Homological Mirror Symmetry for Elliptic Curves

April 25, 2013

We prove homological mirror symmetry for elliptic curves. This is done by applying the methods Seidel developed for quartic surfaces to the much easier one-dimensional case. We state and prove as many preliminary definitions and results as space permits.

After the abstract proof is finished, we determine the mirror map, and illustrate the mirror correspondence.

Contents

List of Figures

1. Introduction

Mirror symmetry was observed by physicists working on so-called superconformal string theories. In particular, given an algebraic Calabi-Yau manifold X, physicists predict the existence of a "mirror" symplectic manifold $X^\vee$ (also Calabi-Yau) and various relations between the two, which can crudely be summarized as

$$\text{algebraic geometry on } X \longleftrightarrow \text{symplectic geometry on } X^\vee. \tag{1}$$

One book attempting to fill in the vagueness of "$\leftrightarrow$" in our crude summary above is [3].

A (superficially) rather different approach to filling in the vagueness was presented by Kontsevich at the ICM in 1994 [10]. He formulated a set of conjectures, nowadays summarised as "homological mirror symmetry", distilling many of the physicists' claims into a categorical framework, which we briefly describe.

We start with an algebraic Calabi-Yau variety X. One category encoding almost all of the geometry of X is the bounded derived category of coherent sheaves $D^b(Coh(X))$. It is the category obtained from the category of bounded cochain complexes of coherent sheaves on X by inverting the quasi-isomorphisms. This is not an abelian category. It is however additive, and comes with certain extra structure known as a shift functor and distinguished triangles. This extra data is usually axiomatized under the heading "triangulated category". (Chapter ten of [23] is a good reference for triangulated and derived categories in this sense, but beware that we will usually mean something else in the main text.) This is Kontsevich's interpretation of the right hand side of "equation" (1): the bounded derived category of coherent sheaves on X, viewed as a triangulated category.

Consider now a symplectic Calabi-Yau Y (that is to say, a real smooth manifold with vanishing first Chern class, together with a choice of closed non-degenerate two-form ω_Y). Building on work of Fukaya, Kontsevich conjectures the existence of the so-called split-closed triangulated derived Fukaya category $H^0D^\pi\mathrm{Fuk}(Y)$, and this is Kontsevich's interpretation of the right hand side of (1). Up to understanding what $H^0D^\pi\mathrm{Fuk}(Y)$ is, the homological mirror symmetry conjecture can be formulated now as follows:

Conjecture (Kontsevich). *Let X be a projective Calabi-Yau manifold. There exists a symplectic manifold $X^\vee$ and an equivalence of triangulated categories*

$$D^b(Coh(X)) \simeq H^0D^\pi\mathrm{Fuk}(X^\vee).$$

As a first approximation, the objects of $\mathrm{Fuk}(Y)$ consist of compact Lagrangian submanifolds of Y (i.e. submanifolds L such that $\omega_Y|_L = 0$). These are in particular half-dimensional, so in good cases (transversal intersection), the set of intersection points of two Lagrangian submanifolds $L_1 \cap L_2$ is finite. We take these intersection points to be (formal) generators of a vector space (usually called Floer chain complex) which we denote $\mathrm{Hom}_{\mathrm{Fuk}(Y)}(L_1, L_2)$. As the name suggests, this will carry a differential, and we will have $\mathrm{Hom}_{H^0D^\pi\mathrm{Fuk}(Y)}(L_1, L_2) = H^0\,\mathrm{Hom}_{\mathrm{Fuk}(Y)}(L_1, L_2)$. That is to say, the derived Fukaya category of Kontsevich's conjecture is obtained from a "chain-level" category $\mathrm{Fuk}(Y)$ by taking cohomology.

This is glossing over a variety of difficulties which we want to at least mention. (1) There is an obvious problem when L_1 and L_2 do not intersect transversely, and this is actually very hard to fix. (2) We have not defined the differential, and also not how to compose morphisms. In nice cases, this can be done using Lagrangian intersection Floer theory. This means that differentials and compositions are computed by counting immersed bigons and triangles (up to a certain equivalence relation). While after taking cohomology this does indeed turn out to yield an ordinary category, the chain level structure is very complicated. In particular, chain-level composition is not associative. This is remedied by higher composition operations (also defined by counting polygons, with more sides) which serve as homotopies measuring the failure of associativity. The algebraic structure we obtain is usually called an A_∞-category. (3) It turns out that the category $H^0\mathrm{Fuk}(Y)$ is not particularly nice. For example, idempotent endomorphisms need not correspond to subobjects. Also $H^0\mathrm{Fuk}(Y)$, while additive, is neither abelian nor triangulated. There is, however a purely algebraic procedure, called "forming the split-closed triangulated envelope" which from $\mathrm{Fuk}(Y)$ constructs a new A_∞-category $D^\pi\mathrm{Fuk}(Y)$ such that $H^0D^\pi\mathrm{Fuk}(Y)$ has these nice properties (i.e. it comes with a natural triangulation, and idempotent endomorphisms correspond to subobjects). It is this final category which is supposed to be triangle-equivalent to the bounded derived category.

There is one shortcoming in Kontsevich's conjecture, at least as stated. This has to do with the observation that results at the level of cohomology should really be shadows of chain-level phenomena. For example the derived category, while a natural category to state results in, is actually rather awkward to work with (e.g. because the derived category of X cannot be recovered from the derived categories of an open cover by "gluing"). One should hence try to find a natural "chain-level" category $D_\infty(X)$ such that $H^0 D_\infty(X) \simeq D^b(Coh(X))$. Fortunately, such categories are well known (and in fact used to prove existence of $D^b(Coh(X))$ in the first place). Moreover, the concept of an A_∞-category applies in an obvious way to $D_\infty(X)$ (we in fact obtain a particularly simple A_∞-structure known as dg-category). We can thus state an improved version of Kontsevich's conjecture as follows:

Conjecture (Kontsevich, chain-level)**.** *Let X be a projective Calabi-Yau manifold. There exists a symplectic manifold $X^\vee$ and a morphism of A_∞-categories*

$$D_\infty(X) \to D^\pi \mathrm{Fuk}(X^\vee)$$

inducing an isomorphism after passing to cohomology.

This conjecture has now been established for all Calabi-Yau hypersurfaces (in projective space). The case of elliptic curves (dimension one) has been known the longest and was established in a series papers of Polishchuk and Zaslow [17] [14] [15] [16]. Quartic threefolds (dimension two) were treated by Seidel [19]. The case of dimension ≥ 3 was established by Sheridan [21].

The main aim of this essay is to present a proof of the first non-trivial case of Kontsevich's conjecture, i.e. for elliptic curves. This is done by applying the refined methods developed by Seidel for quartic surfaces to the much easier one-dimensional case. In doing so we essentially follow [11], at least in spirit (in practice we aim to prove much less, and can thus cut some corners).

This essay has two secondary aims: firstly to develop the relatively large amount of background material needed to state precisely, and eventually prove, Kontsevich's conjecture. Secondly, we wish to illustrate both the techniques used in the proof of homological mirror symmetry, and the extraordinary "coincidences" which are required to identify the symplectic and algebraic geometry sides of the mirror pair.

Organisation of this essay We sketch the basic argument presented in this essay.

In section 2 we define A_∞-categories, and the important notions of A_∞-functor, triangulated A_∞-category and split-closed A_∞-category. We state general existence results for triangulated and split-closed envelopes. All of this is preparation, so that we can state the main theorem from homological algebra we will use:

Theorem. *Suppose $\mathcal{A}$ and $\mathcal{B}$ are two split-closed A_∞-categories, split-generated (respectively) by $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_m$. Let $\mathcal{A}' = \mathrm{Hom}_\mathcal{A}(X_1 \oplus \cdots \oplus X_n, X_1 \oplus \cdots \oplus X_n)$, and similarly for $\mathcal{B}'$. These are A_∞-algebras.*

Then any quasi-isomorphism (i.e. morphism which induces an isomorphism in cohomology) $\mathcal{A}' \to \mathcal{B}'$ extends canonically to a quasi-equivalence $\mathcal{A} \to \mathcal{B}$ (i.e. an A_∞-functor which induces an equivalence in cohomology).

We apologize for the fact that this section consists mainly of definitions and statements of theorems. It provides the homological algebra underpinning the rest of this essay. For reasons of space we cannot elaborate on the proofs or illustrate the diverse uses of A_∞-structures. The reader is advised to read only subsection 2.1 and to refer to the other subsections as needed. If she already knows enough about A_∞-categories to understand the above theorem, the first section can be safely skipped and only referred to as needed.

In section 3 we make first contact with geometry and construct an interesting A_∞-category. Namely, using injective resolutions, we construct a dg-enhancement $D_\infty(X)$ of the bounded derived category of coherent sheaves on a scheme X. We prove that if X is regular projective over an algebraically closed field, then $D_\infty(X)$ is generated by the powers of an ample line bundle. If X is furthermore one-dimensional (i.e. a smooth curve), we deduce that in fact a line bundles $\mathcal{O}_X$ and $\mathcal{O}_X(P)$ suffice.

In particular, if X is an elliptic curve, then we obtain a dg-algebra $\mathcal{Q} = \mathrm{Hom}_{D_\infty(X)}(\mathcal{O}_X \oplus \mathcal{O}_X(P), \mathcal{O}_X \oplus \mathcal{O}_X(P))$ of "relations", and its cohomology algebra $Q = H(\mathcal{Q})$ is a finite-dimensional quiver algebra which is easy to determine explicitly. In order to apply the above theorem, we need to find similar generators for the Fukaya category.

Before that, of course, we have to *define* the Fukaya category. This we do in section 4. We first introduce some basic symplectic geometry, and then describe Seidel's "graded symplectic geometry". This is the basic building block of the Fukaya category. We then describe how to construct the A_∞-structure on subsets of the Fukaya category represented by transverse Lagrangians. We need to apologize in advance to the reader for being incomplete here, even in two regards: firstly we do not deal with the non-transverse case. Doing this would lead us too far astray. For the computations in this essay, ad-hoc methods of resolving this problem suffice, provided we assume that a consistent construction of the Fukaya category exists and has the good properties we claim. Secondly, even for the transverse we do not provide the full definition of certain sign conventions, and no proof of the good properties we claim. The latter is due to space constraints, and the former is because these signs are of no use for us, other than establishing the good properties.

Accepting this description of the Fukaya category, we can press on with the proof of the mirror symmetry conjecture. Using the properties of the Fukaya category, it is easy to show that for a torus T, two meridians A and B generating the homology also generate the Fukaya category. We thus obtain an A_∞-algebra of relations $\mathcal{Q}' = \mathrm{Hom}_{D\mathrm{Fuk}(T)}(A \oplus B, A \oplus B)$. In order to apply the above theorem, we need to find a quasi-isomorphism $\mathcal{Q}' \to \mathcal{Q}$. As on the algebraic side, computing $H(\mathcal{Q}')$ is fairly easy, and one sees by direct comparison that $H(\mathcal{Q}') \simeq Q$.

At this point one of the strengths of the A_∞-picture becomes evident: by the so-called "homological perturbation theory", given any A_∞-algebra $\mathcal{A}$ with cohomology algebra A, we can find an A_∞-structure (m_d) on A with vanishing differential m_1 such that $\mathcal{A}$ is quasi-isomorphic to $(A, (m_d))$. We may thus trade the chain-level information on $\mathcal{A}$ for higher multiplications on A. Applying this to $\mathcal{Q}$ and $\mathcal{Q}'$, we end up with two A_∞-structures (m_d) and (m'_d) on Q, and it remains to show that these are isomorphic.

At this point we turn the problem on its head: in section 5, we present some results about the Hochschild cohomology of graded algebras. In particular we show how Hochschild cohomology computations can be used to classify the A_∞-structures on a graded algebra extending the given multiplication.

Our aim now becomes to classify all A_∞-structures on Q. The key insight of Lekili and Perutz [11] is that almost all of these actually come from elliptic curves by the process described above (there are two exceptions, coming from singular cubic curves). We tie together the argument by proving this key result in section 6. This finishes the abstract proof of Kontsevich's conjecture in the case of elliptic curves.

