F_{\nabla_1}^k)$$

3. A simple verification. □

The homomorphism Ξ is called the Chern-Weil homomorphism. It allows us to define the Chern classes.

Definition A.2.2. Let $E \rightarrow X$ be a complex vector bundle. Then its Chern class $c(E) \in H^{ev}(X)$ is defined by

$$c(E) = \det\left(I - \frac{F_\nabla}{2\pi i}\right)$$

for any choice of connection ∇ . The k -th Chern class $c_k(E)$ is defined to be the coefficient of t^k in the polynomial

$$\det\left(I - \frac{tF_\nabla}{2\pi i}\right) = \sum_k c_k(E) t^k \in H^{ev}(X)[t]$$

Some further remarkable facts about the Chern classes:

1. The Chern classes are integral, meaning that $c_k(E) \in \text{im}(H^{2k}(X, \mathbb{Z}) \rightarrow H^{2k}(X, \mathbb{R}))$. This is ensured by the normalisation of $2\pi i$, as will be demonstrated in the example below.

A.2. CHERN CLASSES

2. The Chern classes are functorial, so that if $f : X \rightarrow Y$ is a smooth map, we have $f^* c_k(E) = c_k(f^* E)$. This simply follows from the fact that $f^* F_\nabla = F_{f^* \nabla}$.

These characteristic classes can be given by a more explicit formula, using the relation between the determinant, trace and logarithm of matrices. Namely, we have

$$c(E) = \det\left(I - t \frac{F}{2\pi i}\right) = \exp \circ \text{tr} \circ \ln\left(I - t \frac{F}{2\pi i}\right) =$$

$$1 - t \frac{\text{tr}(F)}{2\pi i} +$$

$$t^2 \frac{\text{tr}(F \wedge F) - \text{tr}(F) \wedge \text{tr}(F)}{8\pi^2} +$$

$$t^3 \frac{2\text{tr}(F \wedge F \wedge F) - 3\text{tr}(F) \wedge \text{tr}(F \wedge F) + \text{tr}(F) \wedge \text{tr}(F) \wedge \text{tr}(F)}{48\pi^3 i} + \dots$$

This gives a complete expression for Chern classes on manifolds up to dimension 6, as all the higher degrees would have to vanish. We will give an example in which we compute the first Chern class of a line bundle. For this example, we use the following fact: suppose that $E \rightarrow X$ is a vector bundle, and that $E \hookrightarrow X \times \mathbb{C}^N$ is a sub-bundle. Let $\rho : \mathbb{C}^N \rightarrow E$ denote the projection onto E . Then we get a connection on E by setting $\nabla s = \rho(ds)$ for $s \in \Gamma(X, E)$.

Example 35. Consider the complex projective line $\mathbb{C}\mathbb{P}^1$. It parameterises complex lines in \mathbb{C}^2 . There is a tautological line bundle L over $\mathbb{C}\mathbb{P}^1$, which is obtained as a sub-bundle of $\mathbb{C}\mathbb{P}^1 \times \mathbb{C}^2$. Namely, given $x \in \mathbb{C}\mathbb{P}^1$, we take the fibre of L at x to be the line in \mathbb{C}^2 which corresponds to the point x . This clearly defines a smooth sub-bundle $L \subset \mathbb{C}\mathbb{P}^1 \times \mathbb{C}^2$, which then inherits a connection from the trivial bundle. We consider local coordinates on $U_0 = \{[z_0 : z_1] \in \mathbb{C}\mathbb{P}^1 \mid z_0 \neq 0\}$, and write this local coordinate as $z = z_1/z_0$, so that $L_z = \text{span}_{\mathbb{C}}\{e_1 + ze_2\}$. This yields $d(e_1 + ze_2) = dz \otimes e_2$. Using the standard Hermitian inner product on \mathbb{C}^2 , we project e_2 onto the subspace spanned by $e_1 + ze_2$, which yields

$$\frac{\bar{z}}{1 + |z|^2} (e_1 + ze_2)$$

Under the local trivialisation of L over U_0 given by $([z_0 : z_1], \lambda(e_1 + ze_2)) \mapsto (z, \lambda)$, we have that the constant section 1 corresponds to the section $e_1 + ze_2$ and therefore the local 1-form is given by

$$A = \frac{\bar{z} dz}{(1 + |z|^2)}$$

This is an ordinary complex valued 1-form, so $A \wedge A = 0$. It follows that the local curvature form is

$$F = dA = -\frac{dz \wedge d\bar{z}}{(1 + |z|^2)^2}$$

A.2. CHERN CLASSES

The set U_0 excludes only a single point of $\mathbb{C}\mathbb{P}^1$, namely $[0 : 1]$. Therefore, $\int_{\mathbb{C}\mathbb{P}^1} F = \int_{U_0} F = -2\pi i$. This implies that $\int_{\mathbb{C}\mathbb{P}^1} c_1(L) = -1$, and in particular that $c_1(L) \neq 0$. If L was the trivial bundle, then we would have the canonical connection on it which is obviously flat, and therefore would have vanishing curvature. We conclude that L cannot be trivial, because $c_1(L)$ does not depend on the choice of connection.

This example is typical, in the sense that the first Chern class of a line bundle informs us about the triviality of the bundle. More generally, given a vector bundle $\pi : E \rightarrow X$ over a compact oriented manifold X , we can think of $c_1(E)$ as being Poincaré dual to the zero locus of a generic section. By generic, we mean a section which intersects the zero section transversally. Denoting this zero locus by Z , it defines a submanifold $Z \subset X$ because of the assumption that the intersection is transversal. Then Poincaré duality amounts to the statement that given $\alpha \in H^{n-2}(X, \mathbb{R})$, we have

$$\int_Z \alpha = \int_X c_1(E) \wedge \alpha$$

Remark A.2.3. Recall that Poincaré duality asserts that the map

$$H^k(X, \mathbb{R}) \otimes H^{n-k}(X, \mathbb{R}) \rightarrow \mathbb{R} \quad (\alpha, \beta) \mapsto \int_X \alpha \wedge \beta$$

is a perfect pairing, so that $H^{n-k}(X, \mathbb{R}) \cong H^k(X, \mathbb{R})^* = H_k(X, \mathbb{R})$. The Poincaré dual of an element $\alpha \in H^{n-k}(X, \mathbb{R})$ is the cycle $D \in H_k(X, \mathbb{R})$ such that

$$\int_X \alpha \wedge \beta = \int_D \beta$$

The first Chern class for line bundles comes very close to being a full classification of line bundles. To see this, we use some notions from the appendix B. Denote by $C_{\mathbb{C}}^{\infty}$ the sheaf of smooth complex valued functions on a smooth manifold X and by $(C_{\mathbb{C}}^{\infty})^{\times}$ the sheaf of invertible smooth complex valued functions on the same manifold. Then by local invertibility of the exponential map, via the complex logarithm, we have an exact sequence of sheaves

$$0 \rightarrow \mathbb{Z} \rightarrow C_{\mathbb{C}}^{\infty} \xrightarrow{\exp(2\pi i \cdot)} (C_{\mathbb{C}}^{\infty})^{\times} \rightarrow 0$$

Clearly, $C_{\mathbb{C}}^{\infty}$ is a sheaf of C^{∞} -modules, and therefore it is a soft sheaf, so its higher cohomology groups vanish. Consequently, the long exact sequence in cohomology reads

$$\dots \rightarrow 0 \rightarrow H^1(X, (C_{\mathbb{C}}^{\infty})^{\times}) \xrightarrow{\delta} H^2(X, \mathbb{Z}) \rightarrow 0 \rightarrow H^2(X, (C_{\mathbb{C}}^{\infty})^{\times}) \rightarrow \dots$$

This means that the connecting homomorphism δ is in fact an isomorphism. However, the sheaf cohomology group $H^1(X, (C_{\mathbb{C}}^{\infty})^{\times})$ classifies complex line bundles, by sending a line bundle to its collection of transition functions, which are cocycles. So it follows that the image of δ also classifies complex line bundles. The following theorem tells us that the first Chern class classifies line bundles modulo torsion.

A.2. CHERN CLASSES

Theorem A.2.4. [20] *The image of an isomorphism class of line bundles L under δ is $\delta(L) := c_1(L)_{\mathbb{Z}}$ and $c_1(L) = \iota_* c_1(L)_{\mathbb{Z}}$ under the natural inclusion of sheaves $\iota: \mathbb{Z} \rightarrow \mathbb{R}$.*

In fact, the higher Chern classes could be defined axiomatically from the first Chern class, so this characterisation of the first Chern class in terms of Poincaré duality is probably the most important one to keep in mind. Higher Chern classes can be interpreted as being the locus where k generic sections fail to be linearly independent. In this sense, Chern classes give us some crude approximation to whether a given complex vector bundle is trivial or not.

An important property of Chern classes is that they obey the splitting principle. We will use this property later in some calculations. It says that given a short exact sequence of vector bundles

$$0 \rightarrow E' \rightarrow E \rightarrow Q \rightarrow 0$$

we have that $c(E) = c(E') \wedge c(Q)$. It is not difficult to prove this, using that every short exact sequence of vector bundles splits, i.e. $E \cong E' \oplus Q$. It then suffices that we prove $c(E' \oplus Q) = c(E') \wedge c(Q)$. However, the Chern class is defined in terms of the determinant, which obeys the analogous property $\det(A_1 \oplus A_2) = \det(A_1) \det(A_2)$. The result then follows from the fact that $F_{\nabla_1 \oplus \nabla_2} = F_{\nabla_1} \oplus F_{\nabla_2}$, which is a straightforward consequence of the definition of the product connection.

We also have a characteristic class which is additive with respect to exact sequences. This is called the Chern character.

Definition A.2.5. The Chern character of $E \rightarrow X$ is defined as

$$\text{ch}(E) = \text{tr} \exp\left(\frac{-F_{\nabla}}{2\pi i}\right)$$

The k -th Chern character is the degree $2k$ part of $\text{ch}(E) \in H^*(X)$.

The Chern character is so-named because of the fact that

1. $\text{ch}(E_1 \oplus E_2) = \text{ch}(E_1) + \text{ch}(E_2)$
2. $\text{ch}(E_1 \otimes E_2) = \text{ch}(E_1) \wedge \text{ch}(E_2)$

The Chern character may be written in terms of the Chern classes as

$$\text{ch}(E) = \text{rank}(E) + c_1(E) + \frac{1}{2}(c_1(E)^2 - 2c_2(E)) + \frac{1}{6}(c_1(E)^3 - 3c_1(E)c_2(E) + 3c_3(E)) + \dots$$

This also shows that $\text{ch}(E) \in H^{\text{ev}}(X, \mathbb{Q})$. In fact, the Chern character determines a ring homomorphism from the topological K -theory of the manifold X , to its rational cohomology. This will let us extend the Chern character to coherent sheaves. We need this to fix certain topological data, so that we can form some geometric objects (moduli spaces) which parameterise isomorphism classes of sheaves with these discrete invariants.

A.3 Holonomy

Let $\pi : E \rightarrow X$ be a vector bundle over a smooth manifold with a connection ∇ . We fix a basepoint $x_0 \in X$, and denote the fibre $V = \pi^{-1}(E_{x_0})$, which is a real or complex vector space. If a path (resp. loop) starts at the base point x_0 , then the path (resp. loop) will be called based.

Definition A.3.1. A section $s \in \Gamma(E)$ is called parallel if $\nabla s = 0$. If $\gamma : I \rightarrow X$ is a based path in X , then a section $s \in \Gamma(\gamma^*E)$ along γ is called parallel over γ if $(\gamma^*\nabla)s = 0$.

Given $v \in V$, there exists a unique parallel section s along γ such that $s(0) = (x_0, v)$.

Definition A.3.2. Let $\gamma : I \rightarrow X$ be a based path in X , and take $v \in V$. Let s be the parallel section along γ with $s(0) = (x_0, v)$. Then the parallel transport of v along γ is defined as

$$\text{Par}_\gamma(v) = s(\gamma(1)) \in T_{\gamma(1)}X$$

When written in local coordinates, the equation for $(\gamma^*\nabla)s = 0$ is a linear ODE, which means that the parallel transport map is in fact a linear map. It is an automorphism of V , because $\text{Par}_{\gamma^{-1}} = \text{Par}_\gamma^{-1}$. Here, γ^{-1} is the inverse parameterisation of γ . One can also show that concatenating loops results in composition of parallel transport operators:

$$\text{Par}_{\gamma_1 * \gamma_2} = \text{Par}_{\gamma_1} \circ \text{Par}_{\gamma_2}$$

Furthermore, the constant path c obviously has $\text{Par}_c = \text{id}$. As a result, the set

$$\text{Hol}(\nabla) := \{\text{Par}_\gamma \mid \gamma \in \Omega X\}$$

acquires the structure of a group. We denote by $\text{Hol}_0(\nabla)$ the subgroup which is obtained by restricting to based loops which are homotopic to the constant loop.

Definition A.3.3. The group $\text{Hol}(\nabla)$ is called the holonomy group of ∇ . The subgroup $\text{Hol}_0(\nabla)$ is called the reduced holonomy group.

Of course, all of this takes place relative to the fixed based point x_0 , but the conjugacy class of the holonomy group is a well-defined object regardless of the basepoint x_0 .

Theorem A.3.4 ([115]). *The holonomy group is a Lie subgroup of $GL(V)$, whose identity component is the reduced holonomy group. In particular, $\text{Hol}_0(\nabla)$ is a normal subgroup of $\text{Hol}(\nabla)$.*

Observe that $\text{Hol}(\nabla) \subseteq GL(V)$ comes equipped with a natural representation on V . We call this the holonomy representation.

Theorem A.3.5 (The Holonomy Principle). [115] *There is a bijective correspondence between parallel sections of E with respect to a connection ∇ , and vectors $v \in V$ which are invariant under the holonomy representation of ∇ .*

A.4. SPIN STRUCTURES

The correspondence is defined using parallel transport. Namely, given $v \in V$, we define a section $s \in \Gamma(E)$ by setting $s(x) = \text{Par}_\gamma(v)$, where γ is a path from x_0 to x . If the vector v is invariant under the holonomy representation, then we readily see that this is well-defined.

Theorem A.3.6. *Let X be of dimension ≥ 2 , and let $\pi : E \rightarrow X$ be a vector bundle. Suppose $G \subset GL(V)$ is a connected Lie subgroup. Then there exists a connection ∇ on E with $\text{Hol}(\nabla) = G$ if and only if the structure group of E can be reduced to G .*

A.4 Spin Structures

For any Riemannian manifold (X, g) , we get a principal $SO(n)$ -bundle, namely its bundle of oriented orthonormal frames. Denote this bundle by $SO(X, g)$. Given a Lie group homomorphism $H \rightarrow G$, we say that a principal G -bundle P lifts to a principal H -bundle if there exists a principal H -bundle which yields P under the homomorphism $\phi : H \rightarrow G$ (e.g. defining the cocycles of P by $\phi(g_{\alpha\beta})$).

Definition A.4.1. Suppose that $SO(X, g)$ admits a lift to a $\text{Spin}(n)$ -bundle, for the double cover $\text{Spin}(n) \rightarrow SO(n)$. Then (X, g) is said to admit a spin structure.

The obstruction to the existence of a spin structure on a manifold is its second Stiefel-Whitney class $w_2(X) \in H^2(X, \mathbb{Z}/2\mathbb{Z})$. When defining the Floer complex and the Fukaya category, we will make assumptions about the existence of these spin structures. We will also mention spinors when discussing the need for Calabi-Yau manifolds in string theory. Given a representation $\rho : \text{Spin}(n) \rightarrow GL(V)$, we get an associated vector bundle $\Sigma = \text{Spin}(X, g) \times_\rho V$. If ρ is a so-called spin representation, we say that the associated bundle is a spinor bundle, and its sections are spinors (or, more accurately, spinor fields). There are complex spin representations and real spin representations, both of which may be classified completely, and they satisfy certain periodicity conditions. Typically, one would consider complex representations. In even dimensions, there are two inequivalent irreducible complex spin representations. Their direct sum yields an associated bundle whose sections we call Dirac spinors, although we will just refer to them as spinors.

The Levi-Civita connection on $SO(X, g)$ induces a connection on $\text{Spin}(X, g)$ via pull-back. This connection is called the spin connection, and we denote it by ∇^S .

Definition A.4.2. A spinor $\psi \in \Gamma(X, \Sigma)$ is called parallel (or covariantly constant) if $\nabla^S \psi = 0$.

For Kähler manifolds, it is known that the existence of a spin structure is equivalent to the existence of a square root of $K_{\mathcal{X}}$, the canonical line bundle. That is, we require the existence of a holomorphic line bundle \mathcal{L} such that $\mathcal{L} \otimes \mathcal{L} \cong K_{\mathcal{X}}$. When \mathcal{X} is Calabi-Yau, we have $K_{\mathcal{X}} \cong \mathcal{O}_{\mathcal{X}}$, which is its own square root. Therefore, Calabi-Yau manifolds always admit a spin structure.

Appendix B

Appendix: Sheaf Theory

B.1 Coherent Sheaves

Sheaves may be defined on any topological space X . Such a space yields a category $\text{Open}(X)$, in which the objects are open sets, and the morphisms are inclusion maps.

Definition B.1.1. A presheaf \mathcal{F} on X is a functor $\text{Open}(X)^{\text{op}} \rightarrow \text{Set}$.