We then present various extensions and applications in section 7. First, we describe how to actually find the mirror elliptic curve corresponding to the flat torus. This is essentially a polygon counting procedure using a reconstruction theorem we proved along the way in section 6. It has been carried out by Zaslow [24]. We sketch the required computation and report on his result. In the remainder of this section, we illustrate further details of the mirror correspondence. In particular, we explain how the standard $SL_2(\mathbb{Z})$ action on the torus yields a braid action on the Fukaya category, and identify the corresponding braid action on the derived category of coherent sheaves.

Notation and conventions Whenever we say chain complex, we mean cochain complex. Similarly, all our graded vector spaces are index cohomologically. We use Koszul sign conventions for tensor products of complexes and graded vector spaces. In particular, if U, V, W, Z are graded vector spaces and $f : U \to V, g : W \to Z$ are graded maps, then we denote by $f \otimes g : U \otimes W \to V \otimes Z$ the map $(f \otimes g)(x \otimes y) = (-1)^{|g||x|} f(x) \otimes g(y)$, where $|x|$ denotes the degree of the element x, and similarly $|f|$ denotes the degree of the graded map f. For graded vector spaces, Hom denotes the degree-zero maps. The shifts of a graded vector space V are denoted $V(i)$, the shifts of a chain complex C are denoted $C[i]$.

If X is a scheme and P is a point, we denote by $\mathcal{O}_X/P$ the skyscraper sheaf $k(P)$ at P, where $k(P) = \mathcal{O}_{X,P}/m_{X,P}$ is the residue field at P.

For various objects for which it makes sense, we denote by $X^\vee$ the dual. This applies for example to $\mathcal{O}$-modules on a scheme, or vector spaces. Similarly, by $\mathbf{1}$ or $\mathbf{1}_X$ we denote in various contexts the identity map (on X). Isomorphisms (or equivalences of categories) are denoted by $\simeq$ or $\xrightarrow{\sim}$. We use no special notation for quasi-isomorphisms, homotopy-equivalences or the like.

We denote by Λ_t the universal Novikov field over the complex numbers, see section 4.4.

Acknowledgements Many people have helped me in learning this material. I would like to thank in particular Ivan Smith for supervising the essay, Ian Grojnowski for many explanations on homological algebra and algebraic geometry, Yanki Lekili for presenting his paper in a seminar and answering questions I had on symplectic geometry, and Emile Bouaziz for general discussions and pointers to the literature.

2. A_∞-categories

We now define A_∞-categories, the building block of the homological mirror symmetry conjecture. As explained in the introduction, this section will be as brief as possible, essentially without any proofs.

A good introduction to A_∞-algebras is [9]. A more comprehensive account is in the first part of [20]. Beware that these two sources use differing sign conventions. We follow the first.

2.1. Basic definitions

We fix a field (or possibly ring) k. We will say vector space even when we actually mean module (in the unusual case that k is not a field).

Definition. *An A_∞-category $\mathcal{C}$ consist of the following data: a class $Ob(\mathcal{C})$ of* objects *of $\mathcal{C}$, for each pair $X, Y \in \mathcal{C}$ a graded vector space* $\mathrm{Hom}_{\mathcal{C}}(X, Y)$*, and for each set of objects $X_0, \dots, X_n \in Ob(\mathcal{C})$, a linear map*

$$m_n^{\mathcal{C},X_0,\dots,X_n} : \mathrm{Hom}_{\mathcal{C}}(X_{n-1}, X_n) \otimes \dots \otimes \mathrm{Hom}_{\mathcal{C}}(X_0, X_2) \to \mathrm{Hom}_{\mathcal{C}}(X_0, X_n)(2-d).$$

*We require that, for each set of objects $X_0, \dots, X_n \in Ob(\mathcal{A})$, the following A_∞-*relation *has to be satisfied:*

$$\sum_{a+b+c=n} (-1)^{a+bc} m_{a+1+b}^{\mathcal{C},X_0,\dots,X_c,X_{b+c},\dots,X_n}(\mathbf{1}^{\otimes a} \otimes m_b^{\mathcal{C},X_c,\dots,X_{c+b}} \otimes \mathbf{1}^{\otimes c}) = 0.$$

We immediately lighten the notation: instead of $X \in Ob(\mathcal{C})$ we will just write $X \in \mathcal{C}$, and instead of $m_n^{\mathcal{C},X_0,\dots,X_n}$ we will write $m_n^{\mathcal{C}}$ or even just m_n.

The defining equation looks somewhat peculiar (see section 5 for a re-interpretation). We write them out for the first few values of n, in the case $X_0 = X_1 = \dots = X$:

$$m_1 m_1 = 0$$
$$-m_2(\mathbf{1} \otimes m_1 + m_1 \otimes \mathbf{1}) + m_1 m_2 = 0$$
$$m_3(\mathbf{1} \otimes \mathbf{1} \otimes m_1 + \mathbf{1} \otimes m_1 \otimes \mathbf{1} + m_1 \otimes \mathbf{1} \otimes \mathbf{1}) + m_2(-\mathbf{1} \otimes m_2 + m_2 \otimes \mathbf{1}) + m_1 m_3 = 0.$$

The first of these says that $m_1 : \mathrm{Hom}(X, X) \to \mathrm{Hom}(X, X)$ is an ordinary (graded) differential, and the second says that m_2 is a derivation with respect to m_1. In particular, upon passing to cohomology, m_2 descends to a well-defined bilinear operation. The third equation can perhaps be more usefully written as $m_2(\mathbf{1} \otimes m_2) - m_2(m_2 \otimes \mathbf{1}) = m_1 m_3 + m_3(\mathbf{1} \otimes \mathbf{1} \otimes m_1 + \mathbf{1} \otimes m_1 \otimes \mathbf{1} + m_1 \otimes \mathbf{1} \otimes \mathbf{1})$. It implies in particular that m_2 becomes associative in cohomology. Here m_3 is a kind of generalized chain homotopy between the zero map and the associator of m_2.

If $\mathcal{C}$ is an A_∞-category, then we obtain a new "category" $H\mathcal{C}$ by declaring that $Ob(H\mathcal{C}) = Ob(\mathcal{C})$ and that $\mathrm{Hom}_{H\mathcal{C}}(X, Y) = H^* \mathrm{Hom}_{\mathcal{C}}(X, Y)$. This is almost a category, except that it need not have identity morphisms. If it does, then $\mathcal{C}$ is called *cohomologically unital. We will assume this throughout this work.*

We view $H\mathcal{C}$ as a category enriched over graded k-vector spaces. We define $H^0\mathcal{C}$ similarly (enriched over ungraded k-vector spaces).

We now consider some examples. By definition, an A_∞-algebra is a A_∞-category with only one object. A dg-category is an A_∞-category where $m_n = 0$ for $n > 2$. A dg-algebra is defined in the obvious way. Any graded associative algebra is an example of a dg-algebra, with $m_1 = 0$. Similarly a category enriched over graded k-vector spaces is an example of a dg-category, with $m_1 = 0$. An A_∞-category with $m_1 = 0$ is called *minimal.*

Suppose that $\mathcal{B}$ is an abelian category enriched over k-vector spaces. We can form a dg-category C as follows: we declare $Ob(C)$ to consist of the chain complexes in B, and for $M^\bullet, N^\bullet$ chain complexes, we let $\mathrm{Hom}_C(M,N)^n = \prod_{q-p=n} \mathrm{Hom}_{\mathcal{B}}(M^p, N^q)$. This carries a natural differential $m_1(f) = d_N \circ f - (-1)^{|f|} f \circ d_M$. We denote the resulting complex by $\mathrm{Hom}^\bullet(M,N)$. The composition in $\mathcal{B}$ supplies a multiplication m_2 for C, and one may check that this indeed yields a dg-category. A particular example of this is when $\mathcal{B}$ is the category of k-vector spaces. In this case, we denote the dg-category C by $\mathcal{C}h$. A variant of this is its *opposite* $\mathcal{C}h^{op}$, where $\mathrm{Hom}_{\mathcal{C}h^{op}}(M,N) = \mathrm{Hom}_{\mathcal{C}h}(N,M)$ and $m_2^{op}(f,g) = (-1)^{|f||g|} m_2(g,f)$.

Definition. *Let $\mathcal{A}$ and $\mathcal{B}$ be A_∞-categories. By an A_∞-functor $F : \mathcal{A} \to \mathcal{B}$ we mean a map $F : Ob(\mathcal{A}) \to Ob(\mathcal{B})$, and for each set $X_1, \dots, X_n \in Ob(\mathcal{A})$ a linear map F_d :* $\mathrm{Hom}(X_{d-1}, X_d) \otimes \ldots \otimes \mathrm{Hom}(X_0, X_1) \to \mathrm{Hom}(FX_0, FX_d)(1-d)$

We require that

$$\sum_{r, a_1 + \cdots + a_r = d} (-1)^\epsilon m_r^{\mathcal{B}}(F_{a_1} \otimes \ldots \otimes F_{a_r}) = \sum_{a+b+c=d} (-1)^{a+bc} F_{a+1+c}(\mathbf{1}^{\otimes a} \otimes m_b^{\mathcal{A}} \otimes \mathbf{1}^{\otimes c}),$$

where $\epsilon = (r-1)(a_1 - 1) + (r-2)(a_2 - 1) + \cdots + (2)(a_{r-2} - 1) + (1)(a_{r-1} - 1)$.

While the notation m_d is highly ambiguous, it should almost never lead to confusion. Again the required relation looks rather strange. Writing it out for small n as before, one finds that F_1 is a chain map, and that F_2 respects multiplication up to homotopy. In particular, F descends to an ordinary functor $HF : H\mathcal{A} \to H\mathcal{B}$, provided it preserves the identities. We will assume this throughout.

An A_∞-functor F is called a *quasi-isomorphism* if HF is an isomorphism. It is called a *quasi-equivalence* if HF is an equivalence.

There is a natural way of composing A_∞-functors, but we will not need this explicitly. There is in fact a notion of homotopy and natural transformation of A_∞-functors, turning the class of A_∞-functors from $\mathcal{A}$ to $\mathcal{B}$ into an A_∞-category of its own. This is very useful, but the definition is complicated, so we will avoid it.

2.2. Homological perturbation theory

This is perhaps the first non-trivial result about A_∞-categories. For a reference, see [20, (1i)]. The setup is as follows: let $\mathcal{B}$ be a A_∞-category and denote by $\mathcal{A}$ an "A_∞-category to be". By this we mean that $Ob(\mathcal{A}) = Ob(\mathcal{B})$, and for each $X, Y \in \mathcal{A}$ we are given a graded vector space $\mathrm{Hom}_{\mathcal{A}}(X,Y)$.

Theorem 2.1 (Homological perturbation lemma)**.** *Let k be any ring, not necessarily a field.*

In the above situation, suppose we are also given (for each $X, Y \in \mathcal{A}$) a graded differential $m_1^{\mathcal{A}}$ on $\mathrm{Hom}_{\mathcal{A}}(X,Y)$ *of degree $+1$, a linear endomorphism T_1 of* $\mathrm{Hom}_{\mathcal{B}}(X,Y)$ *of degree -1, $F_1 : \mathrm{Hom}_{\mathcal{A}}(X,Y) \to \mathrm{Hom}_{\mathcal{B}}(X,Y)$ and $G_1 : \mathrm{Hom}_{\mathcal{B}}(X,Y) \to \mathrm{Hom}_{\mathcal{A}}(X,Y)$.*

If (for all X, Y) $m_1^{\mathcal{B}} T_1 + T_1 m_1^{\mathcal{B}} = F_1 G_1 - \mathbf{1}$, then it is possible to produce an A_∞-structure (m_d) on $\mathcal{A}$ extending the m_1 given, and to extend F_1, G_1 to A_∞-functors. Moreover, FG is a quasi-isomorphism.

A few remarks are in order. It is suggestive to view, T_1 as the first term of an "A_∞-homotopy". A stronger version of this theorem in fact extends this to a homotopy between FG and $\mathbf{1}$. Secondly, it should be mentioned that the above process is completely functorial: there exist explicit formulas for all the maps this theorem constructs (and they involve no division, hence why this works over any ground ring). See the reference for details.

To illustrate the power of this theorem, let's record two corollaries:

Corollary. *Let $\mathcal{A}$ be an A_∞-category over a field. There exists an A_∞-category $\tilde{\mathcal{A}}$ with $Ob(\tilde{\mathcal{A}}) = Ob(\mathcal{A})$ and a quasi-isomorphism $F : \tilde{\mathcal{A}} \to \mathcal{A}$ with $H(F) = \mathbf{1}$. Moreover, $\tilde{\mathcal{A}}$ is unique up to (non-unique) isomorphism.*

Proof. Since we are working over a field, we may split $\mathrm{Hom}_{\mathcal{A}}(X, Y) = \mathrm{Hom}_{H\mathcal{A}}(X, Y) \oplus C(X, Y)$. This yields a homotopy to feed into the above theorem. □

Corollary. *Let $L : \mathcal{A} \to \mathcal{B}$ be an A_∞-functor between A_∞-categories. If L is a quasi-equivalence, then there exists $M : \mathcal{B} \to \mathcal{A}$ such that $H(ML) \simeq \mathbf{1}_{H(\mathcal{A})}$, $H(LM) \simeq \mathbf{1}_{H(\mathcal{B})}$. (Here by $\simeq$ we mean isomorphism of functors, i.e. invertible natural transformations.) That is to say, L admits a quasi-inverse.*

This is an easy application of the homological perturbation lemma and the following classification of certain automorphisms of A_∞-categories.

Suppose $\mathcal{A}$ is an A_∞-category, and we denote by F the following data: for each $X_0, \ldots, X_n \in \mathcal{A}$, a linear map $F_n : \mathrm{Hom}_{\mathcal{A}}(X_{n-1}, X_n) \otimes \ldots \otimes \mathrm{Hom}_{\mathcal{A}}(X_0, X_1) \to \mathrm{Hom}_{\mathcal{A}}(X_0, X_n)$. We call F a *formal diffeomorphism* if F_1 is an isomorphism for all $X_0 \in \mathcal{A}$.

The following proposition is a relatively easy, explicit calculation.

Proposition 2.2. *Suppose $\mathcal{A}$ is an A_∞-category and F a formal diffeomorphism.*

Then there exists a unique alternative A_∞-structure $F_*\mathcal{A}$ such that F becomes an A_∞-functor $F : \mathcal{A} \to F_*\mathcal{A}$. Moreover, F is automatically an isomorphism.

2.3. Quasi-representation and triangles

By an A_∞-module on an A_∞-category $\mathcal{A}$, we mean an A_∞-functor $M : \mathcal{A} \to \mathcal{C}h^{op}$. As we have mentioned before, using the notion of homotopy of A_∞-functors, one may turn the set of all A_∞-modules on $\mathcal{A}$ into an A_∞-category $Mod(\mathcal{A})$. Given $X \in \mathcal{A}$, there exists an A_∞-module ιX given by $\iota X(Y) = \mathrm{Hom}_{\mathcal{A}}(Y, X)$. In fact, there exists an A_∞-functor $\iota : \mathcal{A} \to Mod(\mathcal{A})$, which is an embedding of A_∞-categories (for an appropriate definition of embedding).

This notion allows us to transfer constructions in the category of chain complexes to arbitrary A_∞-categories. Namely, given an A_∞-module M on $\mathcal{A}$, we say an object $Y \in \mathcal{A}$ quasi-represents M if $\iota Y \simeq M$ in $H^0 Mod(\mathcal{A})$. Since we have not actually defined what this means, here is a more concrete definition. Fix $Y \in \mathcal{A}$. Note that for every $X \in \mathcal{A}$, we have a composition map $m_2 : M(Y) \otimes \iota Y(X) \to M(X)$. Thus, for a fixed $c \in M(Y)$, we get a composition map $m_2(c, \bullet) : \iota Y(X) \to M(X)$.

Definition. *In the above situation, suppose* $[c] \in H^0 M(Y)$. *Then we say* $(Y, [c])$ quasi-represents M *if for each* X, $m_2(c, \bullet)$ *is a quasi-isomorphism – that is to say if* $m_2(c, \bullet)$ *induces an isomorphism* $H(\iota Y(X)) \to HM(X)$. *(This condition is independent of the choice of cocycle* c *representing* $[c]$.*)*

If M, N are $\mathcal{A}$-modules, then we can define the $\mathcal{A}$-module $M \oplus N$ via $(M \oplus N)(X) = M(X) \oplus N(X)$ and the direct sum of the structure maps of M and N. Similarly, if V is a graded vector space (viewed as a complex with zero differentials) and M is an $\mathcal{A}$-module, we define $(V \otimes M)(X) = V \otimes M(X)$, with structure maps just ignoring V, i.e. $m_d(z \otimes f_d, f_{d-1}, \dots, f_1) = z \otimes m_d(f_d, \dots, f_1)$.

For X, Y in $\mathcal{A}$ and V a graded vector space (again viewed as a complex with zero differentials), any object quasi-representing $\iota X \oplus \iota Y$ will be denote $X \oplus Y$ and called *sum* of X and Y. Any object quasi-representing $V \otimes \iota X$ will be denoted $V \otimes X$ and called *tensor product* of V and X. If $V = k[1]$ is one-dimensional concentrated in degree -1, we denote $V \otimes X$ by SX or even $X[1]$, and call it the *shift* of X.

From the above remarks it is clear that the above constructions, if they exist, are unique up to quasi-isomorphism.

Notice that some properties of these constructions in $\mathcal{C}h$ carry over to arbitrary A_∞-categories. For example $(k \oplus k) \otimes X \simeq X \oplus X$, and so on.

Other properties are less obvious. For example, *functorial* constructions in $\mathcal{C}h$ in this way actually yield functors on $\mathcal{A}$, provided they can always be carried out. So if, for example, SX exists for all $X \in \mathcal{A}$, then there exists a functor $S : \mathcal{A} \to \mathcal{A}$, unique up to unique natural transformation in $H^0(\mathcal{A})$. See [20, p. 34] for more details.

Triangulated A_∞-categories

Consider the A_∞-category $\mathcal{D}$ with three objects Z_0, Z_1 and Z_2, such that $\mathrm{Hom}_{\mathcal{D}}(Z_i, Z_i)$ is one-dimensional spanned by e_{Z_i}. The space of morphisms from Z_0 to Z_1 is one-dimensional spanned by z_1 in degree 0, and similarly from Z_1 to Z_2. The space of morphisms from Z_2 to Z_0 is spanned by z_3 in degree *one*. We posit that all m_d are zero, except for m_2 with identities and m_3 involving z_3, z_2 and z_1 in cyclic order, which yields identities in the appropriate space.

Definition. *Let* $\mathcal{A}$ *be an* A_∞*-category. By a* triangle *in* $H(\mathcal{A})$ *we mean three objects* X_0, X_1, X_2 *together with morphisms* $[x_1] \in \mathrm{Hom}^0_{H(\mathcal{A})}(X_0, X_1)$, $[x_2] \in \mathrm{Hom}^0_{H(\mathcal{A})}(X_1, X_2)$ *and* $[x_3] \in \mathrm{Hom}^1_{H(\mathcal{A})}(X_2, X_0)$.

We say that a triangle is exact if there exist an A_∞*-functor* $F : \mathcal{D} \to \mathcal{A}$ *such that* $F(Z_i) = X_i$ *and* $[F(z_i)] = [x_i]$*, where by the last statement we mean equality in* $H(\mathcal{A})$.

We call $\mathcal{A}$ triangulated *if (1) every morphism can be completed to an exact triangle, and (2) the shift functor is essentially surjective on* $H^0(\mathcal{A})$.

We denote an exact triangle formed by objects X_0, X_1 and X_2 as

$$X_0 \to X_1 \to X_2 \to .$$

This is a slightly obscure-looking definition. One may prove that if $X_0 \to X_1 \to X_2 \to$ is an exact triangle, then $\iota(X_1 \to X_2 \to X_0[1])$ is quasi-isomorphic to a natural cone construction in $\mathcal{C}h^{opp}$ on $\iota X_0 \to \iota X_1$. We have no space here to say more, and just summarize the properties of triangulated A_∞-categories we need. The main point is really (1), then (2 - 3) follow easily, and (4) follows directly from the definitions.

Proposition 2.3. *Let $\mathcal{A}$ be a triangulated A_∞-category, and let $Y_0 \xrightarrow{y_1} Y_1 \xrightarrow{y_2} Y_2 \xrightarrow{y_3}$ be an exact triangle.*

1. *$H^0(\mathcal{A})$ is a triangulated category, with shift functor induced from S and distinguished triangles given by exact triangles of $\mathcal{A}$.*
2. *The triangle $Y_1 \to Y_2 \to SY_0 \to$ with edge morphisms $[y_2]$, $[y_3]$ and $[-Sy_1]$ is exact.*
3. *Given exact triangles $Y_0 \xrightarrow{[a_1]} Y_1 \to Z_0 \to$, $Y_1 \xrightarrow{[a_2]} Y_2 \to Z_2 \to$ and $Y_0 \xrightarrow{[a_2 a_1]} Y_2 \to Z_1 \to$, then there is also an exact triangle $Z_0 \to Z_1 \to Z_2 \to$.*
4. *If $F : \mathcal{A} \to \mathcal{B}$ is an A_∞-functor, then images of exact triangles are exact.*

Split-closure

Definition. *Let $\mathcal{A}$ be an A_∞-category. We say $\mathcal{A}$ is split-closed if, given $X \in \mathcal{A}$ and $e \in \mathrm{Hom}_{H^0\mathcal{A}}(X, X)$ such that $e^2 = e$, then X is quasi-isomorphic to $Y \oplus Z$ (for some Y and Z) such that e is the class of the canonical map $X \to Y$.*

Suppose $\mathcal{B} \subset \mathcal{A}$ is a full subcategory. We call the triangulated closure *of $\mathcal{B}$ inside $\mathcal{A}$ the intersection $\tilde{\mathcal{B}}$ of all triangulated full subcategories, closed under quasi-isomorphism, of $\mathcal{A}$, and say $\mathcal{B}$* generates *$\tilde{\mathcal{B}}$.*

If $\mathcal{B}$ is an A_∞-category, $\mathcal{A}$ a triangulated A_∞-category, and $F : \mathcal{B} \to \mathcal{A}$ is an A_∞-functor such that HF is fully faithful, we say that $\mathcal{A}$ is a triangulated envelope of $\mathcal{B}$ if $F(B)$ generates $\mathcal{A}$.

We similarly define the split-triangulated closure*,* split-closed *generation and* split-closed triangulated envelopes*.*

So for example the triangulated closure of $\mathcal{B}$ (inside $\mathcal{A}$) is obtained by adding a zero object, forming cones on all morphisms, including all quasi-isomorphic objects, and then repeating this process infinitely many times. The main point we are getting at is the following theorem. Although stated abstractly, it is usually proved by explicit constructions.

Theorem 2.4. *Both triangulated envelopes and split-closed triangulated envelopes exist and are unique up to quasi-equivalence.*

The theorem stated in the introduction follows directly from this (but is much weaker). Given an A_∞-category $\mathcal{A}$, we denote by $D\mathcal{A}$ any triangulated envelope, and by $D^\pi\mathcal{A}$ any split-closed triangulated envelope.

2.4. Twisting

Let $\mathcal{A}$ be an A_∞-category which is closed under arbitrary shifts and finite direct sums. Let $Y_0, Y_1 \in \mathcal{A}$ be such that $\mathrm{Hom}_{H(\mathcal{A})}(Y_0, Y_1)$ is finite-dimensional. We can then form the tensor product $\mathrm{Hom}_{H(\mathcal{A})}(Y_0, Y_1) \otimes Y_0$. This comes with a natural evaluation homomorphism $[f]$ to Y_1, of degree zero. If $\mathrm{Hom}_{H(\mathcal{A})}(Y_0, Y_1) \otimes Y_0 \xrightarrow{[f]} Y_1$ can be completed to an exact triangle, we denote the third term by $T_{Y_0}(Y_1)$, and call it the *twist of* Y_1 *along* Y_0.

We say that $\mathcal{A}$ is *cohomologically finite* if $\mathrm{Hom}_{H\mathcal{A}}(X, Y)$ is finite-dimensional for all $X, Y \in \mathcal{A}$.