Here, $\text{Open}(X)^{\text{op}}$ is the opposite category of $\text{Open}(X)$. We will generally be interested in refining the category of sets to the category of groups. Then the above definition means that, for every $V \subseteq U \subseteq X$ open in X , we get abelian groups $\mathcal{F}(U)$, $\mathcal{F}(V)$ and a homomorphism $\text{res}_{UV} : \mathcal{F}(U) \rightarrow \mathcal{F}(V)$. Given $s \in \mathcal{F}(U)$, the restriction map is denoted $s|_V \in \mathcal{F}(V)$.

Definition B.1.2. A presheaf \mathcal{F} on X is called a sheaf if the following are satisfied for every $U \subseteq X$, with $\{U_i\}$ and open cover for U :

1. If $s_i \in \mathcal{F}(U_i)$ such that $s_i|_{U_j} = s_j|_{U_i}$ for all i, j , then there exists $s \in \mathcal{F}(U)$ such that $s|_{U_i} = s_i$.
2. If $s, t \in \mathcal{F}(U)$ are such that $s|_{U_i} = t|_{U_i}$ for all i , then $s = t$.

These axioms encode the familiar notions of gluing, and locality. We can also have sheaves of rings, and sheaves of modules over a sheaf of rings. The restriction maps are then required to be compatible with the various multiplication maps.

Definition B.1.3. Let \mathcal{F}, \mathcal{G} be sheaves on X . Then a morphism $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ of presheaves consists of a morphism $\varphi_U : \mathcal{F}(U) \rightarrow \mathcal{G}(U)$ for each $U \subseteq X$ open, such that the following diagram commutes for all $V \subseteq U \subseteq X$:

$$\begin{array}{ccc} \mathcal{F}(U) & \xrightarrow{\varphi_U} & \mathcal{G}(U) \\ \text{res}_{UV} \downarrow & & \downarrow \text{res}_{UV} \\ \mathcal{F}(V) & \xrightarrow{\varphi_V} & \mathcal{G}(V) \end{array}$$

A morphism of sheaves is just a morphism of presheaves between sheaves. Thus, sheaves on X form a category denoted by $\text{Sh}(X)$.

Morphisms of sheaves (of abelian groups) can be added. As such, we get a sheaf $\text{Hom}(\mathcal{F}, \mathcal{G})$ on X by taking

$$U \mapsto \text{Hom}(\mathcal{F}|_U, \mathcal{G}|_U)$$

B.1. COHERENT SHEAVES

Not every presheaf is a sheaf. The simplest counterexample is when $X = \{x, y\}$ consists of two points with the discrete topology. We define a sheaf by $\mathcal{F}(U) = \mathbb{Z}$ for each open subset. Then the first sheaf axiom fails. Indeed, suppose we are given $s \in \mathcal{F}(\{x\})$ and $t \in \mathcal{F}(\{y\})$. Clearly, they agree on the intersection, since the intersection is empty. Hence they should glue to give a section in $\mathcal{F}(X)$. However, this is evidently impossible, since no such section could restrict to s on $\{x\}$ and to t on $\{y\}$ if $s \neq t$. The idea is that the presheaf does not encode the locality of the topological space. This can be corrected by considering the direct limit $\mathcal{F}_x = \varinjlim_{x \in U} \mathcal{F}(U)$. The abelian group \mathcal{F}_x is called the stalk of \mathcal{F} at x , and it allows us to sheafify a presheaf. Elements of \mathcal{F}_x may be seen as equivalence classes (U, s) , with $s \in \mathcal{F}(U)$. The equivalence relation is that

$$(U_1, s_1) \sim (U_2, s_2) \iff \exists V \subseteq U_1 \cap U_2 \mid s_1|_V = s_2|_V$$

Let us use the stalk to give the definition of sheafification. Let $\widetilde{\mathcal{F}}$ be a presheaf on X . We will say that property (\star) holds for $(s_x) \in \prod_{x \in U} \widetilde{\mathcal{F}}_x$ if, for each $x \in U$, there exists an open neighbourhood $x \in V \subseteq U$ and a section $s \in \widetilde{\mathcal{F}}(V)$ such that $s_x = (V, s)$ for all $x \in V$.

Definition B.1.4. Define a sheaf \mathcal{F} by

$$\mathcal{F}(U) = \{(s_x) \in \prod_{x \in U} \widetilde{\mathcal{F}}_x \mid (s_x) \text{ has property } (\star)\}$$

Then \mathcal{F} is called the sheafification of $\widetilde{\mathcal{F}}$.

The sheafification satisfies the universal property that any morphism from $\widetilde{\mathcal{F}}$ to a sheaf (not just a presheaf) factors through the sheafification. Applying the sheafification to the constant sheaf from above yields a sheaf of locally constant sections. That is, because $X = \{x, y\}$ has two components, we would get $\mathcal{F}(X) = \mathbb{Z} \oplus \mathbb{Z}$, which each summand corresponding to a component.

A very important notion is the exactness of a sequence of sheaves.

Definition B.1.5. A sequence of sheaves (of abelian groups) $\cdots \rightarrow \mathcal{F}_k \xrightarrow{d_k} \mathcal{F}_{k+1} \rightarrow \cdots$ is called exact if the induced sequence at the level of stalks is an exact sequence of abelian groups, for all $x \in X$:

$$\cdots (\mathcal{F}_k)_x \xrightarrow{d_k} (\mathcal{F}_{k+1})_x \rightarrow \cdots$$

It is very important to appreciate this definition. It is distinctly **not** the same as asking that $\mathcal{F}_k(U) \rightarrow \mathcal{F}_{k+1}(U)$ is exact for all $U \subseteq X$ open. This would include $U = X$, which is not the right notion of exactness, as illustrated by 39. Rather, this definition requires that the sequence becomes exact once we shrink to a sufficiently small open neighbourhood of each point. This is why lemma's such as the Poincaré lemma are so important when defining cohomology: it tells us that the sequence

$$0 \rightarrow C^\infty \xrightarrow{d} \Omega^1 \xrightarrow{d} \cdots$$

B.1. COHERENT SHEAVES

is an exact sequence of sheaves, because every smooth manifold is locally just \mathbb{R}^n , for which the Poincaré lemma holds. Thus, the Poincaré lemma may be rephrased as stating that the above sequence of sheaves is exact. They are equivalent statements.

If $f : X \rightarrow Y$ is a continuous map, then we may push a sheaf \mathcal{F} on X forward along f . This gives a sheaf $f_*\mathcal{F}$ on Y defined by $U \mapsto \mathcal{F}(f^{-1}(U))$. Conversely, if \mathcal{F} is a sheaf on Y , we may take the inverse image sheaf $f^{-1}\mathcal{F}$ by taking $(f^{-1}\mathcal{F})(U) := \varprojlim_{f(U) \subset V} \mathcal{F}(V)$. These operations are adjoint to each other:

$$\mathrm{Hom}_{\mathrm{Sh}(X)}(f^{-1}\mathcal{F}, \mathcal{G}) \cong \mathrm{Hom}_{\mathrm{Sh}(Y)}(\mathcal{F}, f_*\mathcal{G})$$

if \mathcal{F} and \mathcal{G} are sheaves of modules on the respective spaces.

Definition B.1.6. A locally ringed space is a pair (X, \mathcal{O}_X) where X is a topological space, and \mathcal{O}_X is a sheaf of rings on X such that $\mathcal{O}_{X,x}$ is a local ring for all $x \in X$.

The most important examples of the above definition are (\mathbb{R}^n, C^∞) , $(\mathbb{D}^n, C^{\mathrm{an}})$ ¹ and $(\mathrm{Spec}(R), \mathcal{O}_X)$. In case of the latter, the sheaf \mathcal{O}_X is defined on a basis for the Zariski topology of $\mathrm{Spec}(R)$, by stipulating that $\mathcal{O}_X(D(f)) = R_f$, where R_f denotes localisation at f .

Definition B.1.7. A morphism of ringed spaces $(X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ consists of a continuous map $f : X \rightarrow Y$ together with a morphism of sheaves $\varphi : f^{-1}\mathcal{O}_Y \rightarrow \mathcal{O}_X$. For us, a section $s \in \mathcal{O}_Y(U)$ will be a function (smooth, analytic or algebraic depending on the context) $s : U \rightarrow \mathbb{K}$, and we will always assume that φ is given by $s \mapsto s \circ f|_{f^{-1}(U)}$.

Definition B.1.8. Let (X, \mathcal{O}_X) be a ringed space.

1. If (X, \mathcal{O}_X) is locally isomorphic to (\mathbb{R}^n, C^∞) , then (X, \mathcal{O}_X) is called a smooth manifold.
2. If (X, \mathcal{O}_X) is locally isomorphic to $(\mathbb{D}^n, C^{\mathrm{an}})$, then (X, \mathcal{O}_X) is called a complex manifold.
3. If (X, \mathcal{O}_X) is locally isomorphic to $(\mathrm{Spec}(R), \mathcal{O}_X)$ for some ring R (the ring R could vary depending on $x \in X$), then (X, \mathcal{O}_X) is called a scheme.

Locally isomorphic means that each $x \in X$ has an open neighbourhood $U \subseteq X$ such that $(U, \mathcal{O}_X|_U)$ is isomorphic to the relevant ringed space.

For us, all schemes will be over \mathbb{C} , which means that R will be a \mathbb{C} -algebra. Moreover, we will often be considering Noetherian schemes, which means that each ideal $I \subseteq R$ is finitely generated over R . When referring to points on schemes, we will generally refer to maximal ideals.

We will use the phrase "ringed space" to refer to any of the above three situations. In each of them, we will be interested in sheaves of \mathcal{O}_X -modules. We want to generalise the

¹We denote the unit polydisc in \mathbb{C}^n by \mathbb{D}^n .

B.1. COHERENT SHEAVES

linear algebraic operations that one can perform on modules, to the sheaves of modules. To do so, we will often require sheafification. This will be implicit in the following list.

Definition B.1.9. Let \mathcal{F} and \mathcal{G} be sheaves of \mathcal{O}_X -modules on a ringed space (X, \mathcal{O}_X) , and let $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ be a morphism of \mathcal{O}_X -modules. Then we define the following sheaves of \mathcal{O}_X -modules on X , by sheafifying if necessary:

1. $\mathcal{F} \oplus \mathcal{G}$ by $U \mapsto \mathcal{F}(U) \oplus \mathcal{G}(U)$
2. $\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{G}$ by $U \mapsto \mathcal{F}(U) \otimes_{\mathcal{O}_X(U)} \mathcal{G}(U)$
3. $\mathcal{H}om(\mathcal{F}, \mathcal{G})$ as defined above
4. $\mathcal{F}^* = \mathcal{H}om(\mathcal{F}, \mathcal{O}_X)$
5. $\ker \varphi$ by $U \mapsto \ker \varphi_U$
6. $\operatorname{im} \varphi$ by $U \mapsto \operatorname{im} \varphi_U$
7. \mathcal{G}/\mathcal{F} by $U \mapsto \mathcal{G}(U)/\varphi(\mathcal{F}(U))$

If $f : X \rightarrow Y$ is a morphism of ringed spaces, we define

$$f^* \mathcal{F} = f^{-1} \mathcal{F} \otimes_{f^{-1} \mathcal{O}_Y} \mathcal{O}_X$$

where the sheaf \mathcal{O}_X becomes an $f^{-1} \mathcal{O}_Y$ -module via the morphism of sheaves $\varphi : f^{-1} \mathcal{O}_Y \rightarrow \mathcal{O}_X$ which is part of the definition of a morphism of ringed spaces.

Thus, sheaves of \mathcal{O}_X -modules form a (monoidal) category. Morphisms of ringed spaces induce functors between the respective categories, either by pushforward or by pullback. These are adjoints:

$$\operatorname{Hom}_{\mathcal{O}_X}(f^* \mathcal{F}, \mathcal{G}) \cong \operatorname{Hom}_{\mathcal{O}_Y}(\mathcal{F}, f_* \mathcal{G})$$

Of particular importance will be those \mathcal{O}_X -modules \mathcal{F} which are locally free. That is, every $x \in X$ has an open neighbourhood U such that $\mathcal{F}|_U \cong \mathcal{O}_X|_U \oplus \cdots \oplus \mathcal{O}_X|_U$.

Proposition B.1.10. *There is a bijective correspondence between locally free \mathcal{O}_X -modules of finite rank, and vector bundles on X .*

If X is a smooth manifold, these vector bundles are smooth. If it is a complex manifold, they are holomorphic. We sketch a proof to aid with the intuition.

Proof. Suppose that $E \rightarrow X$ is a vector bundle. Then every $x \in X$ has a local trivialisation $\Phi : \pi^{-1}(U) \rightarrow U \times \mathbb{K}^r$. Hence $\Gamma(U, E) \cong \oplus_{i=1}^r \mathcal{O}_X(U)$. Thus, the sheaf on X defined by $V \mapsto \Gamma(V, E)$ is locally free. Conversely, suppose that \mathcal{F} is a locally free sheaf. Let $\{U_\alpha\}$ be an open cover for X such that $\mathcal{F}|_{U_\alpha}$ is free for all α . Then for all intersections $U_{\alpha\beta}$, there are morphisms $U_{\alpha\beta} \rightarrow \operatorname{GL}(r, \mathbb{K})$ which serve as transition functions. Hence, \mathcal{F} defines a vector bundle. \square

B.1. COHERENT SHEAVES

From this algebraic perspective on vector bundles, we can also recover the fibres of the vector bundle. If \mathfrak{m}_x is the unique maximal ideal of $\mathcal{O}_{X,x}$, then the fibre of a vector bundle (or more generally a sheaf of \mathcal{O}_X -modules) at x is the finite dimensional vector space $\mathcal{F}_x/\mathfrak{m}_x \cdot \mathcal{F}_x$.

Vector bundles on a space X do not form an abelian category. For our purposes, we will need to address this deficiency, by introducing coherent sheaves. These definitions make sense for an arbitrary ringed space, but we will only apply them for complex manifolds and schemes. We denote a complex manifold or scheme by \mathcal{X} instead of X . This is because we view a complex manifold \mathcal{X} as a smooth manifold X with extra structure.

Definition B.1.11. A coherent sheaf on a scheme \mathcal{X} is a sheaf \mathcal{F} of $\mathcal{O}_{\mathcal{X}}$ -modules such that

1. \mathcal{F} is of finite type: for each $x \in X$ there is an open neighbourhood U of x , and a surjective morphism $\mathcal{O}_{\mathcal{X}}^k|_U \rightarrow \mathcal{F}|_U$ for some $k \in \mathbb{N}$.
2. For any open $U \subseteq \mathcal{X}$ and any natural number n , any morphism $\varphi : \mathcal{O}_{\mathcal{X}}^n|_U \rightarrow \mathcal{F}|_U$ has a kernel of finite type.

For our purposes, we will be working on Noetherian schemes when dealing with coherent sheaves. In this case, there is a better intuition given by the following.

Theorem B.1.12. *Let \mathcal{X} be a projective Noetherian scheme. Then \mathcal{F} is a coherent sheaf if and only if \mathcal{F} is of finite type, and every point in \mathcal{X} has a neighbourhood U and an exact sequence of sheaves*

$$\mathcal{O}_{\mathcal{X}}^n|_U \rightarrow \mathcal{O}_{\mathcal{X}}^m|_U \rightarrow \mathcal{F}|_U \rightarrow 0$$

This tells us that coherent sheaves locally look like the cokernels of morphisms between vector bundles. Coherent sheaves on \mathcal{X} form an abelian category denoted by $\text{Coh}(\mathcal{X})$.

Theorem B.1.13 (Serre). *Let \mathcal{X} be a projective algebraic variety over \mathbb{C} , and let \mathcal{X}^{an} be its analytification. Then there is an equivalence of categories*

$$\text{Coh}(\mathcal{X}) \cong \text{Coh}(\mathcal{X}^{\text{an}})$$

Furthermore, the natural morphism $H^q(\mathcal{X}, \mathcal{F}) \rightarrow H^q(\mathcal{X}^{\text{an}}, \mathcal{F}^{\text{an}})$ is an isomorphism for any coherent sheaf \mathcal{F} .

The analytification of an algebraic variety is the locally ringed space $(\mathcal{X}^{\text{an}}, \mathcal{O}_{\mathcal{X}^{\text{an}}})$ which has maximal ideals as points, but with the usual Euclidean topology (obtained by taking affine opens an embedding them into \mathbb{C}^n , from which the Euclidean topology is inherited), and the sheaf of functions assigns to an open subset the complex analytic functions on it, instead of the algebraic functions. We also recall the following theorem, which mimicks the finite dimensionality of (co)homology on a compact manifold. For

example, the analytification of an algebraic curve is a Riemann surface, and the analytification of a smooth projective variety over \mathbb{C} is a compact complex manifold more generally.

Theorem B.1.14. *For a projective Noetherian \mathbb{C} -scheme \mathcal{X} and a coherent sheaf \mathcal{F} on \mathcal{X} , the cohomology groups $H^q(\mathcal{X}, \mathcal{F})$ are finite dimensional \mathbb{C} -vector spaces.*

B.2 Filtrations of Sheaves

We now define the Jordan-Hölder filtration for coherent sheaves, used to define S -equivalence, as well as the Harder-Narasimhan filtration, which is relevant when considering stability conditions. First recall the polynomials $P(\mathcal{E})$ defined by $m \mapsto \chi(\mathcal{E} \otimes \mathcal{O}(m))$ and $p(\mathcal{E}) = P(\mathcal{E})/\alpha_{\dim \mathcal{E}}(\mathcal{E}) = P(\mathcal{E})/\langle [\mathcal{E}], \omega^d \rangle$. By construction, $p(\mathcal{E})$ is a monic polynomial, allowing us to define partial ordering on the set $\{p(\mathcal{E}) \mid \mathcal{E} \text{ is a sheaf on } X\}$. Then a sheaf \mathcal{E} is called semi-stable if $p(\mathcal{F}) \leq p(\mathcal{E})$ for all proper subsheaves $0 \neq \mathcal{F} \subset \mathcal{E}$, and stable if the inequality is strict.