Proposition 2.5 ([20], lemma 5.4)**.** *Let $\mathcal{A}$ and $\mathcal{B}$ be triangulated, cohomologically finite A_∞-categories. Then:*

1. *For any $Y \in A$, there exists a functor $T_Y : \mathcal{A} \to \mathcal{A}$ extending the above construction. It is canonical in $H(\mathcal{A})$.*
2. *If $F : \mathcal{A} \to \mathcal{B}$ has the property that HF is fully faithful, then the functors $T^{\mathcal{B}}_{FX}(F\bullet)$ and $FT^{\mathcal{A}}_X(\bullet)$ are canonically quasi-isomorphic.*

Application one: the total endomorphism ring

Our main application of twisting is as follows: suppose $\mathcal{A}$, $\mathcal{B}$ and F are as in the theorem. Consider objects X and Y of $\mathcal{A}$. We can form the graded vector space $R(X, Y) = \bigoplus_{n \geq 0} \mathrm{Hom}_{H^0(\mathcal{A})}(X, T^n_Y(X))$. This can be given the structure of a (non-commutative) ring: given $f \in \mathrm{Hom}_{H^0(\mathcal{A})}(X, T^n_Y(X))$ and $g \in \mathrm{Hom}_{H^0(\mathcal{A})}(X, T^m_Y(X))$, we can form the composite $T^m_Y(f) \circ g \in \mathrm{Hom}_{H^0(\mathcal{A})}(X, T^{m+n}_Y(X))$.

The canonicity of the isomorphism in the above theorem then immediately implies the following:

Proposition 2.6. *In the above situation, $R(X, Y)$ and $R(FX, FY)$ are isomorphic as graded rings. In particular $R(X, Y)$ is invariant under auto-quasi-equivalences of $\mathcal{A}$ fixing X and Y.*

Application two: generation criteria

One interesting point regarding twists is that when composing several of them, they still compute (increasingly complicated) cones. This follows from parts (2) and (3) of proposition 2.3 summarizing the properties of exact triangles.

Lemma. *Let $\mathcal{A}$ be a triangulated A_∞-categories and $X, Y_1, \ldots, Y_n \in \mathcal{A}$. Let $\tilde{\mathcal{A}}$ denote the triangulated subcategory of $\mathcal{A}$ generated by $Y_1, \ldots, Y_n$.*

Then there exists an object $Z \in \tilde{\mathcal{A}}$ and an exact triangle $X \to T_{Y_n} \ldots T_{Y_1}(X) \to Z \to$.

Theorem 2.7. *Let $\mathcal{A}$ be a split-closed triangulated A_∞-category and $X, Y_1, \ldots, Y_n \in \mathcal{A}$. Let $X' = T_{Y_n} \ldots T_{Y_1}(X)$, and $\tilde{\mathcal{A}}$ the split-closed triangulated subcategory of $\mathcal{A}$ generated by $Y_1, \ldots, Y_n$.*

If $\mathrm{Hom}_{H^0(\mathcal{A})}(X, X') = 0$, then X lies in $\tilde{\mathcal{A}}$.

Proof. By the lemma, there is an exact triangle $X \to X' \to Z \to$ with $Z \in \tilde{\mathcal{A}}$. By assumption, $X \to X'$ is the zero morphism, so $Z \simeq X' \oplus X[1]$. Hence $X[1] \in \tilde{\mathcal{A}}$, and so the same holds for X. □

3. The enhanced derived category

We now construct our first interesting A_∞-category, the derived category of coherent sheaves on a scheme. For a basic reference on abelian categories and homological algebra, see [23]. For a general reference for scheme theory, see [8].

Suppose $\mathcal{B}$ is an abelian category with enough injectives, and $\mathcal{A}$ is a full subcategory. Then for every bounded complex $A^\bullet$ in $\mathcal{A}$, there exists a quasi-isomorphic left-bounded complex of injectives $I_A^\bullet$ *in* $\mathcal{B}$ (for example a Cartan-Eilenberg resolution, see [23, section 5.7]). By standard arguments, any two such complexes are quasi-isomorphic, uniquely up to homotopy.

Suppose such a choice I_A has been made for every bounded complex A. Then we denote the *enhanced derived category* of $\mathcal{A}$ (relative to $\mathcal{B}$) by $D_\infty(\mathcal{A}, \mathcal{B})$. We define $Ob(D_\infty(\mathcal{A}, \mathcal{B}))$ to consist of the bounded complexes in $\mathcal{A}$, and $\mathrm{Hom}_{D_\infty(\mathcal{A},\mathcal{B}}(A^\bullet, B^\bullet) = \mathrm{Hom}^\bullet(I_A, I_B)$. This is a dg-category in a natural way.

Proposition 3.1. *The dg-category $D_\infty(\mathcal{A}, \mathcal{B})$ is independent of the choices of resolving complexes I_A, up to quasi-isomorphism.*

Moreover

$$H^* \mathrm{Hom}_{D_\infty(\mathcal{A},\mathcal{B})}(A, B) \simeq \mathrm{Ext}^*(A, B).$$

Proof. Let I'_A be another set of resolutions. By general results about injective objects, the identity morphism on A lifts to inverse quasi-isomorphisms $\alpha_A : I_A \to I'_A$ and $\beta_A : I'_A \to I_A$. One checks that the morphisms $f_A = (\alpha^*_A, \beta_{B*}) : \mathrm{Hom}^\bullet(I_A, I_B) \to \mathrm{Hom}^\bullet(I'_A, I'_B)$ define the F_1 term of a dg-functor between the two categories (which is the identity on objects). The F_2 term consists of homotopies between $\beta_A\alpha_A$ and $\mathbf{1}$.

It hence suffices to show that the f_A are quasi-isomorphisms. To see this, one can use the weakly convergent spectral sequence $E_1^{p,q} = H^p \mathrm{Hom}_{\mathcal{A}}(C^{-\bullet}, I_B^q) \Rightarrow H^{p+q} \mathrm{Hom}^\bullet(C, I_B)$. A quasi-isomorphism $C \to C'$ induces an isomorphism on the E_1 page, and hence an isomorphism in cohomology. Applying this to $A \to I_A$, we see that $H^* \mathrm{Hom}^\bullet(I_A, I_B) \simeq \mathrm{Ext}^*(A, B)$. Hence f_A is a quasi-isomorphism. □

Definition. *Let X be an algebraic variety (scheme of finite type over a field). Then the* enhanced derived category *of X, denoted $D_\infty(X)$, is the dg-category $D_\infty(Coh(X), QCoh(X))$, well-defined up to quasi-isomorphism.*

Here $Coh(X)$ denotes the category of coherent sheaves, and $QCoh(X)$ the category of quasi-coherent sheaves.

One may show that $D_\infty(X)$ is in fact split-closed and triangulated. The triangulation is essentially obvious, and split-closedness is proved for example in [19, section (5c)]. We do not really need these results, but it means that we do not have to write (or consider) $D^\pi D_\infty(X)$.

One may actually show that $H^0D_\infty(X)$ is the category obtained from $D^b(Coh(X))$ by inverting quasi-isomorphisms (this does not hold for all pairs $(\mathcal{A}, \mathcal{B})$ as above), but we do not actually need this. We now prove two generation results for the enhanced derived category.

Theorem 3.2 ([19], lemma 5.4)**.** *Suppose X is a smooth projective variety with a very ample line bundle $\mathcal{O}(1)$. Then $D_\infty(X)$ is generated (as a triangulated A_∞-category) by the locally free sheaves. Moreover $D_\infty(X)$ is* split-generated *by the line bundles $\mathcal{O}_X(n)$ for $n \in \mathbb{Z}$.*

Proof. It is clear that $D_\infty(X)$ is generated by $Coh(X)$. We show first that every coherent sheaf has a free resolution. By the classification of quasi-coherent sheaves on a projective scheme, every coherent sheaf $\mathcal{F}$ admits a surjection $\mathcal{O}(-n)^m \to \mathcal{F}$ for sufficiently large n and m. Fix $\mathcal{F} \in Coh(X)$. We can form a resolution

$$0 \to \mathcal{R} \to \mathcal{O}(-n_l)^{m_l} \to \cdots \to \mathcal{O}(-n_1)^{m_1} \to \mathcal{F} \to 0. \tag{2}$$

Recall now the sheaf $\mathcal{E}xt$ functor (see [8, III.6]). By a standard (dimension shifting) argument, $\mathcal{E}xt^i(\mathcal{R}, \mathcal{G}) = \mathcal{E}xt^{i+l}(\mathcal{F}, \mathcal{G})$. We claim this is zero for any $\mathcal{G} \in QCoh(X)$ and $i > \dim X$. Indeed, it suffices to prove this stalkwise, and $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})_x = \mathrm{Ext}^i_{\mathcal{O}_{X,x}}(\mathcal{F}_x, \mathcal{G}_x) = 0$ for $i > \dim X \geq \mathrm{cd}\, \mathcal{O}_{X,x}$, by Serre's theorem on the cohomological dimension of regular local rings. It follows that, if $l > \dim X$, then $\mathcal{R}_x$ is a free module for any $x \in X$ (since finitely generated projective modules on a Noetherian local ring are free), and hence $\mathcal{R}$ is a locally free sheaf (because X is Noetherian). This establishes the first claim.

For the second claim, we observe that the exact sequence (2) yields an exact triangle $\mathcal{R}[l-1] \to \mathcal{L}^\bullet \to \mathcal{F} \to$ in $D_\infty(X)$, where $\mathcal{L}^\bullet$ denotes the complex built from $\mathcal{O}(n)$-s. By rotation, we get an exact triangle $\mathcal{F} \to \mathcal{R}[l] \to \mathcal{L}^\bullet[1] \to$. However, $\mathrm{Hom}_{H^0D_\infty(X)}(\mathcal{F}, \mathcal{R}[l]) = \mathrm{Ext}^l(\mathcal{F}, \mathcal{R})$ by the proposition. We claim this is also zero for l sufficiently large. If so, then $\mathcal{L}^\bullet[1]$ is quasi-isomorphic a cone on a zero morphism $\mathcal{F} \to \mathcal{R}[l]$, which is just the direct sum $\mathcal{F} \oplus \mathcal{R}[l]$, proving the theorem.

To establish the claim, using the $\mathcal{E}xt$-to-Ext spectral sequence, one sees that if $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G}) = 0$ for $i > n$, then $\mathrm{Ext}^i(\mathcal{F}, \mathcal{G}) = 0$ for $i > n + \dim X$ (because $H^p(X, \mathcal{G}) = 0$ for $p > \dim X$ and any quasi-coherent sheaf $\mathcal{G}$). But this was already shown above (with $n = \dim X$). □

Theorem 3.3. *Let X be an algebraic curve (smooth projective variety of dimension one) over an algebraically closed field, and $P \in X$ a closed point.*

Then $D_\infty(X)$ is split-generated by $\mathcal{O}_X$ and $\mathcal{O}_X(P)$.

Proof. By the Riemann-Roch theorem and the embedding criteria ([8, remark II.7.8.2]), $\mathcal{O}(P)$ is ample. It hence suffices to show that the category $\mathcal{C}$ split-generated by $\mathcal{O}$ and $\mathcal{O}(P)$ contains $\mathcal{O}(nP)$ for $n \in \mathbb{Z}$. But we have an exact sequence $0 \to \mathcal{O}(nP) \to \mathcal{O}((n+1)P) \to \mathcal{O}/P \to 0$. The case $n = 0$ shows that $\mathcal{O}/P \in \mathcal{C}$. Then induction on n shows that $\mathcal{O}(nP) \in \mathcal{C}$ for any n. □

4. The Fukaya category

In this section we sketch how to construct the Fukaya category. This is very difficult in general, and we will not be able to give all details. We will focus on material which generalizes to higher dimensions, and restrict to the bare minimum of details applying only to the torus. For more complete accounts, see [19, sections 8 and 9] and [20].

The first two subsections review some basic notions from symplectic geometry. For a general overview, see for example [4].

4.1. Symplectic linear spaces

We recall that a *symplectic vector space* V is a finite-dimensional real vector space together with a non-degenerate alternating bilinear form ω. A *compatible complex structure* is an automorphism $I : V \to V$ such that $I^2 = -\mathbf{1}$ and $g(v, w) = \omega(v, Iw)$ is a positive-definite (bilinear symmetric) form. A compatible Riemannian inner product is defined similarly. A triple (ω, I, g) obtained in this way is called compatible, and it is easy to see that compatible triples always exist. If (ω, I, g) is a compatible triple, then $h = g + I\omega$ defines a Hermitian inner product on V considered as a complex vector space.

It is well-known that a symplectic vector space must have even dimension $2n$. A subspace $L \subset V$ is called *Lagrangian* if $\dim L = n$ and $\omega|_L = 0$, i.e. $\omega(v, w) = 0$ for all $v, w \in L$. Recall that there exists a topological space $Gr(m, V)$ of m-dimensional subspaces of V. We denote by $LGr(V) \subset Gr(n, V)$ the *Lagrangian Grassmannian* of all Lagrangian subspaces of V.

Choosing a compatible triple (ω, I, g), we obtain a Hermitian inner product on V. It is then easy to see that every Lagrangian subspace is obtained as the real span of a hermitian basis of V (as a complex vector space). Hence

$$LGr(V) \simeq U(n)/O(n).$$

It follows that $LGr(V)$ is connected (because $U(n)$ is), with fundamental group $\mathbb{Z}$ (consider the evident fibration).

Pick a compatible triple (ω, I, g) and write h for the associated Hermitian metric. Let $\Lambda^n_{\mathbb{C}} V$ denote the top exterior power of V as a complex vector space. This is a complex line, and it carries a hermitian metric. For a Lagrangian subspace $L = \langle e_1, \ldots, e_n \rangle$, where the e_i form a Hermitian basis of V, we write

$$s(L) = \frac{(e_1 \wedge \cdots \wedge e_n)^{\otimes 2}}{|e_1 \wedge \cdots \wedge e_n|^2}.$$

This is a well-defined map $LGr(V) \to S(\Lambda^n_{\mathbb{C}} V^{\otimes 2})$, where the right hand side denotes the unit circle of a complex line. It hence carries a natural orientation, and so a generator of H^1. Pulling back this generator by s we obtain an element $s_V \in H^1(LGr(V), \mathbb{Z})$ called the *Maslov class*. One may show that this is in fact a generator of $H^1(LGr(V), \mathbb{Z})$ and only depends on ω, not I and g. Thus if γ is a loop in $LGr(V)$, we can write $[\gamma] = n s_V$, and thus obtain an integer n.

The so-called *Maslov index* is a grand generalisation of this observation.

Theorem 4.1. *Fix two Lagrangian subspaces L_1 and L_2 of V, and denote by $P(V, L_1, L_2)$ the space of paths $\gamma : [0,1] \to LGr(V)$ with $\gamma(0) = L_1$, $\gamma(1) = L_2$. It carries a natural equivalence relation $\gamma_1 \sim \gamma_2$ given by homotopy relative to $\{0,1\}$.*

It is possible to define a set of functions $\mu : P(V, L_1, L_2) \to \mathbb{Z}$ with the following properties: (1) μ factors through $\sim$, (2) μ reduces to the Maslov index for loops defined above if $L_1 = L_2$, and (3) μ depends only on the symplectic structure of V, not on any other choices.

Suppose now that $\dim V = 2$. *Pick an orientation-preserving homeomorphism $S(\Lambda^n_{\mathbb{C}} V)^{\otimes 2} \simeq S^1$. Then if $\gamma : [0,1] \to LGr(V)$ is a loop, we have*

$$\mu(\gamma) = [f(1)],$$

where f denotes the unique lift of $s\gamma$ to $\mathbb{R}$ with $f(0) = 0$, relative to the covering $\mathbb{R} \to S^1, x \mapsto e^{2\pi i x}$, and $[x]$ is the greatest integer less than or equal to x.

The properties described in the theorem do not determine μ and are meant only for illustration. The first part of this is proved in [2], where the Maslov index is also characterized axiomatically, and constructed in various ways.

The second part of the theorem is an easy exercise, once the Maslov index has been defined.

4.2. Some symplectic geometry

We recall that a *symplectic manifold* is a smooth real manifold M, together with a closed, non-degenerate two-form $\omega = \omega_M$. Then for each point $x \in M$, the tangent space T_xM to M at x is naturally a symplectic vector space. A *symplectomorphism* of symplectic manifolds (M, ω) and (N, η) is a smooth map $f : M \to N$ such that $f^*\eta = \omega$. The group of symplectic automorphisms of M is denoted by $Sym(M) = Sym(M, \omega)$.

As before, a symplectic manifold necessarily has even dimension $2n$. A submanifold $L \subset M$ of dimension n is called *Lagrangian* if $\omega|_L = 0$. Equivalently, for all $x \in L$, $T_xL \subset T_xM$ is a Lagrangian subspace. If $\phi : M \to M$ is a symplectic automorphism, then $\phi(L)$ is also a Lagrangian submanifold.

Suppose now that X_t is a (time-dependent) vector field on M, and that M is compact. Recall that then there exists a unique *flow* $\phi_t : M \to M$ such that $\partial\phi/\partial t = X_t$. We say that X_t is a *symplectic vector field* if ϕ_t is a symplectomorphism. This happens if and only if the one-form $\eta_{X_t}(v) := \omega(X_t, v)$ is closed (this follows from Cartan's formula). In this case, we call the family ϕ_t a *symplectic isotopy* (from $\phi_0 = \mathbf{1}$ to ϕ_1). If η_X is not only closed but also exact, say $\eta_{X_t} = dH_t$, then we call ϕ_t a *Hamiltonian isotopy* corresponding to the *Hamiltonian function* H. Notice that non-degeneracy of ω implies that for each one-form η there exists precisely one vector field X such that $\eta_X = \eta$. Hence for each (time-dependent) Hamiltonian function $M \times [0,1] \to \mathbb{R}$, we get a Hamiltonian isotopy.

4.2.1. Dehn twists

We now describe how to produce particular symplectomorphisms on symplectic surfaces, called *Dehn twists*. We begin with the local model, namely the symplectic cylinder $X_\epsilon := S^1 \times (-\epsilon, \epsilon)$. We use coordinates s and t on S^1 and $(-\epsilon, \epsilon)$ respectively (we put $z = e^{2\pi i s} \in S^1$). The symplectic form is $\omega = \mathrm{d}s \wedge \mathrm{d}t$. Pick a smooth function $f_\epsilon : (-\epsilon, \epsilon) \to \mathbb{R}$ such that

$$\begin{aligned} f'(t) &= 0 \text{ for } t < -\epsilon/2 \\ f'(t) &= 1 \text{ for } t > \epsilon/2 \\ f''(t) &= f''(-t). \end{aligned}$$

Then we define a Hamiltonian function $H_\epsilon(s,t) = -f(t)$. Its time-one flow τ_{S^1} is called a *local Dehn twist*. It is given by

$$\tau_{S^1}(s,t) = \begin{cases} (s,t) & : |t| \geq \epsilon/2 \\ (s - f(t), t) & : |t| < \epsilon/2 \end{cases}.$$

Suppose now that S is a symplectic surface, and L is a Lagrangian circle (i.e. any embedded circle). By the Weinstein neighbourhood theorem [4, theorem 2.10], there exist $\epsilon > 0$, an open neighbourhood U of L in S and a symplectic isomorphism $U \simeq X_\epsilon$. Via this isomorphism, we can transplant τ_{S^1} to a symplectomorphism τ_L on U. Since τ_{S^1} is the identity near the boundary of X_ϵ, τ_L extends to all of S.

A very important observation is that even though we had to make a lot of choices in constructing τ_L (in particular U and f_ϵ), any other set of choices would yield a Hamiltonian isotopic symplectomorphism. Thus we will in the future speak of "the" Dehn twist in L.

Note in particular that τ_L does *not* depend on an orientation of L. Namely, ω induces an orientation of S, and this alone determines "which direction to twist in", as illustrated in figure 1, depicting a Dehn twist of the torus. The twisting direction may seem slightly unnatural, but is dictated by certain less-elementary ways of constructing Dehn twists (using Lefschetz fibrations).

4.3. Graded symplectic geometry

We next describe an extension of the above ideas to incorporate "gradings". All of this material can be found in [18].

First recall that on any manifold M, there exists a natural fibre bundle $\pi : Gr(n, M) \to M$ with fibres $Gr(n, T_xM)$. If M is a symplectic manifold of dimension $2n$, we can form the subbundle $LGr(M) \subset Gr(n, M)$ of Lagrangian subspaces, called the *Lagrangian Grassmannian* of M. For any Lagrangian $L \subset M$, there exists a natural section $s_L : L \to LGr(M)$ (given by $s_L(x) = T_xL$). If $\phi : M \to M$ is a symplectomorphism, then we have a map $\phi_* : LGr(M) \to LGr(M)$ covering ϕ, given by $\phi_*(x, L_x) = (\phi(x), D\phi L_x)$. For any Lagrangian L, it satisfies $\phi_* s_L = s_{\phi(L)} \phi|_L$.

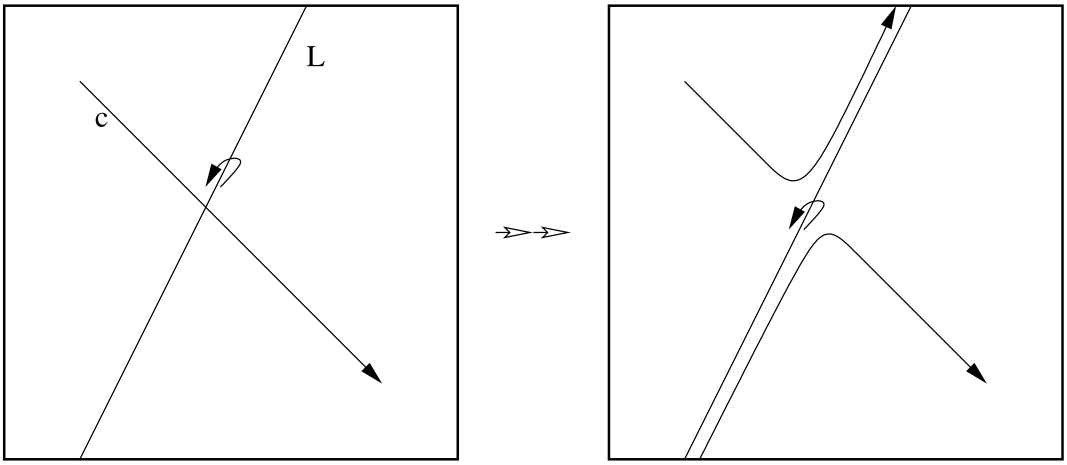

Figure 1: Effect of a Dehn twist τ_L on a test curve c.

Recall that, for a reasonably nice space X and abelian group A, the A-covers of X (i.e. covering spaces with a continuous and fibrewise transitive and faithful action of A) are classified by $H^1(X, A)$.

Definition. *Suppose* (M, ω) *is a symplectic manifold. We call a* Maslov cover *for* M *a* $\mathbb{Z}$*-cover* $\mathcal{L}$ *of* $LGr(M)$ *with the following property: if* $c \in H^1(LGr(M), \mathbb{Z})$ *classifies* $\mathcal{L}$, $x \in M$ *and* $i_x : LGr(T_xM) \to LGr(M)$ *denotes the inclusion of the fibre, then* $i_x^*(c) = s_{T_xM}$.

Recall that s_{T_xM} denotes the Maslov class of $LGr(T_xM)$, which is a canonical generator of $H^1(LGr(T_xM), \mathbb{Z})$. Since $\pi_1(LGr(T_xM)) = \mathbb{Z}$, it follows in particular that $\mathcal{L}$ is a fibre bundle on M with fibres the universal covers of $LGr(T_xM)$. One may prove that Maslov covers exist if and only if $2c_1(M) = 0$ (where $c_1(M)$ denotes the first Chern class of the tangent bundle). In fact, choose a compatible almost-complex structure on M, and write $\Delta_M = \Delta(M, \omega, I) := \Lambda^n(TM, \mathbb{C})^{\otimes 2}$. Then Maslov covers of M are in bijection with homotopy classes of trivialisations of Δ_M (and in particular form an affine space over $H^1(M, \mathbb{Z})$), see [18, lemma 2.2] and the discussion preceding it.

The main reason for grading Lagrangians is as follows: suppose (L_1, t_1) and (L_2, t_2) are graded Lagrangians and $x \in L_1 \cap L_2$. Choose a path γ in $\mathcal{L}_x$ from $t_1(x)$ to $t_2(x)$. Since $\mathcal{L}_x$ is the universal cover of $LGr(T_xM)$, this path is unique up to homotopy. Hence $\mu_x(L_1, L_2) := \mu(\pi\gamma)$ (where the right hand side denotes the Maslov index for paths from theorem 4.1) is well-defined, and called the *Maslov index of* x.