Definition B.2.1. Let \mathcal{E} be a semi-stable sheaf. A Jordan-Hölder filtration of \mathcal{E} is a sequence

$$0 = \mathcal{E}_0 \subset \mathcal{E}_1 \subset \cdots \subset \mathcal{E}_l = \mathcal{E}$$

such that the subquotients $\text{Gr}_k(\mathcal{E}) = \mathcal{E}_k/\mathcal{E}_{k-1}$ are stable and $p(\text{Gr}_k(\mathcal{E})) = p(\mathcal{E})$ for all k .

We may also require that $\mu(\text{Gr}_k(\mathcal{E})) = \mu(\mathcal{E})$ for all k , depending on the kind of stability we are considering.² One can show that such a filtration always exists, and that $\text{Gr}(\mathcal{E}) = \bigoplus_k \text{Gr}_k(\mathcal{E})$ is independent of the chosen filtration.

Definition B.2.2. Let \mathcal{E}_1 and \mathcal{E}_2 be semi-stable sheaves on \mathcal{X} . Then we declare $\mathcal{E}_1 \sim_S \mathcal{E}_2$ if $\text{Gr}(\mathcal{E}_1) \cong \text{Gr}(\mathcal{E}_2)$, and say that they are S -equivalent.

Clearly, if two sheaves are isomorphic then they are S -equivalent. However, the converse is not necessarily true. The next type of filtration that we will use is called the Harder-Narasimhan filtration, which is defined for pure sheaves (see 5.1.1).

Definition B.2.3. Let \mathcal{E} be a pure sheaf of dimension d on \mathcal{X} . A Harder-Narasimhan filtration of \mathcal{E} is a sequence

$$0 = \mathcal{E}_0 \subset \mathcal{E}_1 \subset \cdots \subset \mathcal{E}_l = \mathcal{E}$$

such that $\text{Gr}_k(\mathcal{E})$ are semi-stable of dimension d , and $p(\text{Gr}_1(\mathcal{E})) > \cdots > p(\text{Gr}_l(\mathcal{E}))$.

The same definition can be made, requiring instead that $\mu(\text{Gr}_1(\mathcal{E})) > \cdots > \mu(\text{Gr}_l(\mathcal{E}))$. The version we refer to depends on the notion of stability that we are considering.

Theorem B.2.4. *Every pure sheaf on \mathcal{X} has a unique Harder-Narasimhan filtration.*

²Recall that $\mu(\mathcal{E}) = \frac{1}{\text{rank}(\mathcal{E})} \int_X c_1(\mathcal{E}) \wedge \omega^{n-1}$

Appendix C

Appendix: Floer Theory

C.1 Almost Complex Structures and Holomorphic Curves

First we recap almost complex structures.

Definition C.1.1. An almost complex structure on a vector bundle $E \rightarrow X$ is an automorphism $J : E \rightarrow E$ such that $J^2 = -\text{id}$. If J is an almost complex structure on TX , we call it an almost complex structure on X . An almost complex structure on a symplectic manifold (X, ω) is called compatible with ω if $g(v, w) := \omega(v, Jw)$ defines a Riemannian metric on X .

If we look at an almost complex structure pointwise, we have $J : T_x X := V \rightarrow V$, which extends to a complex linear map $J : V \otimes \mathbb{C} \rightarrow V \otimes \mathbb{C}$. Since $J^2 = -\text{id}$, the eigenvalues of J are $\pm i$. Consequently, we can split $V \otimes \mathbb{C} = V^{1,0} \oplus V^{0,1}$ into the corresponding eigenspaces, which are of equal dimension. Observe that the almost complex structure J induces an almost complex structure on $T_x^* X = V^*$ via pullback. Hence, we get a decomposition

$$V^* \otimes \mathbb{C} = (V^*)^{1,0} \oplus (V^*)^{0,1} = (V^{1,0})^* \oplus (V^{0,1})^*$$

In particular, $(V^*)^{1,0}$ is given by the space of complex linear maps $L : V \rightarrow \mathbb{C}$ w.r.t. J . This decomposition of the cotangent space then gives

$$\wedge^k V^* \otimes \mathbb{C} = \oplus_{p+q=k} \wedge^p (V^*)^{1,0} \otimes_{\mathbb{C}} \wedge^q (V^*)^{0,1}$$

We then make the natural definition: $\wedge^{p,q} V^* = \wedge^p (V^*)^{1,0} \otimes_{\mathbb{C}} \wedge^q (V^*)^{0,1}$. We think of p as the holomorphic index and of q as the anti-holomorphic index, although we note that there is not necessarily a globally defined notion of holomorphicity. There is the following existence result, which guarantees we can always find a (compatible) almost complex structure on a given symplectic manifold.

Theorem C.1.2. *Every symplectic manifold (X, ω) admits a compatible almost complex structure. The space of all compatible almost complex structures is contractible.*

An important class of almost complex manifolds (and in fact complex manifolds, as we shall see later) are orientable surfaces, i.e. orientable smooth manifolds of dimension 2.

Theorem C.1.3. *Let (Σ, g) be an orientable surface, with a Riemannian metric g . Then the symplectic manifold (Σ, ω) admits a compatible almost complex structure, where ω is the volume form of g .*

C.2. THE MASLOV INDEX

Proof. We recall that on an orientable Riemannian manifold, the metric defines the Hodge star operator $\star : \wedge^k TX \rightarrow \wedge^{\dim_{\mathbb{R}} X - k} TX$. When $\dim_{\mathbb{R}} X = 2$, we have $\star : \wedge^1 TX \rightarrow \wedge^1 TX$, and since $\wedge^1 V = V$, we have a bundle automorphism of TX . Furthermore, $\star^2 = (-1)^{k(\dim_{\mathbb{R}} X - k)}$, so that in our case $\star^2 = -\text{id}$. Consequently, every orientable surface admits an almost complex structure, which is compatible with the volume form of g since \star is an isometry. \square

Suppose we find ourselves with an almost complex manifold (X, J) . Then we can consider smooth maps

$$u : (\Sigma, j) \rightarrow (X, J)$$

Because both manifolds have an almost complex structure, we can talk about J -holomorphic maps, in the following sense.

Definition C.1.4. A smooth map $u : \Sigma \rightarrow X$ is called J -holomorphic or pseudo-holomorphic if $du \circ j = J \circ du$. Equivalently, define $\bar{\partial}_J u := du - J \circ du \circ j$. Then u is J -holomorphic if and only if $\bar{\partial}_J u = 0$. The operator $\bar{\partial}_J$ is called the Cauchy-Riemann operator.

When considering maps $\mathbb{C} \rightarrow \mathbb{C}^n$, the operator above is precisely the familiar Cauchy-Riemann operator, which vanishes if and only if the differential of the map intertwines the complex multiplication on the respective tangent spaces.

Definition C.1.5. Let (X, ω) be a symplectic manifold and let $u : \Sigma \rightarrow X$ be a pseudo-holomorphic curve. Then its energy is defined as

$$E(u) := \int_{\Sigma} u^* \omega = \int_{\Sigma} |du|^2$$

C.2 The Maslov Index

We recall the basics of the Lagrangian-Grassmannian of a symplectic vector space, because it gives us the tools we need to define the gradings on the Floer chain complex. By choosing an appropriate basis, we assume that our symplectic vector space is (\mathbb{C}^n, ω) where ω is the standard symplectic form. Then we define the Lagrangian-Grassmannian

$$\text{LGr}(n) := \{L \subseteq \mathbb{C}^n \mid L \text{ is Lagrangian w.r.t. } \omega\}$$

A priori, this is merely a set. However, one readily verifies that the Lie group $U(n)$ acts transitively on $\text{LGr}(n)$, with isotropy group $O(n)$. As such, we can make the identification $\text{LGr}(n) = U(n)/O(n)$, which gives the Lagrangian-Grassmannian the structure of a smooth manifold. Furthermore, we have a principal $O(n)$ -bundle $\pi : U(n) \rightarrow U(n)/O(n)$, which is in particular a fibre bundle over $\text{LGr}(n)$ with fibre $O(n)$. Hence, we get a long exact sequence of homotopy groups

$$\cdots \rightarrow \pi_2(\text{LGr}(n)) \rightarrow \pi_1(O(n)) \rightarrow \pi_1(U(n)) \rightarrow \pi_1(\text{LGr}(n)) \rightarrow 1$$

C.2. THE MASLOV INDEX

As we know, $\pi_1(\mathrm{U}(n)) \cong \mathbb{Z}$ and $\pi_1(\mathrm{O}(n)) \cong \mathbb{Z}/2$. We conclude that the image of $\pi_1(\mathrm{O}(n)) \rightarrow \pi_1(\mathrm{U}(n))$ is trivial, and so exactness yields an isomorphism

$$\pi_1(\mathrm{LGr}(n)) \cong \mathbb{Z}$$

This also holds for the case $n = 2$, although one needs to adapt the argument, because $\pi_1(\mathrm{O}(2)) \cong \mathbb{Z}$. It also holds for the case $n = 1$ because $\mathrm{U}(1)/\mathrm{O}(1) = S^1/\{\pm 1\} \cong S^1$. It can be shown that the fundamental group has a canonical generator.

Definition C.2.1. Let $\gamma : S^1 \rightarrow \mathrm{LGr}(n)$ be a loop. Then its Maslov index $\mu(\gamma)$ is the image of the canonical generator $1 \in \pi_1(S^1)$ under the pushforward $\gamma_* : \pi_1(S^1) \rightarrow \pi_1(\mathrm{LGr}(n))$.

We denote by $\widetilde{\mathrm{LGr}}(n)$ the universal cover of the Lagrangian-Grassmannian. Obviously, these result holds for an arbitrary symplectic vector space $\mathrm{LGr}(V, \omega)$. As such, we have a natural fibre bundle $\mathcal{L} \rightarrow X$, with fibre $\mathcal{L}_x = \mathrm{LGr}(T_x X, \omega_x)$. Every Lagrangian submanifold $L \subseteq X$ defines a natural section $s_L : L \rightarrow \mathcal{L}|_L$, which is also a map $f : L \rightarrow \mathrm{LGr}(n)$. That is, TL is obtained as the pullback of the tautological bundle over $\mathrm{LGr}(n)$, via the map f , because $\mathrm{LGr}(n)$ is the classifying space for Lagrangian sub-bundles of a symplectic vector bundle of rank $2n$.

Definition C.2.2. The Maslov class of a Lagrangian submanifold $L \subseteq X$ is the cohomology class $\mu_L := f^* \mu \in H^1(L, \mathbb{Z})$, where $\mu \in H^1(\mathrm{LGr}(n), \mathbb{Z})$ is the canonical generator.

For the purposes of defining graded Lagrangian submanifolds, we would like to lift $\mathcal{L} \rightarrow X$ to a bundle $\widetilde{\mathcal{L}} \rightarrow X$ whose fibre is $\widetilde{\mathrm{LGr}}(T_x X, \omega_x)$. The Maslov class is the obstruction to lifting the section s_L of $\mathcal{L}|_L$ to a section of $\widetilde{\mathcal{L}}|_L$. We refer to $\widetilde{\mathcal{L}}$ as a Maslov covering. Not every symplectic manifold admits a Maslov cover, in the same way that not every manifold admits a spin structure. This is because we need to lift a principal $\mathrm{Sp}(2n, \mathbb{R})$ -bundle to a $\mathrm{Sp}^\infty(2n, \mathbb{R})$ -bundle, where the latter denotes the group whose elements consist of pairs (A, φ) with $A \in \mathrm{Sp}(2n, \mathbb{R})$ and φ a \mathbb{Z} -equivariant diffeomorphism of $\widetilde{\mathrm{LGr}}(n)$, which lifts the action of A . Therefore, this group is isomorphic to an infinite cyclic cover of $\mathrm{Sp}(2n, \mathbb{R})$. In the way that the second Stiefel-Whitney class is the obstruction to the existence of a lift from a principal $\mathrm{SO}(n)$ -bundle to a $\mathrm{Spin}(n)$ -bundle, the class $2c_1(X) \in H^1(X, \mathbb{Z})$ (viewing TX as a complex vector bundle via a compatible almost complex structure) is the obstruction to the existence of a Maslov covering. To see this, choose a compatible almost complex structure J , which turns TX into a complex vector bundle. Because $\mathrm{U}(n)$ is a deformation retract of $\mathrm{Sp}(2n, \mathbb{R})$, the group $\mathrm{Sp}^\infty(2n, \mathbb{R})$ deformation retracts onto the universal cover of $\mathrm{U}(n)$. As such, there is a canonical bijection between trivialisations of $(\wedge^n TX) \otimes (\wedge^n TX)$, viewed as complex line bundles, and $\mathrm{Sp}^\infty(2n, \mathbb{R})$ -principal bundles. The obstruction to a trivialisation of the latter is, as we described previously, its first Chern class, which is just $2c_1(X)$. The bundle $(\wedge^n TX)^{\otimes 2} = \det_{\mathbb{C}}(TX)^{\otimes 2}$ arises because the universal cover of $\mathrm{U}(n)$ can be described as

$$\{(A, t) \in \mathrm{U}(n) \times \mathbb{R} \mid \det(A)^2 = \exp(2\pi i t)\}$$

Let us take $\Omega \in \Omega^n(X)$ which is nowhere vanishing, and such that $(\Omega \otimes \Omega)_x$ has unit length in $(\wedge^n T_x^* X)^{\otimes 2}$ at each point $x \in X$, with respect to the metric induced by the compatible

C.2. THE MASLOV INDEX

almost complex structure. Then we get a map $\det^2 : \mathcal{L} \rightarrow S^1$, locally defined by

$$\{e_i\} \mapsto \Omega(e_1, \dots, e_n) \otimes \Omega(e_1, \dots, e_n)$$

for an orthonormal frame $\{e_i\}$ of the Lagrangian sub-bundle of TX .

Definition C.2.3. A graded Lagrangian submanifold (L, f) is a Lagrangian submanifold, together with a function $f : L \rightarrow \mathbb{R}$, such that f lifts $\det^2 \circ s_L : L \rightarrow S^1$. Equivalently, it is a choice of lift of $s_L : L \rightarrow \mathcal{L}|_L$ to $\tilde{s}_L : L \rightarrow \tilde{\mathcal{L}}|_L$.

If the Maslov class μ_L vanishes, we can always choose such a lift. The equivalence between these two definitions comes from the fact that, under the assumption $2c_1(X) = 0$, we have

$$\tilde{\mathcal{L}} = \{(l, t) \in \mathcal{L} \times \mathbb{R} \mid \det(l)^2 = \exp(2\pi i t)\}$$

Definition C.2.4. A special Lagrangian submanifold $L \subseteq X$ is a Lagrangian submanifold such that $\det^2 \circ s_L = c \in S^1$ is constant.

Evidently, every special Lagrangian submanifold admits a grading.

Next we discuss the Maslov index. This will be an integer associated to $l_0, l_1 \in \text{LGr}(n)$ such that $l_0 \cap l_1 = \{0\}$ in \mathbb{C}^n , which corresponds to a point $p \in X$ at which two Lagrangian submanifolds L_0 and L_1 intersect transversally (take $l_i = T_x L_i$). We assume that $2c_1(X) = 0$, and $l_0 = \mathbb{R}^n$, and $l_1 = \exp(i\pi c_1)\mathbb{R} \times \dots \times \exp(i\pi c_n)\mathbb{R} \subseteq \mathbb{C}^n$ for some constants $c_j \in (-1, 0]$. Consider the path l_t in LGr given by $l_t = \exp(i\pi t c_1)\mathbb{R} \times \dots \times \exp(i\pi t c_n)\mathbb{R}$. Fix a form $\Omega \in \Omega^n(X)$ such that $\Omega \otimes \Omega$ is nowhere vanishing and has unit length at each point, as before. We also denote by $\Omega_t \in \wedge_{\mathbb{C}}^n T_x^* X \cong \mathbb{C}$ its restriction to the subspace $l_t \subseteq T_x X$, so we get a function $\det^2 : [0, 1] \rightarrow S^1$ via

$$t \mapsto (\Omega_t)^2 \in S^1$$

Choose a lift $g : [0, 1] \rightarrow \mathbb{R}$ of \det^2 , which is always possible since the interval is contractible.

Definition C.2.5. Let (L_0, f_0) and (L_1, f_1) be transversally intersecting graded Lagrangian submanifolds, and take $p \in L_0 \cap L_1$. Let g be as above. The absolute Maslov index of $p \in L_0 \cap L_1$ is defined as

$$I(p) := (f_1(p) - g(1)) - (f_0(p) - g(0)) \in \mathbb{Z}$$

In [28], it is shown that this number is the Maslov index of a loop in $\text{LGr}(n)$. We have a path l_t connecting l_0 to l_1 in $\text{LGr}(n)$. Furthermore, since the Lagrangians are graded, we choose a path in $\widetilde{\text{LGr}}(n)$ which connects the graded lift \tilde{l}_0 to \tilde{l}_1 . This projects to a path in $\text{LGr}(n)$ connecting l_0 and l_1 , let us denote it by $\pi \circ \gamma$. Then the path $l_t^{-1} * (\pi \circ \gamma) : S^1 \rightarrow \text{LGr}(n)$ is canonically associated to the intersection $l_0 \cap l_1$ in $T_x X$, up to homotopy. Then one can equivalently define $I(p) := \mu(l_t^{-1} * (\pi \circ \gamma))$, and verify that these definitions coincide.