We call a *graded manifold* a symplectic manifold (M, ω) together with a choice of Maslov cover $\mathcal{L}$. By a *graded Lagrangian* we mean a Lagrangian submanifold L together with a lift $t_L : L \to \mathcal{L}$ of s_L. By a *graded symplectomorphism* $\phi : (M, \omega_M, \mathcal{L}_M) \to (N, \omega_N, \mathcal{L}_N)$ we mean a symplectomorphism $\phi : M \to N$, together with a $\mathbb{Z}$-equivariant lift $\phi_* : \mathcal{L}_M \to \mathcal{L}_N$ of $\phi_* : LGr(M) \to LGr(N)$.

Graded symlectomorphisms can be composed in the natural way, and we write $Sym^{gr}(M)$ for the group of graded symplectomorphisms of M. There is a central subgroup $\{\mathbf{1}[n] | n \in$

$\mathbb{Z}\} \simeq \mathbb{Z}$, where $\mathbf{1}[n]$ is given by $\phi = \mathbf{1}$ and $\phi_*(x) = x+n$ (using the $\mathbb{Z}$-action on the cover). For any graded automorphism ϕ, we write $\phi[n] := \mathbf{1}[n]\phi = \phi\mathbf{1}[n]$. Note that $Sym^{gr}(M)$ acts naturally on the set of graded Lagrangians. If L is such a graded Lagrangian, then we denote by $L[n]$ the graded Lagrangian $\mathbf{1}[n](L)$. Furthermore, if ϕ_1 and ϕ_2 are joined by a symplectic isotopy, then a grading of ϕ_1 induces a grading of ϕ_2.

One may describe cohomologically which symplectomorphism and Lagrangians are gradable, and how many possible gradings there are, see [18, lemmas 2.3 and 2.4]. We will not need this.

4.3.1. A more explicit description

Let (M, ω) be a symplectic manifold, and choose a compatible almost-complex structure. Suppose we are given a non-vanishing section Ω of Δ_M, with $|\Omega| \equiv 1$ (i.e. a trivialisation of Δ). Put $\mathcal{L}_\Omega = LGr(M) \times_{s,\Delta,c} \mathbb{R} = \{(L,r) | s(L) = c(r)\}$. Here $\mathbb{R}$ denotes the trivial bundle, $c : \mathbb{R} \to S(\Delta)$ the natural cover induced by Ω and $s : LGr(M) \to \Delta$ is the natural map. This is a Maslov cover, and one may show that in fact all Maslov covers are obtained in this way.

Suppose now L is a Lagrangian submanifold of M, and $t_L : L \to \mathcal{L}_\Omega$ is a grading. Then we can write $t_L(x) = (s(s_L(x)), \tilde{s}_L(x))$ for some function $\tilde{s}_L : L \to \mathbb{R}$. One has $e^{2\pi i \tilde{s}_L(x)}\Omega(x) = s(s_L(x))$, and such (continuous) functions are in bijection with gradings of L.

Similarly, if $\phi : M \to M$ is a symplectic automorphism, then giving a grading for ϕ is the same as giving a continuous function $\tilde{s}_\phi : LGr(M) \to \mathbb{R}$ such that $s(D\phi(L)) = e^{2\pi i \tilde{s}_\phi(L)}\Omega(x)$, for any $x \in M$ and $L \in LGr(M)_x$.

Suppose now that $(\phi, \tilde{s}_\phi)$ and $(\psi, \tilde{s}_\psi)$ are graded symplectomorphisms, and that $(L, \tilde{s}_L)$ is a graded Lagrangian. Then one easily verifies that

$$\tilde{s}_{\psi\phi} = \tilde{s}_\phi + \tilde{s}_\psi \circ D\phi$$
$$\tilde{s}_{\phi(L)} = \tilde{s}_L \circ \phi^{-1} + \tilde{s}_\phi \circ s_L \circ \phi^{-1}.$$

4.3.2. Example: the flat torus

We consider now the torus $T = \mathbb{C}/\mathbb{Z}^2$. This carries a natural (constant) symplectic form and holomorphic structure. We obtain a trivialisation $\Omega = \partial/\partial z^{\otimes 2}$ of Δ_T. Up to homotopy, all others are of the form $\Omega_{m,n}(x,y) = e^{2\pi i(mx+ny)}\Omega$. We thus get corresponding Maslov covers $\mathcal{L}_{m,n}$ for $m, n \in \mathbb{Z}$, and by the above discussion this yields all of them. It is now easy to check that, with respect to the cover $\mathcal{L} = \mathcal{L}_{0,0}$ corresponding to the holomorphic trivialisation Ω, a Lagrangian circle is gradable if and only if it is not contractible. In contrast, if $(m,n) \neq (0,0)$, then contractible circles are still not gradable, but there always exist further (non-contractible) circles which are not gradable either.

Hence the holomorphic trivialisation is a natural choice of Maslov cover. Moreover, for this choice of Maslov cover, *any* symplectic automorphism of T is gradable. (All of these statements are easy exercises.)

4.4. The Fukaya category of symplectic surfaces

We now describe the Fukaya category a compact symplectic surface with choice of Maslov cover (i.e. the torus, with some choice of symplectic form and Maslov cover). A closely related situation is described in [1], and also in our main reference [19].

Suppose $L_1, L_2, \dots, L_n$ are graded Lagrangian circles in T (i.e. compact connected Lagrangian subspaces with choices of grading) which are pairwise transverse. Then for $i \neq j$, the set $L_i \cap L_j$ is *finite*. Denote by Λ_t the field of formal power series $\sum_{n>0} a_n t^{b_n}$ in t, with complex coefficients $a_n \in \mathbb{C}$ and *real* exponents $b_n \in \mathbb{R}$ tending to infinity as $n \to \infty$. We then form the *Floer chain space*

$$CF^*(L_i, L_j) = \bigoplus_{x \in L_i \cap L_j} [x]\Lambda_t,$$

where $[x]$ is a formal generator in degree $\mu_x(L_i, L_j)$.

We will describe how to define an operation $m_n : CF^*(L_n, L_{n-1}) \otimes \dots \otimes CF^*(L_1, L_2) \to CF^*(L_1, L_n)$ in such a way that the A_∞-relations are satisfied whenever they make sense. Specifically, we will put

$$m_n(p_{n-1}, \dots, p_1) = \sum_{q \in L_1 \cap L_n} C(q; p_1, \dots, p_{n-1})[q].$$

Here $p_i \in L_i \cap L_{i+1}$. We now explain how to compute the coefficients $C(q; p_1, \dots, p_{n-1})$.

Let D_n denote the unit disc, with marked points $P_1, P_2, \dots, P_{n-1}$ and Q on the boundary (in that order). We denote the boundary segment ending at P_i by γ_i. We write $\mathcal{M}(q; p_1, \dots, p_{n-1})$ for the set of equivalence classes of immersed polygons $u : D_n \to X$ with the following properties: (1) u is orientation-preserving and an embedding away from P_i or Q, (2) u has convex corners at P_i and Q, (3) $u(P_i) = p_i$ and $u(Q) = q$, and (4) $u(\gamma_i) \subset L_i$. Two embedded polygons are viewed as equivalent if they differ by a diffeomorphism of D_n.

We then put $C(q; p_1, \dots, p_{n-1}) = \sum_{u \in \mathcal{M}(q; p_1, \dots, p_{n-1})} \pm t^{\int_{D_n} u^*\omega}$. Here $\pm$ denotes a sign which we are not going to explain[1]

Unfortunately we have no space here to explain why this yields maps satisfying the A_∞-relations. The following theorem summarizes the properties of the Fukaya category we are going to use.

Theorem 4.2. *Let $(T, \omega, \mathcal{L})$ be a graded manifold, where T is the torus. The Fukaya category $\mathrm{Fuk}(T)$ has as objects the graded Lagrangians (L, t_L). Moreover:*

1. *If $L_1, L_2 \in \mathrm{Fuk}(T)$ are transverse, then $\mathrm{Hom}_{\mathrm{Fuk}(T)}(L_1, L_2) = CF^*(L_1, L_2)$. If $L_1, \dots, L_n$ are transverse, then $m_n^{\mathrm{Fuk}(T)} : \mathrm{Hom}(L_{n-1}, L_n) \otimes \dots \otimes \mathrm{Hom}(L_1, L_2) \to \mathrm{Hom}(L_1, L_n)$ is the map constructed above.*

[1] We will not be needing the sign, only the fact that there exists a sign rule which makes the theorem below true.

2. *There is an action of* $Sym^{gr}(T)$ *on* $\mathrm{Fuk}(T)$. *On objects, it is the action constructed previously. Moreover, if* ϕ *is a graded symplectomorphism which is Hamiltonian isotopic to the identity, then its action on* $\mathrm{Fuk}(T)$ *is quasi-isomorphic to the identity functor.*

3. *If* L *is a Lagrangian circle,* $\tilde{\tau_L}$ *denotes the Symplectic Dehn twist in* L *with its standard grading, and* $L' \in \mathrm{Fuk}(T)$, *then* $T_L(L') \simeq \tilde{\tau}_L(L')$. *Here* T_L *denotes the algebraic twist functor from section 2.4.*

5. Hochschild cohomology and classification of A_∞-structures

In this section we show how Hochschild cohomology can be used to classify minimal A_∞-structures. In the first subsection, we reinterpret A_∞-algebras in terms of graded coalgebras. This material is summarized most succinctly in [9, 3.6]. In the second subsection we define Hochschild cohomology and explain some of its properties.

Much of the material can be found in [5] and also in [19, section 3].

5.1. A_∞-structures and coalgebra codifferentials

We first reinterpret the A_∞-equations in a somewhat less ad-hoc way. For this, recall that a *graded coalgebra* is a graded vector space V, together with a comultiplication map $\Delta : V \to V \otimes V$, satisfying a certain "associativity condition" dual to the definition of an algebra. Similarly, a coderivation of degree d is an element of $\mathrm{Hom}(V, V(d))$ satisfying a certain relation dual to the definition of a derivation on an algebra. Finally, we recall that the notion of a coalgebra morphism is also defined in the evident way. We remark that both Δ and coalgebra morphisms are automatically of degree zero.

Our main example of a coalgebra is as follows: let V be a graded vector space. Then the *reduced cotensor coalgebra* on V is

$$\overline{T}V = \bigoplus_{n>0} V^{\otimes n}.$$

We will write $v_1 \dots v_n$ for $v_1 \otimes \dots \otimes v_n$, and grade $\overline{T}V$ by $|v_1 \dots v_n| = |v_1| + \cdots + |v_n|$. We write $\overline{T}^n V = V^{\otimes n}$, seen as a subspace of $\overline{T}V$. The comultiplication is defined by

$$\Delta(v_1 \dots v_n) = \sum_{k=1}^{n} (v_1 \dots v_k) \otimes (v_{k+1} \dots v_n).$$

This coalgebra behaves in many ways "dually" the ordinary tensor algebra. Its properties are summarized in the following proposition, the proof of which is an easy exercise.

Proposition 5.1. *Let* V *and* W *be graded vector spaces. Write* $\pi : \overline{T}V \to V$ *for the projection onto the first summand. For any graded vector space* U *and linear map* $f : U \to \overline{T}V$, *denote by* $\pi_*(f)$ *the composite* $\pi \circ f$.

1. *The map*

$$\{\textit{coderivations } \overline{T}V \to \overline{T}V \textit{ of degree } d\} \xrightarrow{\pi_*} \mathrm{Hom}_k(\overline{T}V, V(d)) = \prod_{n>0} \mathrm{Hom}_k(V^{\otimes n}, V(d))$$

is a bijection, with inverse $(b_i) \in \prod_{n>0} \mathrm{Hom}_k(V^{\otimes n}, V(d)) \mapsto b$*, where*

$$b|_{\overline{T}^n V} = \sum_{a+b+c=n} \mathbf{1}^{\otimes a} \otimes b_i \otimes \mathbf{1}^{\otimes b}.$$

2. *The map*

$$\{\textit{coalgebra morphisms } \overline{T}V \to \overline{T}W\} \xrightarrow{\pi_*} \mathrm{Hom}_k(\overline{T}V, W) = \prod_{n>0} \mathrm{Hom}_k(V^{\otimes n}, W)$$

is a bijection, with inverse $(f_i) \in \prod_{n>0} \mathrm{Hom}_k(V^{\otimes n}, W) \mapsto f$*, where*

$$f|_{\overline{T}^n V} = \sum_{r, a_1+\cdots+a_r=n} f_{a_1} \otimes \ldots \otimes f_{a_r}.$$

We will in the sequel also write $\pi_*(f)_i \in \mathrm{Hom}(V^{\otimes}i, V(d))$ for the restriction of $\pi_*(f)$ to $\overline{T}^i V$.

The main reason for us to introduce these coalgebras is the relation to A_∞-algebras. Again, the following proposition is an easy exercise in writing out definitions. For a graded vector space V, we denote by SV the graded vector space $V(1)$ and by $s : V \to SV$ the natural (degree -1) *shift map.*

Proposition 5.2. 1. *Let A be a graded vector space. An A_∞-structure on A is the same as a degree one codifferential on $\overline{T}SA$. Namely, given a codifferential $(b_i) \in \prod_{n>0} \mathrm{Hom}_k((SA)^{\otimes n}, SA(1))$, the maps $m_i = s^{-1} b_i s^{\otimes i} \in \mathrm{Hom}_k(A^{\otimes n}, A(2-i))$ define an A_∞-structure if and only if $b^2 = 0$.*

2. *Let A and B be A_∞-algebras, with associated coderivations b_A and b_B. Then A_∞-morphisms $A \to B$ are the same as coalgebra morphisms $f : \overline{T}SA \to \overline{T}SB$ such that $fb_A = b_B f$.*

For the remainder of this section, we fix a graded vector space A. Let us write $C^{m,n} = \mathrm{Hom}_k((SA)^{\otimes m}, SA(n))$ and $C^n = \mathrm{Hom}_k(\overline{T}SA, SA(n)) \simeq \prod_r C^{r,n}$. The above proposition allows us to view C^n as a space of coderivations on $\overline{T}SA$. For b_1 and b_2 graded coderivations, we write $[b_1, b_2] = b_1 b_2 - (-1)^{|b_1||b_2|} b_2 b_1$ and call it the *(super-)Lie bracket.* It is easy to check that it satisfies the *graded Jacobi identity*

$$(-1)^{|x||z|}[x,[y,z]] + (-1)^{|y||x|}[y,[z,x]] + (-1)^{|z||y|}[z,[x,y]] = 0. \tag{3}$$

One may also check that the Lie bracket is compatible with the gradings, in the sense that $[C^{m,n}, C^{p,q}] \subset C^{m+p-1,n+q}$.

5.2. Hochschild cohomology

Definition. *Let A be a graded vector space. We write* $\mathit{CC}^{m,n} = \mathit{CC}^{m,n}(A) = \mathrm{Hom}_k(A^{\otimes m}, SA(n))$ *for the* Hochschild chain spaces*, and* $\mathit{CC}^n = \prod_{p+q=n} \mathit{CC}^{p,q}$.

Via $s : A \to SA$, we have isomorphisms (as a graded vector spaces) $\mathit{CC}^{m,n} \simeq C^{m,m+n-1}$ and $\mathit{CC}^n \simeq C^n$. This allows us to transplant the super-Lie structure to $\mathit{CC}^{\bullet,\bullet}$.

Suppose now that A is a graded algebra. Recall that this is just a special case of an A_∞-algebra, where $m_i = 0$ for $i \neq 2$, and so proposition 5.2 yields a degree one codifferential $b \in \mathit{CC}^{2,0}$. We define, for $x \in \mathit{CC}^{m,n}$, the *Hochschild differential* $D(x) = [b, x]$. We have $D(\mathit{CC}^{m,n}) \simeq [b, C^{m,m+n-1}] \subset C^{m+1,m+n} \simeq \mathit{CC}^{m+1,n}$. The graded Jacobi identity (3) implies that D is in fact a differential, and that it satisfies a *graded Leinbniz rule.*

Definition. *The bigraded Hochschild cohomology of A is the cohomology $HH^{m,n}(A)$ of $\mathit{CC}^{\bullet,n}(A)$ with respect to D.*

One important observation is that the Hochschild differential defined in this way coincides precisely with the differential on the cohomological bar complex of A. This implies the important formula (where by A^{op} we denote the opposite graded algebra of A)

$$HH^{m,n}(A) \simeq \mathrm{Ext}^m_{A\otimes A^{op}}(A, A(n)).$$

We are now interested in the following problem: we say that an A_∞-structure (m_i) on A (equivalently, a codifferential (d_i)) *extends the multiplication* if $m_1 = 0$ and m_2 is the multiplication coming from A (equivalently, $d_1 = 0$ and $d_2 = b$). Recall from section 2.2 on homological perturbation theory that a formal diffeomorphism is specified by a family $f_i : A^{\otimes i} \to A(1-i)$ with f_1 invertible. We call it a *gauge transformation* if $f_1 = \mathbf{1}$. Notice that a formal diffeomorphism is the same as a coalgebra automorphism $\Phi : \overline{T}SA \to \overline{T}SA$ (by proposition 5.2 part (2), and proposition 2.2). We can then define a codifferential $\Phi^*(d) = \Phi^{-1} d \Phi$, which defines a new A_∞-structure on A. If Φ comes from a gauge transformation, we call these A_∞-structures gauge-equivalent.

We would like to classify A_∞-structures on A extending the multiplication, up to gauge-equivalence. The following theorem is another exercise in writing out the definitions.

Theorem 5.3. *Let A be an associative graded algebra, and d, d' be degree one codifferentials extending the multiplication.*

1. *Let $x \in \mathrm{Hom}_k((SA)^{\otimes r}, SA)$. If Φ denotes the formal diffeomorphism specified by*

$$\pi_*(\Phi)_n = \begin{cases} \mathbf{1} & \text{if } n = 1 \\ x & \text{if } n = r \\ 0 & \text{else,} \end{cases}$$

then we have

$$\Phi^*(d)_n = \begin{cases} b_k & \text{if } n \le r \\ b_{r+1} + D(x) & \text{if } n = r+1. \end{cases}$$

2. Suppose that there exist integers $r \geq 2$ and $0 \leq s \leq r-2$ such that

$$\pi_*(d)_n = \pi_*(d')_n \text{ for } n \leq r$$
$$\pi_*(d)_{2+n} = 0 \text{ for } 0 < n \leq s.$$

Then $\pi_*(d-d')_n \in CC^{n,2-n}$ is a Hochschild cocycle for $r < n \leq r+s+1$.

This is a fairly effective tool. Suppose, for example, that $HH^{d,2-d}(A) = 0$ for $d > 2, d \neq R$ and $\dim_k HH^{R,2-R} = 1$. Then if (m_i) is an A_∞-structure extending the multiplication, we can apply part (2) of the theorem, with $d' = b$ and $r = 2$, to conclude that m_3 defines a cocycle. If say $R > 3$, so that $HH^{3,-1}(A) = 0$, we can apply part (1) of the theorem to find a gauge-equivalent A_∞-structure with $m_3 = 0$. We can thus trivialize all higher multiplications up to m_R, at which point we find an obstruction class $[m_R]$ in $HH^{R,2-R}(A)$. If $[m_R] = 0$ we can keep trivialising, and thus show that the A_∞-structure is formal. Even if not, then if (m'_i) is another A_∞-structure, trivialised up to order $R-1$, and if $[m_R] = [m'_R]$ we can keep applying the theorem with d' corresponding to (m'_i) and thus show that the two A_∞-structures are in fact gauge-equivalent. So there is in fact at most a one-dimensional family of A_∞-structures (up to gauge equivalence) extending the multiplication.

If we do not restrict our attention to only gauge-equivalences, but allow all A_∞-isomorphisms, then we can do even better. Indeed for $\tilde{\epsilon} \in k^\times$, the formal diffeomorphism ϵ, with $\epsilon_1 : A \to A, x \mapsto \tilde{\epsilon}^{|x|}x$ and $\epsilon_n = 0$ else, has the property that

$$\epsilon^*(d)_i = \epsilon^{i-2} d_i.$$

Suppose now that k is algebraically closed. Then in the above situation, where $HH^2(A)$ is one-dimensional, there are at most two isomorphism classes of A_∞-structures extending the multiplication: the formal one (where $m_i = 0$ for all $i > 2$) and (provided it exists) another structure where $m_i = 0$ for $i < r$ but $[m_r] \neq 0 \in HH^{r,2-r}(A)$. Namely, given any two non-formal A_∞-structures trivialised up to order $r-1$, we can find $\tilde{\epsilon}$ such that $[\epsilon^*(m)_r] = [m'_r]$, and then the argument from above shows that the A_∞-structures are isomorphic.

6. Mirror symmetry for elliptic curves over Λ_t

We are now ready to state and prove homological mirror symmetry for elliptic curves. We will be working over the algebraically closed field Λ_t of characteristic zero. We let T denote the flat torus with its standard symplectic structure and the Maslov cover coming from the standard (holomorphic) trivialisation of Δ (see section 4.3.2). In this section we aim to prove the following theorem:

Theorem 6.1. *There exists an elliptic curve E over Λ_t and a quasi-equivalence of A_∞-categories*

$$D^\pi \mathrm{Fuk}(T) \to D_\infty(E).$$

We will use the method of proof sketched in the introduction. As said before, this basically applies the strategy developed in [19] to the one-dimensional case. Namely, we will show in subsection 6.3 that the two standard meridians A and B of T split-generate $\mathrm{Fuk}(T)$ (and hence a fortiori $D^\pi\mathrm{Fuk}(T)$). We thus need to study the A_∞-algebra of relations $\mathcal{Q} = \mathrm{Hom}_{D^\pi\mathrm{Fuk}(T)}(A \oplus B, A \oplus B)$. We denote its cohomology algebra by $Q = H(\mathcal{Q})$.

In subsection 6.1 we show how to compute its Hochschild cohomology. Here we will be working over an arbitrary field k of characteristic zero. In particular we will show that

$$HH^{d,2-d}(Q) = \begin{cases} k: & d = 6, 8 \\ 0: & d > 2, d \neq 6, 8 \end{cases}.$$

Theorem 5.3 about the relationship between Hochschild cohomology and A_∞-structures, together with the homological perturbation theory of theorem 2.1 now implies the following: Given any A_∞-algebra $\mathcal{Q}'$ with cohomology algebra Q, we can find a quasi-isomorphic A_∞-algebra $(Q, (m_d))$ with $m_1 = 0$ and m_2 the multiplication on Q, and additionally $m_d = 0$ for $d = 3, 4, 5, 7$. Furthermore, m_6 and m_8 will be cocycles, giving elements $a_4(\mathcal{Q}') \in HH^{6,-4}(Q)$ and $a_6(\mathcal{Q}') \in HH^{8,-6}(Q)$. Beware that even though we are using functional notation $a_i(\mathcal{Q}')$, a priori these cohomology classes depend on the choices we made in trivializing the A_∞-structure, and not only on $\mathcal{Q}'$. However, theorem 5.3 certainly implies that if $\mathcal{Q}''$ is another such A_∞-algebra, and we end up with $a_4(\mathcal{Q}') = a_4(\mathcal{Q}'')$ and $a_6(\mathcal{Q}') = a_6(\mathcal{Q}'')$, then $\mathcal{Q}'$ and $\mathcal{Q}''$ are quasi-isomorphic.

We now turn to the algebraic geometry side of mirror symmetry. In subsection 6.2 we will be working over an arbitrary algebraically closed field of characteristic zero. Theorem 3.3 about generators for the derived category of curves immediately implies that for any elliptic curve E and any (closed) point $P \in E$, the two sheaves $\mathcal{O}$ and $\mathcal{O}(P)$ split-generate $D_\infty(E)$. The cohomology algebra of the relations algebra, i.e. $\mathrm{Ext}^*(\mathcal{O} \oplus \mathcal{O}(P), \mathcal{O} \oplus \mathcal{O}(P))$ is easy to compute and seen to coincide with Q.