C.3. LAGRANGIAN-FLOER COHOMOLOGY

One last incarnation of the Maslov index needs to be constructed, and this time, for homotopy classes of maps $[u] \in \pi_2(X, L)$. Suppose $u : D^2 \rightarrow X$ represents such a class. Then we can consider the pullback bundle u^*TX , which is a symplectic vector bundle over D^2 . Since D^2 is contractible, we can give a symplectic trivialisation for u^*TX . Since $u(\partial D^2) \subseteq L$, where L is some Lagrangian submanifold, we have the sub-bundle

$$u|_{\partial D^2}^* TL \subseteq u^*TX|_{\partial D^2}$$

which is a Lagrangian sub-bundle. The trivialisation of u^*TX then gives us a loop $\gamma : S^1 \rightarrow \text{LGr}(n)$, by taking $t \mapsto (u|_{\partial D^2}^* TL)_t$.

Definition C.2.6. The Maslov index $\text{ind}([u])$ of a homotopy class $[u] \in \pi_2(X, L)$ is the Maslov index of the loop γ . Equivalently, it is the degree of the map $\det^2 \circ \gamma : S^1 \rightarrow S^1$.

C.3 Lagrangian-Floer Cohomology

We are now going to outline a cohomology theory, called Lagrangian-Floer cohomology, which describes the intersection theory of Lagrangian submanifolds on a given symplectic manifold (X, ω) . When constructing this cohomology theory, we run into some technical problems. As such, we will always assume that the intersection $L_0 \cap L_1$ of Lagrangian submanifolds is transversal. In this case, we have that $L_0 \cap L_1$ consists of a (finite, if X is compact) set of points, since both L_0 and L_1 have codimension $\frac{1}{2} \dim X$. Then we set $CF^*(L_0, L_1) := \Lambda \langle L_0 \cap L_1 \rangle$ the Λ -vector space generated by the intersection points, where Λ denotes the Novikov field:

$$\Lambda := \left\{ \sum a_i T^{\lambda_i} \mid a_i \in \mathbb{C}, \quad \lim_{i \rightarrow \infty} \lambda_i = \infty \right\}$$

We wish to turn the vector space $CF^*(L_0, L_1)$ into a chain complex, and so we need to endow it with a grading, and with a differential.

Definition C.3.1. The grading on $CF^*(L_0, L_1)$ is defined by specifying the degrees of the generators $p \in L_0 \cap L_1$ to be

$$\text{deg}(p) := I(p)$$

where $I(p)$ is the Maslov index of p from the previous subsection.

Next, we want to define the differential. In order to do this, we use a compatible almost complex structure on X (see C.1.1). We let (X, J, ω) denote the symplectic manifold together with a regular almost complex structure J which is compatible with the symplectic form. We say that a pseudo-holomorphic disk $u : D^2 \rightarrow (X, J)$ connects p to q if $u(-1) = p$ and $u(1) = q$. By the Riemann mapping theorem, we can also take $D^2 \setminus \{-1, 1\} \cong \mathbb{R} \times [0, 1]$. Then we are interested in maps $u(s, t) : \mathbb{R} \times [0, 1] \rightarrow (X, J)$ such that

1. $\bar{\partial}_J u = 0$
2. $\lim_{s \rightarrow -\infty} u(s, t) = p$ and $\lim_{s \rightarrow \infty} u(s, t) = q$ for all $t \in [0, 1]$

C.3. LAGRANGIAN-FLOER COHOMOLOGY

3. $u(\mathbb{R} \times \{0\}) \subseteq L_0$ and $u(\mathbb{R} \times \{1\}) \subseteq L_1$

4. $\int_{D^2} u^* \omega < \infty$

We denote by $M(p, q, L_0, L_1, [u])$ the space of maps which satisfy these conditions and which represent a homotopy class $[u] \in \pi_2(X, L_0 \cup L_1)$, and

$$\mathcal{M}(p, q, L_0, L_1, [u]) := M(p, q, L_0, L_1, [u]) / \mathbb{R}$$

where the action of \mathbb{R} is given by reparameterisation of the s -coordinate. The space $M(p, q, L_0, L_1, [u])$ is the solution space to a Fredholm problem, i.e. the kernel of the linearisation of $\bar{\partial}_J$, and this operator is Fredholm. It follows that the dimension of $M(p, q, L_0, L_1, [u])$ is given by the index of the operator, which is the Maslov index (see [28]).

Lemma C.3.2. *Let $[u]$ be a pseudo-holomorphic disk connecting p and q as above. Then*

$$\text{ind}([u]) = I(q) - I(p)$$

It follows that

$$\dim \mathcal{M}(p, q, L_0, L_1, [u]) = \text{ind}([u]) - 1 = I(q) - I(p) - 1$$

Assume that the Lagrangians come with spin structures on them, giving the moduli space an orientation. We will assume that the moduli space can be compactified. This requires some technical issues to be addressed, which we will omit. Let us denote by $\overline{\mathcal{M}}(p, q, L_0, L_1, [u])$ its compactification. Then we can count its points with orientations. In this way, one defines the Floer differential ∂ on generators $p \in L_0 \cap L_1$ by setting

$$\partial(p) := \sum_{\substack{q \in L_0 \cap L_1, \\ [u] \in \pi_2(X, L_0 \cup L_1): \\ \text{ind}([u])=1}} \# \overline{\mathcal{M}}(p, q, [u], J) T^{\langle \omega, [u] \rangle} \cdot q$$

where $\#$ denotes the signed count (i.e. with orientations), and $\langle \omega, [u] \rangle := \int_{D^2} u^* \omega$ is the energy of u . Evidently, summing over homotopy classes $[u]$ with $\text{ind}([u])$ has two important implications: the moduli spaces are 0-dimensional, and the Floer differential indeed has degree 1.

Theorem C.3.3. *The Floer differential satisfies $\partial^2 = 0$, i.e. turns $(CF^\bullet(L_0, L_1), \partial)$ into a chain complex*

$$\dots \xrightarrow{\partial} CF^{k-1}(L_0, L_1) \xrightarrow{\partial} CF^k(L_0, L_1) \xrightarrow{\partial} \dots$$

The cohomology of the chain complex is independent of the choice of almost complex structure.

We refer to [28] for the proof, as well as the technical conditions that have to be satisfied for the theorem to hold. In this text, we will simply assume that the theorem above holds whenever we talk about the Floer chain complex (or its cohomology).

C.3. LAGRANGIAN-FLOER COHOMOLOGY

Definition C.3.4. Let L_0 and L_1 be two transversally intersecting Lagrangian submanifolds. Then their Floer cohomology is defined as $HF^\bullet(L_0, L_1) := H^\bullet(CF^\bullet(L_0, L_1), \partial)$.

One observes that the Floer cohomology is invariant under Hamiltonian isotopies. Thus, if an intersection $L_0 \cap L_1$ is not transversal, we can use a time-dependent Hamiltonian to perturb L_1 , say $\varphi(L_1)$, such that the intersection $L_0 \cap \varphi(L_1)$ is transversal. Then we define $HF^\bullet(L_0, L_1) := HF^\bullet(L_0, \varphi(L_1))$. An important special case of this occurs when $L_0 = L_1$. Under suitable conditions (i.e. when L is aspherical), one finds that $HF^\bullet(L, L) \cong H^\bullet(L, \Lambda)$, the ordinary cohomology with coefficients in the Novikov field.

Appendix D

Homological Algebra Preliminaries

This appendix contains the preliminary notions from homological algebra that we need, to rigorously formulate the homological mirror symmetry conjecture. Before we begin, recall the following definitions.

Definition D.0.1. A functor $L : \mathcal{A} \rightarrow \mathcal{A}'$ is called an equivalence of categories if there exists a functor $L' : \mathcal{A}' \rightarrow \mathcal{A}$ such that $L \circ L'$ and $L' \circ L$ are both naturally isomorphic to the identity functor on the respective categories.

Notice that we do not require $L \circ L' = \text{id}_{\mathcal{A}'}$ and $L' \circ L = \text{id}_{\mathcal{A}}$. Put differently, we only require that each object in \mathcal{A}' is isomorphic to some object $F(A)$ (meaning that the functor is essentially surjective), and that $L : \text{Hom}_{\mathcal{A}}(A, B) \rightarrow \text{Hom}_{\mathcal{A}'}(L(A), L(B))$ is an isomorphism. This is a subtle difference, but an important one.

D.1 Abelian categories and cohomology

The starting point for all of homological algebra is the notion of an abelian category, which is an additive category with some extra convenient properties.

Definition D.1.1. An additive category is a category \mathcal{A} such that

1. For all $A, B \in \text{Ob}(\mathcal{A})$, $\text{Hom}_{\mathcal{A}}(A, B)$ is an abelian group
2. Composition of morphisms distributes over the addition:

$$f \circ (g_1 + g_2) = f \circ g_1 + f \circ g_2 \quad (f_1 + f_2) \circ g = f_1 \circ g + f_2 \circ g$$

3. \mathcal{A} has an object Z called the zero object: for every $A \in \text{Ob}(\mathcal{A})$, the groups $\text{Hom}_{\mathcal{A}}(Z, A)$ and $\text{Hom}_{\mathcal{A}}(A, Z)$ are trivial.
4. For every pair $A, B \in \text{Ob}(\mathcal{A})$, there exists a $C \in \text{Ob}(\mathcal{A})$ which is both the sum and the product of A and B .

We also simply write $A \in \mathcal{A}$ instead of $A \in \text{Ob}(\mathcal{A})$.

Of course, as with anything that is presented to us in terms of category theory, it is important to keep in mind examples which inspire the definitions that are made.

Example 36.

1. The category of abelian groups.
2. The category of vector spaces over a given field.
3. The category of modules over a ring
4. The category of (holomorphic) vector bundles over a (complex) manifold.

The final category we listed lacks a crucial property, which is one of the main reasons why we consider coherent sheaves, rather than only holomorphic vector bundles. Namely, the kernel of a morphism between holomorphic vector bundles need not be a vector bundle. Instead, it is a coherent sheaf. This is why we would like more structure than just that of an additive category. We want abelian categories, so that taking kernels and cokernels does not take us outside of the category. This will be important, if we wish to do homological algebra.

Definition D.1.2. An abelian category \mathcal{A} is an additive category such that

1. Every morphism has a kernel and a cokernel
2. Every monomorphism is the kernel of some morphism
3. Every epimorphism is the cokernel of some morphism

The first condition is clear: it states precisely what we wanted, that our category is closed under taking kernels and cokernels of morphisms. The second condition requires that, given a monomorphism $A \rightarrow B$, there exists a morphism $f : B \rightarrow C$ which makes

$$0 \rightarrow A \rightarrow B \xrightarrow{f} C$$

an exact sequence. In fact, we need the existence of kernels of morphisms, to be able to talk about exact sequences in the first place. Of course, the third condition is just the dual implication to the second. As alluded to, this is a crucial difference between the category of vector bundles over a projective complex manifold or scheme, and the category of coherent sheaves over it. They are both additive categories, but in positive dimensions, only the category of coherent sheaves is an abelian category. Besides the category of holomorphic vector bundles, the other three additive categories we listed above are also abelian categories.

Given any abelian category \mathcal{A} , we define $\text{Ch}(\mathcal{A})$ to be the category whose objects are complexes (i.e. $d_{i+1} \circ d_i = 0$) indexed by integers

$$\dots \xrightarrow{d_{i-1}} E_i \xrightarrow{d_i} E_{i+1} \xrightarrow{d_{i+1}} \dots$$

D.1. ABELIAN CATEGORIES AND COHOMOLOGY

with $E_i \in \mathcal{A}$ and $d_i \in \text{Hom}_{\mathcal{A}}(E_i, E_{i+1})$. Objects in this category will be denoted E_{\bullet} and we will usually omit the index from the differential d . A morphism $f : E_{\bullet} \rightarrow F_{\bullet}$ in the category of chain complexes is called a chain map, and it consists of a family of morphisms $f_i : E_i \rightarrow F_i$ such that $d_F \circ f_i = f_{i+1} \circ d_E$. We say that $E_{\bullet} \in \text{Ch}(\mathcal{A})$ is bounded if there exists $k \in \mathbb{N}$ such that $E_i = 0$ for all $|i| \geq k$. This category of chain complexes comes equipped with a natural shift functor, denoted $[n] : \text{Ch}(\mathcal{A}) \rightarrow \text{Ch}(\mathcal{A})$. Rather than $[n]E_{\bullet}$, we write $E_{\bullet}[n]$. The name already defines the operation: it shifts the degrees of the complex by n , so that $E[n]_i = E_{i+n}$ and $d[n]_i = (-1)^n d_{i+n}$. The use of this operation may not be immediately obvious, but it will be of crucial importance later on. Observe that the original category \mathcal{A} embeds into $\text{Ch}(\mathcal{A})$ via

$$E \mapsto E[0] := \cdots \rightarrow 0 \rightarrow E \rightarrow 0 \rightarrow \cdots$$

Given $\text{Ch}(\mathcal{A})$, we can define the cohomology functors

$$H^i : \text{Ch}(\mathcal{A}) \rightarrow \mathcal{A} \quad E_{\bullet} \mapsto \frac{\ker(d_i : E_i \rightarrow E_{i+1})}{\text{im}(d_{i-1} : E_{i-1} \rightarrow E_i)}$$

Definition D.1.3. A complex E_{\bullet} is said to be concentrated in degree i if $H^j(E_{\bullet}, d_E) = 0$ whenever $j \neq i$.

Example 37. For any smooth manifold X , we take $(X \times \mathbb{R}, d)$ to get the de Rham complex

$$0 \rightarrow \Omega^0(X) \xrightarrow{d} \Omega^1(X) \xrightarrow{d} \cdots \xrightarrow{d} \Omega^n(X) \rightarrow 0$$

This is an object in the category of chain complexes of (infinite dimensional) real vector spaces, and its cohomology objects are the de Rham cohomology groups of the manifold. The Dolbeault complex also arises in this way, yielding the Dolbeault cohomology groups. However, the way in which we arrive at these complexes is not quite so straightforward. It is merely the final step in the process of sheaf cohomology, as we will outline below.

Another example is given by the notion of a cohomology sheaf. Suppose we have a complex manifold \mathcal{X} and a complex of coherent sheaves \mathcal{E}_{\bullet} over \mathcal{X} . Then the i -th cohomology sheaf is defined as

$$H^i(\mathcal{E}_{\bullet}) = \frac{\ker(d_i : \mathcal{E}_i \rightarrow \mathcal{E}_{i+1})}{\text{im}(d_{i-1} : \mathcal{E}_{i-1} \rightarrow \mathcal{E}_i)} \in \text{Coh}(\mathcal{X})$$

That is, the cohomology object is itself a sheaf. This idea will be rather important when we discuss the derived category of coherent sheaves later.

D.2 Injective Objects and Sheaf Cohomology

Just as we want the morphisms within a category to preserve certain structures on the objects of that category, so too do we want functors between additive or abelian categories to respect the extra structure on them, which leads us to the following definition.

Definition D.2.1. Let $L : \mathcal{A}_1 \rightarrow \mathcal{A}_2$ be a functor between additive categories. Then L is called additive if, for all $A, B \in \mathcal{A}_1$, the map

$$L : \text{Hom}_{\mathcal{A}_1}(A, B) \rightarrow \text{Hom}_{\mathcal{A}_2}(L(A), L(B))$$

is a homomorphism of abelian groups.

If the Hom-spaces of an additive category are \mathbb{C} -vector spaces, the category is called \mathbb{C} -linear. In this case, a functor is called linear when the same map is a linear map between vector spaces. When we consider abelian categories, the main objects of interests are short exact sequences

$$0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$$

This is because they simultaneously capture the notions of kernels and cokernels, as well as direct sums. Indeed, we recall that a short exact sequence is called split if it is isomorphic to

$$0 \rightarrow A \rightarrow A \oplus B \rightarrow B \rightarrow 0$$

where the maps are the natural maps, i.e. inclusion and projection. So what is the kind of functor that preserves the structure of an abelian category? It would have to be an additive functor which takes exact sequences to exact sequences. If a functor L has this property, we say that L is an exact functor. However, we will often be interested in a functor which only exact from one side, in the following sense: if we have a short exact sequence as above, then the functor L will only guarantee that

$$0 \rightarrow L(A) \rightarrow L(B) \rightarrow L(C)$$

is exact, in which case L is called left exact. Likewise, if it can only guarantee that

$$L(A) \rightarrow L(B) \rightarrow L(C) \rightarrow 0$$

is exact, then L is called right exact. An important case of a (contravariant) left exact functor on an abelian category is the left Hom-functor. Fix $I \in \mathcal{A}$ and define

$$\text{Hom}_{\mathcal{A}}(\cdot, I) : \mathcal{A} \rightarrow \text{Ab} \quad A \mapsto \text{Hom}_{\mathcal{A}}(A, I)$$

We know from algebraic considerations that this is a left exact functor. However, if I is such that $\text{Hom}_{\mathcal{A}}(\cdot, I)$ is an exact functor, then I is called an injective object.

Example 38. Not all objects in an abelian category are exact: taking $\mathcal{A} = \text{Ab}$, let $I = \mathbb{Z}$. Consider the short exact sequence

$$0 \rightarrow \mathbb{Z} \xrightarrow{\cdot 2} \mathbb{Z} \rightarrow \mathbb{Z}/2 \rightarrow 0$$

Applying the functor $\text{Hom}(\cdot, \mathbb{Z})$, we get

$$0 \leftarrow \mathbb{Z} \xleftarrow{\cdot 2} \mathbb{Z} \leftarrow 0 \leftarrow 0$$

This obviously cannot be exact, because multiplication by 2 is not surjective.

Definition D.2.2. An injective resolution of an object $A \in \mathcal{A}$ is an exact sequence

$$0 \rightarrow A \rightarrow F_1 \rightarrow F_2 \rightarrow \dots$$

such that all the $F_i \in \mathcal{A}$ are injective. When every object of \mathcal{A} admits an injective resolution, the category is said to have enough injectives.

It is proved e.g. in [116] that the category of sheaves of abelian groups has enough injectives.

We can now define sheaf cohomology, which is a special case of the notion of a derived functor. The construction is as follows. Given a sheaf \mathcal{F} , we find an injective resolution

$$0 \rightarrow \mathcal{F} \rightarrow \mathcal{F}_1 \rightarrow \dots$$

which we write as $0 \rightarrow \mathcal{F} \rightarrow \mathcal{F}_\bullet$. Then, we apply a left exact functor L (we will use the global sections functor) to get a chain complex (not necessarily an exact sequence)

$$0 \rightarrow \Gamma(X, \mathcal{F}) \rightarrow \Gamma(X, \mathcal{F}_1) \rightarrow \dots = 0 \rightarrow \Gamma(X, \mathcal{F}) \rightarrow \Gamma(X, \mathcal{F}_\bullet)$$

Subsequently, we can apply the cohomology functors, and we make the following definition.

Definition D.2.3. The group $H^i(X, \mathcal{F}) := H^i(\Gamma(X, \mathcal{F}_\bullet))$ is called the i -th sheaf cohomology of \mathcal{F} .

More generally, if $L: \mathcal{A}_1 \rightarrow \mathcal{A}_2$ is a left exact exact functor between abelian categories with enough injectives, then we define its right derived functors $R^i L$ as follows. Let $0 \rightarrow A \rightarrow F_1 \rightarrow \dots$ be an injective resolution so that $0 \rightarrow L(A) \rightarrow L(F_1) \rightarrow \dots$ is a chain complex. Then $R^i L: \mathcal{A}_1 \rightarrow \mathcal{A}_2$ is defined by $R^i L(A) := H^i(L(F_\bullet))$. The name derived functor will make more sense in the context of derived categories, to be discussed later. A basic result from sheaf theory is that this definition is independent of the chosen injective resolution. A proof of this fact, as well as the following proposition, can be found in [116].

D.2. INJECTIVE OBJECTS AND SHEAF COHOMOLOGY

Proposition D.2.4. *Let \mathcal{R} be a sheaf of rings on a smooth manifold X , such that \mathcal{R} admits a partition of unity subordinate to any open cover. Then any sheaf of \mathcal{R} -modules is acyclic, i.e. $H^i(X, \mathcal{F}) = 0$ for $i > 0$.*

In particular, any sheaf of C^∞ -modules on a smooth manifold, such as the sheaf of smooth sections of a vector bundle, is an acyclic sheaf. For computations, one often uses acyclic resolutions, rather than injective resolutions. The resulting cohomology is the same (which is proved in the same reference).

Example 39. In very rough terms, sheaf cohomology is the answer to the following question: given an exact sequence of sheaves on X

$$0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{F}_2 \rightarrow \mathcal{F}_3 \rightarrow 0$$

what is the obstruction to the map $\Gamma(X, \mathcal{F}_2) \rightarrow \Gamma(X, \mathcal{F}_3)$ being surjective? We recall that the exactness of a sequence of sheaves can be checked locally. For example, consider the sequence of sheaves on S^1 , which is an acyclic resolution of the constant sheaf \mathbb{R} :

$$0 \rightarrow \mathbb{R} \rightarrow C^\infty \xrightarrow{d} \Omega^1 \rightarrow 0$$

This sequence is exact, because that is the statement of the Poincaré lemma in dimension 1. However, the map $d : C^\infty(S^1) = \Gamma(S^1, C^\infty) \rightarrow \Gamma(S^1, \Omega^1) = \Omega^1(S^1)$ is evidently not surjective. Indeed, if we view $S^1 = \mathbb{R}/\mathbb{Z}$, then $dx \in \Omega^1(S^1)$ is not the differential of a periodic function. This follows immediately from Stokes's theorem. In fact, dx is a generator for $H^1(S^1) \cong \mathbb{R}$, so we can already see how the language of sheaves, abelian categories and cohomology functors let us connect these more abstract notions to the de Rham cohomology of a manifold. As we know, the cohomology (integer or real coefficients) of a manifold is an important invariant. In fact, the sheaf cohomology of the constant sheaf G for some abelian group G is just $H^k(X, G)$ in the ordinary sense of algebraic topology. The de Rham complex is an acyclic resolution of the constant sheaf \mathbb{R} , since we have an exact sequence of sheaves:

$$0 \rightarrow \mathbb{R} \hookrightarrow C^\infty \xrightarrow{d} \Omega^1 \xrightarrow{d} \Omega^2 \xrightarrow{d} \dots$$

Each sheaf in this resolution is a sheaf of C^∞ -modules, meaning the resolution is indeed acyclic. As such, we see that the sheaf cohomology of \mathbb{R} can be computed via de Rham cohomology. Suppose that $\mathcal{E} \rightarrow \mathcal{X}$ is a holomorphic vector bundle. Then we have an acyclic resolution of sheaves

$$0 \rightarrow \Omega_{\mathcal{X}}^p \otimes \mathcal{E} \hookrightarrow \Omega^{p,0} \otimes \mathcal{E} \xrightarrow{\bar{\partial}} \Omega^{p,1} \otimes \mathcal{E} \xrightarrow{\bar{\partial}} \Omega^{p,2} \otimes \mathcal{E} \xrightarrow{\bar{\partial}} \dots$$

So we see that the Dolbeault cohomology groups from 2.2.5 just compute the sheaf cohomology of the sheaf of holomorphic p -forms with values in \mathcal{E} . This is known as the Dolbeault theorem, which states $H^{p,q}(\mathcal{X}, \mathcal{E}) = H^q(\mathcal{X}, \Omega_{\mathcal{X}}^p \otimes \mathcal{E})$.

So morally, we think of an exact sequence of sheaves $0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{F}_2 \rightarrow \mathcal{F}_3 \rightarrow 0$ as follows: the morphism $\mathcal{F}_2 \rightarrow \mathcal{F}_3$ asks us to solve a certain equation. Exactness of the sequence implies that this equation can be solved around a point, if we restrict our attention to a sufficiently small neighbourhood of each point. Whether global solutions exist, is the question of sheaf cohomology, which we might think of as the obstruction to existence of global solutions.


D.3 Triangulated Categories

Recall that we constructed the category $\text{Ch}(\mathcal{A})$ of chain complexes of an abelian category \mathcal{A} . We noted that there exists a shift functor $[n] : \text{Ch}(\mathcal{A}) \rightarrow \text{Ch}(\mathcal{A})$ which is an equivalence of categories. This is going to play an important role in the definition of a triangulated category (although $\text{Ch}(\mathcal{A})$ will not turn out to be a triangulated category). Observe that $[n] = [1] \circ \dots \circ [1]$ if $n > 0$, and $[n] = [1]^{-1} \circ \dots \circ [1]^{-1}$ if $n < 0$. Thus, the shift functor which gives us the most important information is $[1]$.

Definition D.3.1. Suppose a category \mathcal{C} has a shift functor $[1]$. Then a triangle in \mathcal{C} is a diagram

$$A \rightarrow B \rightarrow C \rightarrow A[1]$$

Alternatively, such a diagram is denoted by a triangle



The dashed line indicates the morphism has its degree shifted. Note: there is no notion of exactness in this diagram, which is important. We are now going to introduce distinguished triangles, which will be a weaker substitute for short exact sequences in certain categories which are not abelian. Let us note in passing that a morphism between triangles is a collection of morphisms between the objects in \mathcal{C} which make the obvious diagram commute.


Definition D.3.2. A triangulated category is an additive category \mathcal{C} together with a shift functor $[1]$ and a set of distinguished triangles, which are subject to the TR axioms, outlined below.

The intuition for why we might want to consider such categories will be provided after we have discussed the defining properties of distinguished triangles.


TR1 The first axiom states that:

D.3. TRIANGULATED CATEGORIES


1. For every $A \in \mathcal{C}$, there exists a distinguished triangle




2. If two triangles are isomorphic, then either they are both distinguished, or neither is distinguished.
3. For every $f \in \text{Hom}_{\mathcal{C}}(A, B)$, there exists an object $C(f)$ called the mapping cone of f , and a distinguished triangle



TR2 The second axiom states that, given a distinguished triangle




then the following triangle is also distinguished:



The converse statement should also hold.

TR3 The third axiom states that, if two distinguished triangles fit into a commutative diagram



D.3. TRIANGULATED CATEGORIES

then there should exist a morphism $h \in \text{Hom}_{\mathcal{C}}(C, C')$, which is not necessarily unique, which completes the diagram:

$$\begin{array}{ccccccc}
 A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & A[1] \\
 \downarrow & & \downarrow & & \downarrow h & & \downarrow \\
 A' & \longrightarrow & B' & \longrightarrow & C' & \longrightarrow & A'[1]
 \end{array}$$

TR4 So far, the axioms have seen quite reasonable. The final axiom may seem a bit less reasonable. It states the following. Let $f \in \text{Hom}_{\mathcal{C}}(A, B)$ and $g \in \text{Hom}_{\mathcal{C}}(B, C)$, which also results in $g \circ f \in \text{Hom}_{\mathcal{C}}(A, C)$. Each of these morphisms can be made into a distinguished triangle using the mapping cone from TR1. The final axiom requires the existence of a distinguished triangle

$$\begin{array}{ccc}
 C(f) & \longrightarrow & C(g \circ f) \\
 & \swarrow \text{dashed} & \searrow \\
 & C(g) &
 \end{array}$$

which makes the following diagram commute (which is why this axiom is called the octahedral axiom):

$$\begin{array}{ccccc}
 & & C(g \circ f) & & \\
 & \nearrow & \uparrow & \searrow & \\
 C(f) & \xleftarrow{\text{dashed}} & C(g) & \xrightarrow{\text{dashed}} & \\
 \downarrow & \nearrow \text{dashed} & \uparrow & \searrow \text{dashed} & \\
 A & \xrightarrow{\text{dashed}} & B & \xrightarrow{\text{dashed}} & C
 \end{array}$$

This concludes the list of axioms which a triangulated category needs to satisfy. Next, let us try to shed some light on them, as is done in [38].

Example 40. Given a continuous map between topological spaces $f : X \rightarrow Y$, the mapping cone is defined via

$$C(f) := (X \times I) \sqcup Y / \sim \quad (x, 0) \sim (x', 0), (x, 1) \sim f(x)$$

We can view $C(f)$ as providing a "(co)kernel up to homotopy equivalence", in the following sense. Let us say that for an inclusion map $i : X \hookrightarrow Y$, the quotient Y/X is the cokernel. Then $C(i) \simeq \text{coker } i$, so up to homotopy equivalence, $C(i)$ plays the role of the cokernel in topology. In a moment, we will see how this can be made more precise at the level of chain complexes. We can fit $C(i)$ into the following sequence

$$X \xrightarrow{i} Y \xrightarrow{l} C(i) \xrightarrow{\kappa} C(l) \rightarrow \dots$$

D.3. TRIANGULATED CATEGORIES

which is, up to homotopy equivalence,

$$X \rightarrow Y \rightarrow Y/X \rightarrow \Sigma X \rightarrow \dots$$

The suspension isomorphism from algebraic topology yields $H_i(\Sigma X) \cong H_{i-1}(X)$. Therefore, applying the homology functor to the above sequence, one obtains the long exact sequence of the pair (Y, X) :

$$H_n(X) \rightarrow H_n(Y) \rightarrow H_n(Y, X) \xrightarrow{\delta} H_{n-1}(X) \rightarrow \dots$$

A similar heuristic argument reveals that, if $X \rightarrow Y$ is a fibration, then $C(f)$ can be thought of as $\ker f$, up to homotopy equivalence, in some topological sense. In particular (when restricting attention to triangulations and CW complexes), we note that $C(f)$ is homotopy equivalent to a point if and only if f is a homotopy equivalence. From this, we can see a relation with short exact sequences and their induced long exact sequences in (co)homology. But how does it relate to distinguished triangles?

It can be shown that the chain complex of singular (co)chains of $C(f)$ is homotopy equivalent to the mapping cone of homological algebra, which is defined as follows.

Definition D.3.3. Let $f \in \text{Hom}_{\text{Ch}(\mathcal{A})}(E_\bullet, F_\bullet)$. The mapping cone $C(f)$ of f is defined as $C(f) := E_\bullet[1] \oplus F_\bullet$, and the differential is given by

$$d^i(a, b) = (-d_E^{i+1}(a), f^{i+1}(a) + d_F^i(b))$$

Clearly, the mapping cone of homological algebra naturally fits into a sequence

$$E_\bullet \xrightarrow{f} F_\bullet \xrightarrow{\iota} C(f) \xrightarrow{\pi} E_\bullet[1]$$

Proposition D.3.4. [117] For any short exact sequence of chain complexes $0 \rightarrow A_\bullet \rightarrow B_\bullet \rightarrow C_\bullet \rightarrow 0$, the natural maps $B_i \oplus A_{i-1} \rightarrow B_i \rightarrow C_i$ define a chain map $h : C(f) \rightarrow C_\bullet$. This map induces an isomorphism $H_i(h) : H_i(C(f)) \rightarrow H_i(C_\bullet)$.

Using this, one may prove the following result.

Theorem D.3.5. [117] Let $0 \rightarrow A_\bullet \rightarrow B_\bullet \rightarrow C_\bullet \rightarrow 0$ be a short exact sequence of chain complexes. Then up to isomorphism, the homology functor takes the distinguished triangle

$$A_\bullet \rightarrow B_\bullet \rightarrow C(f) \rightarrow A_\bullet[1]$$

to the long exact sequence

$$H_n(A_\bullet) \rightarrow H_n(B_\bullet) \rightarrow H_n(C_\bullet) \rightarrow H_{n-1}(A_\bullet) \rightarrow \dots$$

There are two main points in the above example. The first is that the mapping cone generalises the notion of kernels and cokernels, in the sense that we can think about this construction as providing (co)kernels up to homotopy equivalence (whenever such a notion exists). Let $f \in \text{Hom}_{\mathcal{A}}(E, F)$, which we also consider as a morphism between the complexes $E[0] \rightarrow F[0]$. If f is injective, then there is an isomorphism $H^i(C(f)) \cong H^i(\text{coker } f)$ for all i , and if f is surjective, then there exists an isomorphism $H^i(C(f)) \cong H^i(\ker f[1])$ for all i . The second point is that distinguished triangles in a triangulated category are the avatars of long exact sequences at the level of cohomology.

D.4 Derived Categories

In the appendix, we define the cohomology functors $H^i : \text{Ch}(\mathcal{A}) \rightarrow \mathcal{A}$ for abelian categories \mathcal{A} . In what is to follow, we would like to say that two morphisms $f, g \in \text{Hom}_{\text{Ch}(\mathcal{A})}(E_{\bullet}, F_{\bullet})$ are equivalent if the induced maps on cohomology $H^i(f) : H^i(E_{\bullet}) \rightarrow H^i(F_{\bullet})$ coincide.

Definition D.4.1. A morphism $f \in \text{Hom}_{\text{Ch}(\mathcal{A})}(E_{\bullet}, F_{\bullet})$ is called a quasi-isomorphism if $H^i(f) : H^i(E_{\bullet}) \rightarrow H^i(F_{\bullet})$ is an isomorphism for all i .

The basic idea of the derived category, is to apply the notion of localisation (as in ring theory) to the category of chain complexes, so that quasi-isomorphisms become invertible.

Theorem D.4.2. [30] *There exists a category $D(\mathcal{A})$ together with a functor $Q : \text{Ch}(\mathcal{A}) \rightarrow D(\mathcal{A})$, so that whenever $f \in \text{Hom}_{\text{Ch}(\mathcal{A})}(E_{\bullet}, F_{\bullet})$ is a quasi-isomorphism, its image $Q(f) \in \text{Hom}_{D(\mathcal{A})}(E_{\bullet}, F_{\bullet})$ is an isomorphism.*

Of course, the bounded derived category $D^b(\mathcal{A})$ appearing in homological mirror symmetry is defined as

$$\text{Ob}(D^b(\mathcal{A})) := \{A \in D(\mathcal{A}) \mid \exists n \in \mathbb{N} : H^i(A) = 0 \quad \forall |i| > n\}$$

Remark D.4.3. Note that the cohomology functors $H^i : \text{Ch}(\mathcal{A}) \rightarrow \mathcal{A}$ descend to give functors $H^i : D(\mathcal{A}) \rightarrow \mathcal{A}$. This is trivial, but also important.

We have written $\text{Hom}_{D(\mathcal{A})}(E_{\bullet}, F_{\bullet})$ because the functor Q does nothing to the objects, they are still chain complexes of objects in \mathcal{A} . As for the category $\text{Ch}(\mathcal{A})$, the original category \mathcal{A} embeds as a full subcategory into $D(\mathcal{A})$. Namely as the full subcategory of complexes with cohomology concentrated in degree 0. This is going to be important later on, for the following reason. Whenever we have an equivalence between two derived categories $D(\mathcal{A}_1) \cong D(\mathcal{A}_2)$, this does not typically arise from an equivalence between \mathcal{A}_1 and \mathcal{A}_2 . However, we can view both \mathcal{A}_1 and \mathcal{A}_2 as sitting inside of the same triangulated category, and giving us a different way to "generate" this triangulated category. Putting this into more formal terms, for a full subcategory $\mathcal{A} \subset \mathcal{B}$ we define the right-orthogonal of \mathcal{A} to be the full subcategory \mathcal{A}^{\perp} via

$$\text{Ob}(\mathcal{A}^{\perp}) := \{B \in \mathcal{B} \mid \text{Hom}_{\mathcal{B}}(A, B) = 0 \quad \forall A \in \mathcal{A}\}$$

D.4. DERIVED CATEGORIES

Definition D.4.4. Let \mathcal{B} be a triangulated category. A t -structure on \mathcal{B} is a full subcategory $\mathcal{A} \subset \mathcal{B}$ such that $\mathcal{A}[1] \subset \mathcal{A}$, and such that every $B \in \mathcal{B}$ fits into a distinguished triangle

$$A \rightarrow B \rightarrow C \rightarrow A[1]$$

where $A \in \mathcal{A}$ and $C \in \mathcal{A}^\perp$. The heart of the t -structure is defined by $\mathcal{A}^\heartsuit := \mathcal{A} \cap \mathcal{A}^\perp[1]$.

In our case, we would like to define a t -structure on $D(\mathcal{A})$ for which the heart is \mathcal{A} . This is not difficult to do. We define $\mathcal{T} := \{A \in D(\mathcal{A}) \mid H^i(A) = 0 \quad \forall i > 0\}$. One readily verifies that $\mathcal{T}^\perp = \{A \in D(\mathcal{A}) \mid H^i(A) = 0 \quad \forall i < 0\}$. Hence, we find that \mathcal{T}^\heartsuit consists of the complexes which are concentrated in degree 0, which is equivalent to \mathcal{A} , as desired. Since we will be mostly interested in the bounded derived category $D^b(\mathcal{A})$, the notion of a bounded t -structure is also useful. A t -structure \mathcal{T} on \mathcal{B} is said to be bounded if $\mathcal{B} = \cup_{i,j} (\mathcal{T}[i] \cap \mathcal{T}^\perp[j])$. The t -structure we gave above for $D(\mathcal{A})$ is evidently a bounded t -structure for $D^b(\mathcal{A})$.

Theorem D.4.5. [30] *A bounded t -structure \mathcal{T} is determined by its heart \mathcal{T}^\heartsuit .*

Now that we know this, we can at least philosophically understand what it means to say that $D^b(\mathcal{A}_1) \cong D^b(\mathcal{A}_2)$. It means that we have a triangulated category which is, in some sense, determined in two different ways, by two different abelian categories. Studying how these two t -structures interact might lead to interesting insights into the respective abelian categories, which are (as we said) typically not equivalent themselves. Having given this motivation, let us carry out the traditional method of constructing the derived category via the homotopy category $K(\mathcal{A})$.

The basic picture to keep in mind throughout this construction, is the singular cohomology of topological spaces. Suppose that two continuous maps $f, g : X \rightarrow Y$ are homotopic. Then the induced maps on cohomology are equivalent, $f^* = g^* : H^i(Y) \rightarrow H^i(X)$. This can be seen at the level of chain complexes because the induced maps are chain homotopic, in the sense of the following definition:

Definition D.4.6. Two morphisms between chain complexes $f, g : E_\bullet \rightarrow F_\bullet$ are called chain homotopic if there exists a morphism $h : E_\bullet \rightarrow F_\bullet[-1]$ such that, for all i ,

$$f_i - g_i = d_{i-1} \circ h_i + h_{i+1} d_i$$

For example, the typical proof that the de Rham cohomology is homotopy invariant involves the construction of a chain homotopy operator. In the homotopy category of an abelian category, we want to declare that chain homotopic morphisms are equivalent. If two morphisms of chain complexes are chain homotopic, we denote this by $f \sim g$. One easily verifies that this is indeed an equivalence relation.

Definition D.4.7. Let \mathcal{A} be an abelian category. Then the homotopy category $K(\mathcal{A})$ has $\text{Ob}(K(\mathcal{A})) = \text{Ob}(\text{Ch}(\mathcal{A}))$, and the morphisms are

$$\text{Hom}_{K(\mathcal{A})}(E_\bullet, F_\bullet) = \text{Hom}_{\text{Ch}(\mathcal{A})}(E_\bullet, F_\bullet) / \sim$$

D.4. DERIVED CATEGORIES

Thus, if we have two morphisms

$$f \in \text{Hom}_{\text{Ch}(\mathcal{A})}(E_\bullet, F_\bullet) \quad g \in \text{Hom}_{\text{Ch}(\mathcal{A})}(F_\bullet, E_\bullet)$$

and $f \circ g \sim \text{id}_F$, as well as $g \circ f \sim \text{id}_E$, then the equivalence class of f in $\text{Hom}_{K(\mathcal{A})}(E_\bullet, F_\bullet)$ is the inverse of the equivalence class of g in $\text{Hom}_{K(\mathcal{A})}(F_\bullet, E_\bullet)$. The construction inverts homotopy equivalences. Of course, homotopy equivalences induce isomorphisms on the respective cohomologies, so these would also be inverted in the derived category.

Definition D.4.8. In $D(\mathcal{A})$ or $K(\mathcal{A})$, we define a distinguished triangle to be a triangle isomorphic to one of the form

$$\begin{array}{ccc} E_\bullet & \xrightarrow{f} & F_\bullet \\ & \swarrow \text{---} & \searrow \\ & C(f) & \end{array}$$

The important result, which can be found in [30], states that this turns both categories into triangulated categories. With this result in hand, we can construct the derived category from $K(\mathcal{A})$ as follows. Morphisms in the derived category can be presented as equivalence classes of diagrams:

$$\begin{array}{ccc} & C_\bullet & \\ \swarrow \cong & & \searrow \\ A_\bullet & & B_\bullet \end{array}$$

We have denoted a quasi-isomorphism by writing \cong over the arrow. Two such diagrams are equivalent if there exists a commutative diagram in $K(\mathcal{A})$ of the form

$$\begin{array}{ccccc} & & D_\bullet & & \\ & & \swarrow \cong & \searrow & \\ & C_\bullet^1 & & C_\bullet^2 & \\ \swarrow \cong & & \searrow & \swarrow & \\ A_\bullet & & & & B_\bullet \end{array}$$

It remains to define a composition of morphisms in $D(\mathcal{A})$. In other words, given two diagrams

$$\begin{array}{ccc} & D_\bullet^1 & \\ \swarrow \cong & & \searrow \\ A_\bullet & & B_\bullet \end{array} \quad \begin{array}{ccc} & D_\bullet^2 & \\ \swarrow \cong & & \searrow \\ B_\bullet & & C_\bullet \end{array}$$

D.4. DERIVED CATEGORIES

We want the existence of a commutative diagram

$$\begin{array}{ccccc}
 & & E_{\bullet} & & \\
 & & \swarrow \cong & \searrow & \\
 & D^1_{\bullet} & & & D^2_{\bullet} \\
 & \swarrow \cong & & \swarrow \cong & \searrow \\
 A_{\bullet} & & & B_{\bullet} & & C_{\bullet}
 \end{array}$$

This composition is well-defined and unique in $D(\mathcal{A})$, by the following result.

Lemma D.4.9. [30] *Let $f : A_{\bullet} \xrightarrow{\cong} B_{\bullet}$, and $g : C_{\bullet} \rightarrow B_{\bullet}$. Then there exists a commutative diagram in $K(\mathcal{A})$ of the form*

$$\begin{array}{ccc}
 D_{\bullet} & \xrightarrow{\cong} & C_{\bullet} \\
 \downarrow & & \downarrow g \\
 A_{\bullet} & \xrightarrow{f} & B_{\bullet}
 \end{array}$$

The proof uses TR3 of the triangulated category axioms. As a corollary, the composition of morphisms is uniquely defined up to quasi-isomorphism in $K(\mathcal{A})$, and is hence defined uniquely in $D(\mathcal{A})$, as can be verified by applying the lemma iteratively. Because we have defined distinguished triangles in $D(\mathcal{A})$ to be isomorphic to the distinguished triangles coming from mapping cones, we immediately get the following result:

Proposition D.4.10. *Every distinguished triangle $A_{\bullet} \rightarrow B_{\bullet} \rightarrow C_{\bullet} \rightarrow A_{\bullet}[1]$ results in a long exact sequence in cohomology*

$$\dots \rightarrow H^i(A_{\bullet}) \rightarrow H^i(B_{\bullet}) \rightarrow H^i(C_{\bullet}) \rightarrow H^{i+1}(A_{\bullet}) \rightarrow \dots$$

Let us contemplate why we want to work with these categories in the first place. In the words of R.P. Thomas [38]:

"Chain complexes good, (co)homology bad."

What is meant by this? Cohomology is clearly an incredibly useful tool in all kinds of topology, geometry and algebra. The point is that the cohomology is obtained from some chain complex, and the chain complex itself contains more information than its (co)homology. This is obvious since the cohomology is obtained as the quotient of subspaces of terms in the chain complex.

Example 41. Let us look at some smooth manifold X . Then a triangulation $T(X)$ of X yields a chain complex of simplicial (co)chains $C_{\bullet}(T(X), \mathbb{Z})$. The chain complex contains more information than its (co)homology, but it is not canonically associated to X , because we had to make a choice of triangulation $T(X)$. However,

D.5. A_∞ -CATEGORIES

different choices of triangulations lead to isomorphic (co)homologies. Let us suppose that X is simply connected. Then there is the following theorem by Whitehead.

Theorem D.4.11. *Let X_1 and X_2 be simply connected simplicial complexes. Then $X_1 \simeq X_2$ if and only if there is a simplicial complex Y and simplicial maps $f_i : Y \rightarrow X_i$ which induce quasi-isomorphisms on the simplicial chain complexes.*

As such, if we are concerned with studying objects up to homotopy equivalence, it appears that the right notion to consider at the level of homological algebra is actually that of a quasi-isomorphism. If we then consider different triangulations of the (simply connected) manifold, we get quasi-isomorphic chain complexes. Thus, we can make the chain complex into a topological invariant by inverting quasi-isomorphisms, and this is a reason for us to be interested in the derived category.

We summarise this by thinking about derived categories through the following heuristics:

1. The derived category is the right setting in which to do homological algebra, retaining more information and yielding the right notion of isomorphism.
2. Short exact sequences of chain complexes give long exact sequences in (co)homology, which lift to distinguished triangles in the derived category (using TR2 D.3). Every distinguished triangle comes from a short exact sequence.
3. In particular, if $0 \rightarrow A_\bullet \rightarrow B_\bullet \rightarrow C_\bullet \rightarrow 0$ is a short exact sequence in $\text{Ch}(\mathcal{A})$, then this corresponds to a distinguished triangle $A_\bullet \rightarrow B_\bullet \rightarrow C_\bullet \rightarrow A_\bullet[1]$ in $D(\mathcal{A})$. If $A_\bullet \rightarrow B_\bullet \rightarrow C(f) \rightarrow A_\bullet[1]$ is a distinguished triangle in $D(\mathcal{A})$, there is a short exact sequence $0 \rightarrow B_\bullet \rightarrow C(f) \rightarrow A_\bullet[-1] \rightarrow 0$ in $\text{Ch}(\mathcal{A})$. Their long exact sequences in (co)homology are isomorphic.

D.5 A_∞ -Categories

We are not going to give any sort of geometric intuition for A_∞ -structures as we did for derived categories, referring instead to our discussion on the derived category of coherent sheaves, and the Fukaya category. The highly abstract theory presented in this section should fall into place, as there is a very concrete geometric interpretation for what we discuss here.

Definition D.5.1. An A_∞ -category \mathcal{A} is a class of objects $\text{Ob}(\mathcal{A})$, together with \mathbb{Z} -graded vector spaces $\text{Hom}_{\mathcal{A}}(A, B)$ for all $A, B \in \text{Ob}(\mathcal{A})$ and for all $k \geq 1$ and $A_0, \dots, A_k \in \text{Ob}(\mathcal{A})$ a graded morphism of degree $2 - k$

$$m_k : \text{Hom}_{\mathcal{A}}(A_{k-1}, A_k) \otimes \cdots \otimes \text{Hom}_{\mathcal{A}}(A_0, A_1) \rightarrow \text{Hom}_{\mathcal{A}}(A_0, A_k)$$

D.5. A_∞ -CATEGORIES

which is subject to the A_∞ -relations

$$\sum_{k=r+s+t} (-1)^{r+st} m_{r+1+t} \circ (\text{id}^{\otimes r} \otimes m_s \otimes \text{id}^{\otimes t}) \quad (\text{D.1})$$

Remark D.5.2. We are using the so-called Koszul sign convention for graded vector spaces. This means that if $f \otimes g$ is a bilinear map for some graded vector spaces V, W , then

$$(f \otimes g)(v \otimes w) = (-1)^{|v||g|} f(v) \otimes g(w)$$

As such, if we were to evaluate the maps m_k on elements of the Hom-spaces, we would pick up additional signs.

Remark D.5.3. There exists a "natural" interpretation of A_∞ -algebras in terms of coalgebras and codifferentials, which is generally used in the work of Kontsevich and Soibelman [118]. We are not going to be using that interpretation here, but this language does make it more evident that these objects do not come completely out of the blue.

Example 42. The A_∞ -relations are best understood when looking at A_∞ -algebras, which are A_∞ -categories with a single object. In this case, we obtain a single Hom-space which we denote A . It is a \mathbb{Z} -graded vector space, and the m_k become graded morphisms

$$m_k : A^{\otimes k} \rightarrow A$$

which satisfy the A_∞ -relations, that we now investigate for small k , i.e. $k = 1, 2, 3$. We need to look at possible non-negative integers r, s, t which sum to k , subject to the constraint $s \geq 1$.

Case $k = 1$ Clearly, there is but one option: $s = 1, r = t = 0$. Thus, the first A_∞ -relation dictates that $m_1 \circ m_1 = 0$. In other words, if $A = \bigoplus_{i \in \mathbb{Z}} A_i$ is the decomposition of A into its homogeneous degree parts, then m_1 yields a chain complex

$$\dots \xrightarrow{m_1} A_{i-1} \xrightarrow{m_1} A_i \xrightarrow{m_1} A_{i+1} \xrightarrow{m_1} \dots$$

This follows immediately from the fact that m_1 is a degree $2 - k = 1$ map.

Case $k = 2$ Keeping the same constraint in mind, there are now a few more options:

1. $r = 0, s = 2, t = 0$
2. $r = 0, s = 1, t = 1$
3. $r = 1, s = 1, t = 0$

D.5. A_∞ -CATEGORIES

The A_∞ -relation becomes

$$\begin{aligned} (-1)^{0+2\cdot 0} m_1 \circ m_2 + (-1)^{0+1\cdot 1} m_2 \circ (m_1 \otimes \text{id}) + (-1)^{1+1\cdot 0} m_2 \circ (\text{id} \otimes m_1) &= 0 \iff \\ m_1 \circ m_2 - m_2 \circ (m_1 \otimes \text{id}) - m_2 \circ (\text{id} \otimes m_1) &= 0 \iff \\ m_1 \circ m_2 &= m_2 \circ (m_1 \otimes \text{id} + \text{id} \otimes m_1) \end{aligned}$$

If we view m_2 as defining a multiplication operation $m_2 : A \otimes A \rightarrow A$, then this tells us that the differential m_1 satisfies the (graded) Leibniz rule. However, the multiplication defined by m_2 does not turn A into an algebra, because it may not be associative, and the existence of an identity is not guaranteed; see the case $k = 3$.

Case $k = 3$ We again investigate the possible choices:

1. $r = 0, s = 3, t = 0$
2. $r = 0, s = 2, t = 1$
3. $r = 0, s = 1, t = 2$
4. $r = 1, s = 2, t = 0$
5. $r = 1, s = 1, t = 1$
6. $r = 2, s = 1, t = 0$

The A_∞ -relation becomes

$$\begin{aligned} (-1)^{0+3\cdot 0} m_1 \circ m_3 + (-1)^{0+2\cdot 1} m_2 \circ (m_2 \otimes \text{id}) + (-1)^{0+1\cdot 2} m_3 \circ (m_1 \otimes \text{id} \otimes \text{id}) + \\ (-1)^{1+2\cdot 0} m_2 \circ (\text{id} \otimes m_2) + (-1)^{1+1\cdot 1} m_3 \circ (\text{id} \otimes m_1 \otimes \text{id}) + (-1)^{2+1\cdot 0} m_3 \circ (\text{id} \otimes \text{id} \otimes m_1) &\iff \\ m_1 \circ m_3 + m_3 \circ (m_1 \otimes \text{id} \otimes \text{id} + \text{id} \otimes m_1 \otimes \text{id} + \text{id} \otimes \text{id} \otimes m_1) &= \\ m_2 \circ (\text{id} \otimes m_2 - m_2 \otimes \text{id}) & \end{aligned}$$

What does this say? If we are given a vector space A and a linear map $* : A \otimes A \rightarrow A$, then we define the associator $\text{Ass} : A^{\otimes 3} \rightarrow A$ as

$$\text{Ass}(a \otimes b \otimes c) = a * (b * c) - (a * b) * c$$

Evidently, $*$ defines an associative multiplication map on A if and only if $\text{Ass} \equiv 0$. For us, the operation $*$ is given by $a * b = m_2(a \otimes b)$, so the associator is

$$\text{Ass} = m_2 \circ (\text{id} \otimes m_2 - m_2 \otimes \text{id})$$

This is one side of the third A_∞ -relation. The other side of the equation is not necessarily zero, which is why m_2 does not turn A into a genuine algebra. However, it contains two terms, one of which is a boundary in the chain complex defined by m_1 .

D.5. A_∞ -CATEGORIES

Furthermore, if we take cycles of the chain complex (A, m_1) as arguments (i.e. elements such that $m_1(a) = 0$), then we see that at the level of cohomology, m_2 actually gives us a well-defined, associative product

$$m_2 : H^\bullet(A) \otimes H^\bullet(A) \rightarrow H^\bullet(A) \quad [a] \otimes [b] \mapsto [m_2(a \otimes b)]$$

The higher A_∞ -relations are interpreted similarly, as associativity up to higher homotopy. This is also how we view the A_∞ -relations on a category, but in a more general setting.

In general, it is easier to think about A_∞ -categories at the level of algebras, and try to extrapolate relevant results to the more general setting. Let us apply this to the notion of a morphism between A_∞ -algebras, which should be generalised to an A_∞ -functor. It stands to reason that such a morphism $f : A \rightarrow B$ should consist of a family of linear maps

$$f_k : A^{\otimes k} \rightarrow B$$

The correct grading to impose, is that f_k has degree $1 - k$. Each f_k should satisfy the k -th A_∞ -relation on the respective A_∞ -algebras, so we require that

$$\sum_{r+s+t=k} (-1)^{r+st} f_{r+1+t} \circ (\text{id}^{\otimes r} \otimes m_s^A \otimes \text{id}^{\otimes t}) = \sum_{1 \leq r \leq k=i_1+\dots+i_r} (-1)^u m_r^B \circ (f_{i_1} \otimes \dots \otimes f_{i_r})$$

where $u = (r-1)(i_1-1) + (r-2)(i_2-1) + \dots + (i_{r-1}-1)$. If we were looking for beauty and elegance, this is probably not the right place. A morphism $f : A \rightarrow B$ is called strict if $f_k = 0$ for $k \geq 2$, because in this case, the equations above reduce to

$$m_k^B \circ (f_1 \otimes \dots \otimes f_1) = f_1 \circ m_k^A \circ (\text{id} \otimes \dots \otimes \text{id})$$

for all k . The f_k for $k \geq 2$ could be interpreted as measuring the failure of the above equality. On the other hand, the relation on f_1 just stipulates that $f_1 : (A, m_1^A) \rightarrow (B, m_1^B)$ should be a morphism of chain complexes. Using the A_∞ -morphisms for A_∞ -algebra, we can define A_∞ -functors.

Definition D.5.4. An A_∞ -functor $F : \mathcal{A}_1 \rightarrow \mathcal{A}_2$ is a map $\phi : \text{Ob}(\mathcal{A}_1) \rightarrow \text{Ob}(\mathcal{A}_2)$, and for any finite sequence of objects $A_0, \dots, A_k \in \text{Ob}(\mathcal{A}_1)$, a morphism of graded vector spaces of degree $1 - k$

$$f_k : \text{Hom}_{\mathcal{A}_\infty}(A_0, A_1) \otimes \dots \otimes \text{Hom}_{\mathcal{A}_1}(A_{k-1}, A_k) \rightarrow \text{Hom}_{\mathcal{A}_2}(\phi(A_0), \phi(A_k))$$

such that for any sequence $A_1, \dots, A_N \in \mathcal{A}_1$, the sequence f_k defines an A_∞ -morphism

$$f : \oplus_{i,j} \text{Hom}_{\mathcal{A}_1}(A_i, A_j) \rightarrow \oplus_{i,j} \text{Hom}_{\mathcal{A}_2}(\phi(A_i), \phi(A_j))$$

D.5. A_∞ -CATEGORIES

Such functors can be composed. Let $F : \mathcal{A}_1 \rightarrow \mathcal{A}_2$ and $G : \mathcal{A}_2 \rightarrow \mathcal{A}_3$. We set $H := G \circ F$ and take

$$h^k = \sum_{r \geq 1} \sum_{s_1 + \dots + s_r = k} g^r \circ (f^{s_1} \otimes \dots \otimes f^{s_r})$$

The map on objects should be clear. In this way, we obtain the category of A_∞ -categories, which is in fact a genuine category.

The statement of homological mirror symmetry involves an A_∞ -quasi-equivalence between A_∞ -categories. To state what this is, we note that, just as the cohomology of an A_∞ -algebra becomes a (not necessarily unital) non-commutative algebra, so too can we recover a category from an A_∞ -category (with the exceptions that there may be no identity morphisms). In particular, given an A_∞ -category \mathcal{A} , we define $H^\bullet(\mathcal{A})$ to be the category whose objects are those of \mathcal{A} , and whose morphisms are given by

$$\mathrm{Hom}_{H^\bullet(\mathcal{A})}(A_0, A_1) := H^\bullet(\mathrm{Hom}_{\mathcal{A}}(A_0, A_1))$$

Because we have passed to cohomology, the composition of morphisms is now associative, as it was for the A_∞ -algebra case. We similarly define $H^0(\mathcal{A})$, except by taking

$$\mathrm{Hom}_{H^0(\mathcal{A})}(A_0, A_1) := H^0(\mathrm{Hom}_{\mathcal{A}}(A_0, A_1))$$

Clearly, an A_∞ -functor $F : \mathcal{A}_1 \rightarrow \mathcal{A}_2$ induces a functor

$$H^\bullet(F) : H^\bullet(\mathcal{A}_1) \rightarrow H^\bullet(\mathcal{A}_2)$$

by taking a morphism $[a] \in H^\bullet(\mathrm{Hom}_{\mathcal{A}_1}(A_0, A_1))$ to $[F_1(a)]$. For us, the categories $H^\bullet(\mathcal{A}_i)$ will be genuine categories, i.e. each object will have an identity morphism. We make the following definition with this understanding in mind.

Definition D.5.5. Let \mathcal{A}_1 and \mathcal{A}_2 be A_∞ -categories, and suppose that $H^\bullet(\mathcal{A}_i)$ are both categories. Let $F : \mathcal{A}_1 \rightarrow \mathcal{A}_2$ be an A_∞ -functor. Then F is called a quasi-equivalence of A_∞ categories if

$$H^\bullet(F) : H^\bullet(\mathcal{A}_1) \rightarrow H^\bullet(\mathcal{A}_2)$$

is an equivalence of categories.

In some sense, we only care about A_∞ -structures up to quasi-equivalence. As such, we want a good model for a certain A_∞ -structure, i.e. the most convenient representative of a quasi-isomorphism class of A_∞ -categories. This leads to the notion of a minimal model.

Theorem D.5.6 (Kontsevich-Soibelman [118]). *Let \mathcal{A} be an A_∞ -category. Then there exists an A_∞ -structure on $H^\bullet(\mathcal{A})$ such that the differential vanishes, and a quasi-isomorphism $H^\bullet(\mathcal{A}) \simeq \mathcal{A}$.*

There is one final ingredient that we need. Namely, we need to construct $D^b(\mathcal{A})$ for an A_∞ -category, or something which plays the role of the bounded derived category. This is done via twisted complexes, for which we need to define the A_∞ version of the Yoneda embedding, for which we need some additional machinery.

D.5. A_∞ -CATEGORIES

Definition D.5.7. Let \mathcal{A} be an A_∞ -category. Then a (right) A_∞ -module \mathcal{M} over \mathcal{A} consists of a graded vector space $\mathcal{M}(A)$ for each $A \in \mathcal{A}$, together with multiplication maps of degree $2 - k$ for $k \geq 1$

$$m_k^{\mathcal{M}} : \mathcal{M}(A_{k-1}) \otimes \text{Hom}_{\mathcal{A}}(A_{k-2}, A_{k-1}) \otimes \cdots \otimes \text{Hom}_{\mathcal{A}}(A_0, A_1) \rightarrow \mathcal{M}(A_0)$$

which satisfy the relations D.1. However, the term

$$m_{r+1+t} \circ (\text{id}^{\otimes r} \otimes m_s \otimes \text{id}^{\otimes t})$$

is to be interpreted as

$$m_{r+1+t}^{\mathcal{M}} \circ (m_s^{\mathcal{M}} \otimes \text{id}^{\otimes t})$$

when $r = 0$ and as

$$m_{r+1+t}^{\mathcal{M}} \circ (\text{id}^{\otimes r} \otimes m_s \otimes \text{id}^{\otimes t})$$

when $r > 0$. Of course, m_j are the multiplication maps of the A_∞ -category.

Once again, the relation for $k = 1$ makes each $\mathcal{M}(A)$ into a chain complex with differential $m_2^{\mathcal{M}}$, the second one tells us that the graded Leibniz rule is satisfied, and the third one measures the failure of the associativity of the map $m_2^{\mathcal{M}}$. We want to produce an A_∞ -category of A_∞ -modules, denoted $\text{Mod}(\mathcal{A})$. To do this, we need to specify what a morphism between A_∞ -modules looks like. Let \mathcal{M}_1 and \mathcal{M}_2 be \mathcal{A} -modules. A pre-morphism of degree d between them is a sequence of maps of degree $d - k + 1$ for $k \geq 1$

$$f_k : \mathcal{M}_1(A_{k-1}) \otimes \text{Hom}_{\mathcal{A}}(A_{k-2}, A_{k-1}) \otimes \cdots \otimes \text{Hom}_{\mathcal{A}}(A_0, A_1) \rightarrow \mathcal{M}_2(A_0)$$

Let $\text{Hom}_{\text{Mod}(\mathcal{A})}(\mathcal{M}_1, \mathcal{M}_2)$ be the graded vector space of pre-morphisms between \mathcal{M}_1 and \mathcal{M}_2 . We will define a differential on this vector space, and a multiplication map. All higher m_k will vanish. We define m_1 on the degree d homogeneous part by

$$\begin{aligned} m_1(f)_d := & \sum_n (-1)^\star m_{n+1}^{\mathcal{M}_2} \circ (f_{d-n} \otimes \text{id}^{\otimes n}) + \sum_n (-1)^\star f_{n+1} \circ (m_{d-n}^{\mathcal{M}_1} \otimes \text{id}^{\otimes n}) \\ & + \sum_{m,n} (-1)^\star f_{d-m+1} (\text{id}^{\otimes (d-m-n)} \otimes m_n^{\mathcal{A}} \otimes \text{id}^{\otimes n}) \end{aligned}$$

where $\star = n + 1 - d$. Given $f, g \in \text{Hom}_{\text{Mod}(\mathcal{A})}(\mathcal{M}_1, \mathcal{M}_2)$, we define their composition as

$$m_2(f, g)_d := \sum_n (-1)^\star g_{n+1} \circ (f_{d-n} \otimes \text{id}^{\otimes n})$$

We cite [5] for the fact that this yields a differential graded category. That is, we indeed have $m_k = 0$ for $k \geq 3$, while m_1 is a differential which satisfies the graded Leibniz rule w.r.t. m_2 .

Definition D.5.8. An element $f \in \text{Hom}_{\text{Mod}(\mathcal{A})}(\mathcal{M}_1, \mathcal{M}_2)$ is called a homomorphism if $m_1(f) = 0$. It is called an isomorphism if it is a homomorphism such that $[f]$ is an isomorphism in $H^\bullet(\text{Mod}(\mathcal{A}))$.

D.5. A_∞ -CATEGORIES

Now, given any $B \in \mathcal{A}$, we can define an A_∞ -module \mathcal{M}_B by taking $\mathcal{M}_B(A) := \text{Hom}_{\mathcal{A}}(A, B)$, and setting $m_k^{\mathcal{M}_B} := m_k^A$. This gives us a functor

$$Y : \mathcal{A} \rightarrow \text{Mod}(\mathcal{A}) \quad B \mapsto (A \mapsto \text{Hom}_{\mathcal{A}}(A, B))$$

which is called the A_∞ -Yoneda embedding. To make this into a functor, we need to define

$$Y_k : \text{Hom}_{\mathcal{A}}(A_{k-1}, A_k) \otimes \cdots \otimes \text{Hom}_{\mathcal{A}}(A_0, A_1) \rightarrow \text{Hom}_{\text{Mod}(\mathcal{A})}(Y(A_0), Y(A_k))$$

of degree $1 - k$. We do this by setting

$$(Y_k(a_k, \dots, a_1))_d(b, b_{d-1}, \dots, b_1) := m_{k+d}^A(a_d, \dots, a_1, b, b_{d-1}, \dots, b_1)$$

An object $A \in \mathcal{A}$ is said to quasi-represent an A_∞ -module \mathcal{M} when there is an isomorphism of A_∞ -functors $Y(A) \cong \mathcal{M}$. Technically, this definition requires us to define natural transformations between A_∞ -functors, but we leave this to the imagination of the reader. We prefer to work with the category $\text{Mod}(\mathcal{A})$ because it allows us to use more algebraic constructions, such as the shift functor. Define an endofunctor $\Sigma : \text{Mod}(\mathcal{A}) \rightarrow \text{Mod}(\mathcal{A})$ by

$$(\Sigma \mathcal{M})(A) := \mathcal{M}(A)[1]$$

If we suppose that, for any $A \in \mathcal{A}$, the functor $\Sigma Y(A)$ is quasi-representable by some $\Sigma A \in \mathcal{A}$, then we in fact get a shift functor $\Sigma : \mathcal{A} \rightarrow \mathcal{A}$, taking A to ΣA . This is going to be important when we discuss the Fukaya category, because it will allow us to define a shift functor in an easy way. However, our present goal is to make \mathcal{A} into a triangulated category, so we need an analogue of the mapping cone construction as well. Again, we construct a sensible \mathcal{A} -module, which is in this case defined for $f \in \text{Hom}_{\mathcal{A}}^0(B_1, B_2)$ by setting

$$C(f)(A) := \text{Hom}_{\mathcal{A}}(A, B_1)[1] \oplus \text{Hom}_{\mathcal{A}}(A, B_2)$$

The morphisms are defined by

$$m_k^{C(f)}((b_0, b_1), a_{k-1}, \dots, a_1) = (m_k^A(b_0, a_{k-1}, \dots, a_1, m_k^A(b_1, a_{k-1}, \dots, a_1) + m_{k+1}^A(f, b_0, a_{k-1}, \dots, a_1))$$

As for the chain complex version of the mapping cone, we get canonical (pre-)morphisms

$$\iota \in \text{Hom}_{\text{Mod}(\mathcal{A})}^0(Y(B_2), C(f)) \quad \pi \in \text{Hom}_{\text{Mod}(\mathcal{A})}^1(C(f), Y(B_1))$$

We do not distinguish between the \mathcal{A} -module $C(f)$ and an object which represents it. Then the above morphisms give us a triangle in $H^\bullet(\text{Mod}(\mathcal{A}))$.

$$\begin{array}{ccc} Y(B_1) & \xrightarrow{[Y_1(f)]} & Y(B_2) \\ & \swarrow [\pi] \text{ (dashed)} & \searrow [\iota] \\ & & C(f) \end{array}$$

This finally allows us to define the notion of a triangulated A_∞ -category.

D.5. A_∞ -CATEGORIES

Definition D.5.9. A distinguished triangle in an A_∞ -category \mathcal{A} is a diagram in $H^\bullet(\mathcal{A})$ which is isomorphic to the triangle above under the Yoneda embedding. An A_∞ -category is said to be triangulated if

1. There is a shift functor $\Sigma : \mathcal{A} \rightarrow \mathcal{A}$.
2. Every $[f] \in \text{Hom}_{H^0(\mathcal{A})}(A, B)$ can be extended to a distinguished triangle.
3. For every $A \in \mathcal{A}$, there exists $A' \in \mathcal{A}$ such that $\Sigma A' \cong A \in H^0(\mathcal{A})$.

Not all A_∞ -categories are triangulated, and in particular, the Fukaya category is not (or at least, is not known to be). There is a way to modify the category \mathcal{A} to produce a triangulated category, by using twisted complexes.

Definition D.5.10. A twisted complex in an A_∞ -category is a sequence of objects A_1, \dots, A_n together with a strictly lower diagonal matrix μ of morphisms $\mu_{ij} \in \text{Hom}^1(A_j, A_i)$ satisfying $\sum_k m_k(\mu, \dots, \mu) = 0$.

A twisted object is thus a pair (A, μ^A) . The category $\text{Tw}(\mathcal{A})$ has these as objects, and the Hom-spaces are defined by

$$\text{Hom}_{\text{Tw}(\mathcal{A})}((A, \mu^A), (B, \mu^B)) := \oplus_{i,j} \text{Hom}_{\mathcal{A}}(A_i, B_j)$$

which makes $\text{Tw}(\mathcal{A})$ into an A_∞ -category. We assume that \mathcal{A} has a shift functor Σ , which (as mentioned) the Fukaya category does. Then given $f \in \text{Hom}_{\text{Tw}(\mathcal{A})}^0((A_1, \mu^{A_1}), (A_2, \mu^{A_2}))$, we define

$$C(f) := ((\Sigma A_1, A_2), \begin{pmatrix} \Sigma(\mu^{A_1}) & 0 \\ -\Sigma(f) & \mu^{A_2} \end{pmatrix})$$

This gives us an abstract mapping cone of f , which turns $\text{Tw}(\mathcal{A})$ into a triangulated category. We assume that $H^0(\text{Tw}(\mathcal{A}))$ has identity morphisms, which turns it into a genuine category.

Definition D.5.11. Let \mathcal{A} be an A_∞ -category with a shift functor $\Sigma : \mathcal{A} \rightarrow \mathcal{A}$. Then the mapping cone constructed above turns $\text{Tw}(\mathcal{A})$ into a triangulated A_∞ -category. We define $D^b(\mathcal{A}) := H^0(\text{Tw}(\mathcal{A}))$ and say that it is the bounded derived category of \mathcal{A} .

At this point, we have established the theory of A_∞ -categories that we need for the statement of homological mirror symmetry, so we will discuss derived category of coherent sheaves, as well as the twisted Fukaya category. Hopefully, these concrete examples will provide the reader with an idea as to why these constructions were worthwhile, in spite of their rather inelegant (at times) formulations.

Bibliography

- [1] R. P. Thomas, “Moment maps, monodromy and mirror manifolds,” 2001.
- [2] R. P. Thomas and S. T. Yau, “Special Lagrangians, stable bundles and mean curvature flow,” 2001.
- [3] P. Candelas, X. C. De La Ossa, P. S. Green, and L. Parkes, “A Pair of Calabi-Yau manifolds as an exactly soluble superconformal theory,” *Nucl. Phys. B*, vol. 359, pp. 21–74, 1991.
- [4] M. Kontsevich, “Homological algebra of mirror symmetry,” 1994.
- [5] P. Aspinwall *et al.*, *Dirichlet Branes and Mirror Symmetry*. American Mathematical Society, 2009.
- [6] M. R. Douglas, B. Fiol, and C. Römelsberger, “Stability and BPS branes,” *Journal of High Energy Physics*, vol. 2005, pp. 006–006, sep 2005.
- [7] K. Uhlenbeck and S. T. Yau, “On the existence of Hermitian-Yang-Mills connections in stable vector bundles,” *Communications on Pure and Applied Mathematics*, vol. 39, no. S1, pp. S257–S293, 1986.
- [8] S. K. Donaldson and P. B. Kronheimer, *The Geometry of Four-Manifolds*. Oxford Science Publications, 1997.
- [9] J. D. Lotay and G. Oliveira, “Special Lagrangians, Lagrangian mean curvature flow and the Gibbons-Hawking ansatz,” 2020.
- [10] B. Fang, “Homological mirror symmetry is T-duality for \mathbb{P}^n ,” 2008.
- [11] G. Oliveira and R. Sena-Dias, “Minimal Lagrangian tori and action-angle coordinates,” 2020.
- [12] D. Joyce, “Conjectures on Bridgeland stability for Fukaya categories of Calabi-Yau manifolds, special Lagrangians, and Lagrangian mean curvature flow,” *EMS Surveys in Mathematical Sciences*, vol. 2, no. 1, pp. 1–62, 2015.
- [13] M. Kontsevich and Y. Soibelman, “Stability structures, motivic Donaldson-Thomas invariants and cluster transformations,” 2008.
- [14] R. P. Thomas, “A holomorphic Casson invariant for Calabi-Yau 3-folds, and bundles on K3 fibrations,” 1998.
- [15] B. Bolognese, A. Küronya, and M. Ulirsch, “ $P = W$ phenomena on abelian varieties,” 2023.

BIBLIOGRAPHY

- [16] M. A. de Cataldo, T. Hausel, and L. Migliorini, “Topology of Hitchin systems and Hodge theory of character varieties: the case A_1 ,” 2011.
- [17] A. da Silva, *Lectures on Symplectic Geometry*. Springer, 2008.
- [18] M. Seri, *Lecture Notes on Hamiltonian Mechanics*. github.io, 2020.
- [19] A. Strominger, S.-T. Yau, and E. Zaslow, “Mirror symmetry is T-duality,” *Nuclear Physics B*, vol. 479, pp. 243–259, nov 1996.
- [20] D. Huybrechts, *Complex Geometry*. Springer, 2005.
- [21] M. Gross *et al.*, *Calabi-Yau Manifolds and Related Geometries*. Springer, 2003.
- [22] R. Wells, *Differential Analysis on Complex Manifolds*. Springer, 2008.
- [23] J.-P. Demailly, “*Complex Analytic and Differential Geometry*”. Université de Grenoble, 1997.
- [24] K. Kodaira and D. C. Spencer, “On Deformations of Complex Analytic Structures, I,” *Annals of Mathematics*, vol. 67, no. 2, pp. 328–401, 1958.
- [25] R. Miranda, *Algebraic Curves and Riemann Surfaces*. American Mathematical Society, 1995.
- [26] S.-T. Yau, *Open Problems in Geometry*. American Mathematical Society, 1993.
- [27] R. Harvey and H. B. Lawson, “Calibrated geometries,” *Acta Mathematica*, vol. 148, no. none, pp. 47 – 157, 1982.
- [28] D. Auroux, “A beginner’s introduction to Fukaya categories,” 2013.
- [29] D. Huybrechts and M. Lehn, *The Geometry of the Moduli Space of Sheaves*. Cambridge University Press, 2010.
- [30] M. Ballard, C. Doran, D. Favero, and E. Sharpe, *Superschool on Derived Categories and D-branes*. Springer, 2016.
- [31] <https://ncatlab.org/nlab/show/higher+category+theory+and+physics>.
- [32] B. Kreussler, “Homological Mirror Symmetry in Dimension One,” 2000.
- [33] M. F. Atiyah, “Vector Bundles Over an Elliptic Curve,” *Proceedings of the London Mathematical Society*, vol. s3-7, no. 1, pp. 414–452, 1957.
- [34] S. Mukai, “Duality between $D(X)$ and $D(\hat{X})$ with its application to Picard sheaves,” *Nagoya Mathematical Journal*, vol. 81, no. none, pp. 153 – 175, 1981.
- [35] J. Tu, “Homological mirror symmetry and Fourier-Mukai transform,” 2012.
- [36] A. Bondal and D. Orlov *Compositio Mathematica*, vol. 125, no. 3, pp. 327–344, 2001.
- [37] T. Bridgeland, “Flops and derived categories,” 2000.

BIBLIOGRAPHY

- [38] R. P. Thomas, “Derived categories for the working mathematician,” 2000.
- [39] R. Hartshorne, *Algebraic Geometry*. Springer, 1977.
- [40] <https://stacks.math.columbia.edu/tag/0656>.
- [41] B. Greene and M. Plesser, “Duality in Calabi-Yau moduli space,” *Nuclear Physics B*, vol. 338, no. 1, pp. 15–37, 1990.
- [42] K. Hori *et al.*, *Mirror Symmetry*. American Mathematical Society, 2000.
- [43] P. Deligne, P. Etingof, D. S. Freed, L. C. Jeffrey, D. Kazhdan, J. W. Morgan, D. R. Morrison, and E. Witten, eds., *Quantum fields and strings: A course for mathematicians. Vol. 1, 2*. 1999.
- [44] M. Vonk, “A mini-course on topological strings,” 2005.
- [45] P. Deligne and D. S. Freed, “Supersolutions,” 1999.
- [46] L. Alvarez-Gaumé and D. Z. Freedman, “Geometrical structure and ultraviolet finiteness in the supersymmetric sigma-model,” *Communications in Mathematical Physics*, vol. 80, no. 3, pp. 443 – 451, 1981.
- [47] P. Candelas, G. T. Horowitz, A. Strominger, and E. Witten, “Vacuum configurations for superstrings,” *Nucl. Phys. B*, vol. 258, pp. 46–74, 1985.
- [48] A. Kapustin and D. Orlov, “Vertex algebras, mirror symmetry, and d-branes: The case of complex tori,” *Communications in Mathematical Physics*, vol. 233, pp. 79–136, feb 2003.
- [49] E. Witten, “Topological quantum field theory,” *Communications in Mathematical Physics*, vol. 117, no. 3, pp. 353 – 386, 1988.
- [50] E. Witten, “Mirror manifolds and topological field theory,” 1991.
- [51] D. A. Cox and S. Katz, “Mirror symmetry and algebraic geometry,” 1999.
- [52] P. S. Aspinwall and D. R. Morrison, “String Theory on K3 Surfaces,” 1994.
- [53] D. Huybrechts, “Moduli spaces of hyperkaehler manifolds and mirror symmetry,” 2002.
- [54] P. S. Aspinwall and A. Lawrence, “Derived categories and zero-brane stability,” *Journal of High Energy Physics*, vol. 2001, pp. 004–004, aug 2001.
- [55] P. S. Aspinwall, “D-Branes on Calabi-Yau Manifolds,” in *Progress In String Theory*, WORLD SCIENTIFIC, jul 2005.
- [56] M. Herbst, C. I. Lazaroiu, and W. Lerche, “Superpotentials, A_∞ Relations and WDVV Equations for Open Topological Strings,” *Journal of High Energy Physics*, vol. 2005, pp. 071–071, 2004.

BIBLIOGRAPHY

- [57] S. Katz and E. Sharpe, “D-branes, open string vertex operators, and ext groups,” *Advances in Theoretical and Mathematical Physics*, vol. 6, no. 6, pp. 979–1030, 2002.
- [58] C. I. Lazaroiu, “D-Brane Categories,” *International Journal of Modern Physics A*, vol. 18, pp. 5299–5335, nov 2003.
- [59] A. Polishchuk and E. Zaslow, “Categorical Mirror Symmetry: The Elliptic Curve,” 1998.
- [60] P. Seidel, “Homological mirror symmetry for the quartic surface,” 2003.
- [61] N. Sheridan, “Homological mirror symmetry for Calabi–Yau hypersurfaces in projective space,” *Inventiones mathematicae*, vol. 199, pp. 1–186, feb 2014.
- [62] M. Abouzaid and I. Smith, “Homological mirror symmetry for the 4-torus,” *Duke Mathematical Journal*, vol. 152, apr 2010.
- [63] K. Fukaya, “Mirror symmetry of Abelian varieties and multi-theta functions,” *Journal of Algebraic Geometry*, vol. 11, 09 2002.
- [64] A. Polishchuk, *Abelian varieties, theta functions and the Fourier transform*. Cambridge University Press, 2003.
- [65] D. Bejleri, “The SYZ conjecture via homological mirror symmetry,” 2017.
- [66] M. Gross, “Topological Mirror Symmetry,” 1999.
- [67] N. C. Leung, S.-T. Yau, and E. Zaslow, “From Special Lagrangian to Hermitian-Yang-Mills via Fourier-Mukai Transform,” 2000.
- [68] D. Arinkin and A. Polishchuk, “Fukaya category and Fourier transform,” 1998.
- [69] R. Pandharipande and R. P. Thomas, “13/2 ways of counting curves,” in *Moduli Spaces*, pp. 282–333, Cambridge University Press, mar 2014.
- [70] D. Maulik, N. Nekrasov, A. Okounkov, and R. Pandharipande, “Gromov-Witten theory and Donaldson-Thomas theory, I,” 2004.
- [71] J. Li and G. Tian, “Virtual moduli cycles and Gromov-Witten invariants of algebraic varieties,” 1996.
- [72] S. Feyzbakhsh and R. P. Thomas, “Rank r DT theory from rank 0,” 2022.
- [73] S. K. Donaldson and R. P. Thomas, “Gauge theory in higher dimensions,” in *Conference on Geometric Issues in Foundations of Science in honor of Sir Roger Penrose’s 65th Birthday*, pp. 31–47, 6 1996.
- [74] S. K. Donaldson, “Anti Self-Dual Yang-Mills Connections Over Complex Algebraic Surfaces and Stable Vector Bundles,” *Proceedings of the London Mathematical Society*, vol. s3-50, no. 1, pp. 1–26, 1985.

BIBLIOGRAPHY

- [75] D. Maulik and J. Shen, “The $P = W$ conjecture for GL_n ,” 2022.
- [76] N. J. Hitchin, “Stable bundles and integrable systems,” *Duke Math. J.*, vol. 54, pp. 91–114, 1987.
- [77] N. J. Hitchin, “The Selfduality equations on a Riemann surface,” *Proc. Lond. Math. Soc.*, vol. 55, pp. 59–131, 1987.
- [78] *Integrable systems: Twistors, loop groups, and Riemann surfaces. Proceedings, Conference, Oxford, UK, September 1997*, 1999.
- [79] *Nonabelian Hodge Theory. Proceedings of the International Congress of Mathematicians, Kyoto*, 1990.
- [80] R. Donagi and T. Pantev, “Langlands duality for Hitchin systems,” 2006.
- [81] Z. Zhang, “Multiplicativity of Perverse Filtration for Hilbert Schemes of Fibered Surfaces,” 2017.
- [82] J. A. M. Silva, “Hodge-Deligne Polynomials of Symmetric Products of Algebraic Groups,” 2019.
- [83] W. Y. Chuang, D. E. Diaconescu, and G. Pan, “BPS states and the $P=W$ conjecture,” 2012.
- [84] R. Pandharipande and R. P. Thomas, “Curve counting via stable pairs in the derived category,” *Inventiones mathematicae*, vol. 178, pp. 407–447, may 2009.
- [85] Y. Toda, “Curve counting theories via stable objects I. DT/PT correspondence,” 2009.
- [86] T. Bridgeland, “Hall algebras and curve-counting invariants,” 2020.
- [87] D. Maulik and Y. Toda, “Gopakumar-Vafa invariants via vanishing cycles,” 2018.
- [88] S. Hosono, M.-H. Saito, and A. Takahashi, “Relative Lefschetz Action and BPS State Counting,” *International Mathematics Research Notices*, vol. 2001, no. 15, p. 783, 2001.
- [89] J. Shen and Q. Yin, “Topology of Lagrangian fibrations and Hodge theory of hyper-Kähler manifolds,” 2021.
- [90] L. Göttsche, “Hilbert schemes of points on surfaces,” 2003.
- [91] K. Becker, M. Becker, and A. Strominger, “Fivebranes, membranes and non-perturbative string theory,” *Nuclear Physics B*, vol. 456, pp. 130–152, dec 1995.
- [92] P. S. Aspinwall, “D-Branes, Pi-Stability and Theta-Stability,” 2004.
- [93] T. Bridgeland, “Stability conditions on triangulated categories,” 2002.

BIBLIOGRAPHY

- [94] S. K. Donaldson, “Moment maps and diffeomorphisms,” *Surveys in differential geometry*, vol. 3, pp. 107–127, 2002.
- [95] D. Joyce, “On counting special Lagrangian homology 3-spheres,” 1999.
- [96] M. Kontsevich and Y. Soibelman, “Wall-crossing structures in Donaldson-Thomas invariants, integrable systems and Mirror Symmetry,” 2013.
- [97] Y. Zhu, “Donaldson-Thomas invariants and wall-crossing formulas,” 2014.
- [98] P. Seidel, “Graded lagrangian submanifolds,” 1999.
- [99] G. Lawlor, “The angle criterion.,” *Inventiones mathematicae*, vol. 95, no. 2, pp. 437–446, 1989.
- [100] J. Lotay, “Introduction to Lagrangian Mean Curvature Flow,” <http://people.maths.ox.ac.uk/lotay>.
- [101] Y. Li, “Thomas-yau conjecture and holomorphic curves,” 2022.
- [102] M. A. Grayson, “Shortening embedded curves,” *Annals of Mathematics*, vol. 129, no. 1, pp. 71–111, 1989.
- [103] M. Gross and P. M. H. Wilson, “Large Complex Structure Limits of K3 Surfaces,” 2000.
- [104] D. Auroux, “Mirror symmetry and T-duality in the complement of an anticanonical divisor,” 2007.
- [105] N. C. Leung and M. Symington, “Almost toric symplectic four-manifolds,” 2003.
- [106] G. Tian, “On Kähler-Einstein metrics on certain Kähler manifolds with $c_1(X) > 0$,” *Inventiones mathematicae*, vol. 89, pp. 225–246, 1987.
- [107] A. Bayer, E. Macri, and Y. Toda, “Bridgeland Stability conditions on threefolds I: Bogomolov-Gieseker type inequalities,” 2011.
- [108] A. Maciocia and D. Piyaratne, “Fourier-Mukai Transforms and Bridgeland Stability Conditions on Abelian Threefolds,” 2013.
- [109] J. D. Lotay and G. Oliveira, “Neck pinch singularities and joyce conjectures in lagrangian mean curvature flow with circle symmetry,” 2023.
- [110] K. Fukaya, “Cyclic symmetry and adic convergence in Lagrangian Floer theory,” *Kyoto Journal of Mathematics*, vol. 50, jan 2010.
- [111] D. Joyce and Y. Song, “A theory of generalized Donaldson-Thomas invariants,” 2008.
- [112] Y. Toda, “Limit stable objects on Calabi-Yau 3-folds,” *Duke Mathematical Journal*, vol. 149, jul 2009.

BIBLIOGRAPHY

- [113] T. C. Collins and Y. Shi, “Stability and the deformed Hermitian-Yang-Mills equation,” 4 2020.
- [114] M. R. Douglas, R. Reinbacher, and S.-T. Yau, “Branes, Bundles and Attractors: Bogomolov and Beyond,” 2006.
- [115] G. Rudolph and M. Schmidt, *Differential Geometry and Mathematical Physics, Part 2: Fibre Bundles, Topology and Gauge Fields*. Springer, 2017.
- [116] D. Arapura, *Algebraic Geometry over the Complex Numbers*. Springer, 2012.
- [117] <https://ncatlab.org/nlab/show/mapping+cone>.
- [118] M. Kontsevich and Y. Soibelman, “Notes on A_∞ -algebras, A_∞ -categories and non-commutative geometry. i,” 2006.