We then turn to the key insight of Lekili and Perutz [11]. Namely we will consider the entire Weierstrass family $\mathcal{X} = \{X_0 X_2^2 = X_1^3 + aX_0^2 X_1 + bX_0^3\} \subset \mathbb{P}^2 \times \mathbb{A}^2 \to \mathbb{A}^2$. The fibres $E_{a,b}$ are elliptic curves when smooth, and exhaust all elliptic curves over k. The group $\mathbb{G}_m$ acts on $\mathcal{X}$ in a natural way (by Weierstrass reparametrizations $x \mapsto u^2 x$ and $y \mapsto u^3 y$). We will show that the trivialisation procedure on the A_∞-algebra of relations can be carried out for all $E_{a,b}$ (including the singular ones!) uniformly. One then obtains classes $a_4(a, b) \in HH^{6,-4}(Q)$ and $a_6(a, b) \in HH^{8,-6}(Q)$ which depend polynomially on a, b and moreover intertwine the action of $\mathbb{G}_m$ (recall from the end of 5.2 that $HH^{*,2-*}$ also has a natural $\mathbb{G}_m$-action). This is then shown to imply, for purely formal reasons, that, up to choosing an appropriate basis, we have $a_4(a, b) = a$ and $a_6(a, b) = b$. In particular, (possibly singular) Weierstrass cubics yield all possible A_∞-structures on Q.

An important observation is that given the derived category D_∞ of an elliptic curve, and the two generators $\mathcal{O}$ and $\mathcal{O}(P)$, we can actually reconstruct the curve E. This is because the ring $\bigoplus_n \mathrm{Hom}_{H^0 D_\infty}(\mathcal{O}, \mathcal{O}(3nP))$ is in fact the projective coordinate ring of an embedding of E. One may identify this ring as of the form considered in proposition

2.6 about invariant rings of certain twist functors, so it is in fact an invariant of the quasi-equivalence class of D_∞.

Combining these ideas yields the important reconstruction theorem 6.4, which is the heart of the algebraic geometry and homological algebra results proved in this essay.

We should mention that in fact the classes a_4 and a_6 are (in a precise sense) independent of the choices made. This is shown in [11], but we do not need such a strong result here.

We can finally turn to the symplectic geometry side. In section 6.3 we will heavily use the properties of the Fukaya category stated in theorem 4.2. Using the identification of Dehn twists with algebraic twists, and theorem 2.7 on generation criteria using twist functors, we show that the two standard meridians A and B are indeed split-generators of $D^\pi \mathrm{Fuk}(T)$. We then compute the cohomology algebra of the A_∞-algebra of relations $\mathcal{Q}$ by counting (degenerate) polygons. This turns out to coincide with Q, as claimed. Specialising our algebraic work to $k = \Lambda_t$, this concludes the abstract proof of mirror symmetry.

It still remains to use the reconstruction theorem 6.4 to find out which elliptic curve over Λ_t is actually the mirror of T (and in particular to rule out the possibility that the mirror might be singular). This we will take up in the next section. The remainder of this section will be used to fill in the details of the proof of abstract mirror symmetry sketched above.

One final introductory remark. All that we have said would go through just as stated. However, it turns out that the mirror correspondence becomes most natural if we identify the two standard meridians A and B not with the sheaves $\mathcal{O}$ and $\mathcal{O}(P)$, but instead with $\mathcal{O}$ and $\mathcal{O}/P$. As a matter of fact, there exists an automorphism of $D_\infty(E)$ fixing $\mathcal{O}$ and interchanging $\mathcal{O}(P)$ and $\mathcal{O}/P$. The reason why we prefer to work with $\mathcal{O}(P)$ instead of $\mathcal{O}/P$ is purely for technical convenience (the former being a line bundle).

We will show later that $\mathcal{O}(P) = T_{\mathcal{O}/P}(\mathcal{O})$. Hence under the "correct" mirror correspondence, the mirror of $\mathcal{O}(P)$ is $B' = \tau_B(A)$, which is a line of slope $\pi/4$. Of course, since Dehn twists correspond to taking cones, the set $\{A, B'\}$ generates $\mathrm{Fuk}(T)$ if and only if the set $\{A, B\}$ generates. In fact, $\tilde{\tau}_B$ is an autoequivalence of $\mathrm{Fuk}(T)$ interchanging the two generating sets, and so we find that there is a natural quasi-isomorphism $\mathrm{Hom}_{D^\pi \mathrm{Fuk}}(A \oplus B, A \oplus B) \to \mathrm{Hom}_{D^\pi \mathrm{Fuk}}(A \oplus B', A \oplus B')$. So working with $\{A, B\}$ instead of $\{A, B'\}$ makes no difference.

6.1. The algebra Q

We begin by investigating the graded associative algebra Q which will turn out to be the cohomology algebra of our A_∞-algebras of relations. For this, we fix a field k of characteristic 0 (in fact any characteristic different from 2 and 3 would also work). We can most easily describe Q as the six-dimensional subalgebra of the matrix algebra $M_2(k[s]/s^2)$, where s is placed in degree one, spanned as a vector space by the following elements:

$$e = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad e_p = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad g_0 = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \quad f = \begin{pmatrix} s & 0 \\ 0 & 0 \end{pmatrix}, \quad f_p = \begin{pmatrix} 0 & 0 \\ 0 & s \end{pmatrix}, \quad g_1 = \begin{pmatrix} 0 & s \\ 0 & 0 \end{pmatrix}.$$

This *is* an algebra, with the following non-vanishing compositions of basis elements: $e_p f_p = f_p = f_p e_p$, $e_p e_p = e_p$, $ef = f = fe$, $ee = e$, $eg_1 = g_1 = g_1 e_p$, $e_p g_0 = g_0 = g_0 e$, $g_1 g_0 = f$ and $g_0 g_1 = f_p$. Its identity is $1 = e + e_p$.

The algebra Q is most easily visualised as a two-object quiver, as shown below:

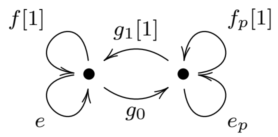

In the remainder of this section, we sketch how to establish the following result:

Proposition 6.2. *With the above notation, we have*

$$HH^{d,2-d}(Q) = \begin{cases} k: & d = 6, 8 \\ 0: & d > 2, d \neq 6, 8 \end{cases}.$$

Before that, we make plain the following important corollary, which motivates the entire computation. It spells out the implications of our Hochschild cohomology computation which can be obtained by applying the theory about the relationship between Hochschild cohomology and A_∞-structures of theorem 5.3, together with the homological perturbation theory of theorem 2.1.

Corollary. *Every A_∞-algebra $(\mathcal{Q}', m_d)$ with $H(\mathcal{Q}') \simeq Q$ is quasi-isomorphic to a minimal A_∞-algebra $(\tilde{\mathcal{Q}}, \tilde{m}_d)$ such that $\tilde{m}_d = 0$ for $d \leq 5$ or $d = 7$. The structure maps $\tilde{m}_6$ and $\tilde{m}_8$ define cocycles in $HH^{6,-4}(A)$ and $HH^{8,-6}(A)$.*

Furthermore if $(\mathcal{Q}', m_d)$ and $(\mathcal{Q}'', m'_d)$ are two minimal A_∞-algebras as above such that $[m_6] = [m'_6]$ and $[m_8] = [m'_8]$, then $\mathcal{Q}'$ and $\mathcal{Q}''$ are gauge-equivalent.

In particular, the space of A_∞-structures on Q up to quasi-isomorphism is at most one-dimensional.

The starting point of the proof is the observation that $HH^{m,n}(A) = \mathrm{Ext}^m_{A \otimes A^{op}}(A, A(n))$ (see section 5.2). One may thus establish this theorem by finding an explicit projective resolution of the finite-dimensional graded algebra $Q^e = Q \otimes Q^{op}$. This seems rather unenlightening. So we will first see what can be established by more abstract arguments.

The crucial observation now is that Q is a finite-dimensional algebra over a field. This opens up quite a lot of general theory, see e.g. [22]. (While this reference works with ungraded algebras and modules, the adaptations to the graded case are straightforward.)

The minimal idempotents of Q are $e_1 = e$ and $e_2 = e_p$. It follows that $P_i = Qe_i$ $(i = 1, 2)$ are minimal indecomposable projective Q-modules, and all of these are of the form $P_i(n)$ for some n. These projectives have unique non-trivial simple quotients S_i, and again in this way we obtain all the simple Q-modules. In fact, the S_i are one-dimensional k-vector spaces, spanned by the classes of e and e_p, respectively.

It follows from general theory that the minimal projective modules of Q^e are $P_{ij}(m) = (e_iQe_j)(m)$, and the simple modules are $S_{ij}(m) = \mathrm{Hom}_k(S_i, S_j)(m)$. Then, adapting the argument of lemma 1.5 in [7] to the graded case, we obtain the following:

Lemma. *Let* $\cdots \to R_n \to R_{n-1} \to \cdots \to R_1 \to R_0 \to Q$ *be a minimal projective resolution of* Q *as a* Q^e*-module. Then*

$$R_n = \bigoplus_{i,j,m} P_{i,j}(m)^{\mathrm{Ext}^n_Q(S_i,S_j(m))}.$$

Now it is easy to verify that there exist exact sequences

$$\begin{array}{ccccccccccc} 0 & \longleftarrow & S_1 & \longleftarrow & P_1 & \xleftarrow{\times g_0} & P_1 & \xleftarrow{\times f_p} & S_2(-1) & \longleftarrow & 0 \\ 0 & \longleftarrow & S_2 & \longleftarrow & P_2 & \xleftarrow{\times g_1} & P_2(-1) & \xleftarrow{\times f} & S_1(-2) & \longleftarrow & 0 \end{array}$$

and hence there exist projective resolutions as follows:

$$\cdots \to P_1(-3) \to P_1(-3) \dot{\to} P_2(-2) \to P_2(-1) \to P_1 \to P_1 \to S_1 \to 0$$
$$\cdots \to P_2(-4) \to P_2(-3) \dot{\to} P_1(-2) \to P_1(-2) \to P_2(-1) \to P_2 \to S_2 \to 0$$

Ignoring gradings, these resolutions are periodic of period four. The arrows with dots above them indicate the end of the first period, and of course the end bits $S_i \to 0$ are not repeated. After each period, gradings are increased by three.

Observing now that $\mathrm{Hom}_Q(P_i(n), S_j(m)) = \mathrm{Hom}_Q(S_i(n), S_j(m))$ is zero unless $i = j = n = m$, we deduce from the lemma that similarly Q admits over Q^e a projective resolution which, when ignoring gradings, is periodic of period eight, with gradings increasing by six after each period. In particular it follows that for any l, we have $HH^{n,l-n}(Q) = 0$ for n sufficiently large. Indeed it is a subquotient of $\mathrm{Hom}_{Q^e}(R_n, Q(l-n)) = \mathrm{Hom}_{Q^e}(R_n([6n/8]), Q(l-n+[6n/8]))$, and $R_n([6n/8])$ only depends on $n \pmod 8$ by periodicity of the resolution, whereas $Q(l-n+[6n/8])$ will eventually live in arbitrarily high degrees. Here $[x]$ denotes the greatest integer less than or equal to x.

Working a little more carefully, it is easy to find exactly the modules occurring in a minimal projective resolution of Q, and to give an explicit lower bound for n such that $HH^{n,2-n} = 0$. Then proving the proposition only boils down to finding a projective resolution of Q to finite length (e.g. actually determining the *maps* between the modules in the resolution) and computing the low-dimensional Hochschild cohomology. This is a pain best left to a computer. If a particularly masochistic reader wishes to verify this by hand, she will be aided by the fact that not only are the modules in the minimal projective resolution periodic, but the maps turn out to be periodic, too. Hence the whole proposition can be established by finding this periodic resolution. The computation has been carried out in detail in [12, theorem 4].

The proposition has also been reproved (by different means) in [11].

6.2. The derived category of an elliptic curve

Fix an algebraically closed field k of characteristic zero. Let E be a smooth projective curve and $P \in E$ a closed point. By theorem 3.3 on generators for derived categories of curves, we know that $\mathcal{L} = \mathcal{O}_E \oplus \mathcal{O}_E(P)$ is a split-generator. By the construction of $D_\infty(E)$ in section 3 and in particular proposition 3.1, we know that $\mathrm{Hom}_{HD_\infty(E)}(\mathcal{L}, \mathcal{L}) = \mathrm{Ext}^*(\mathcal{L}, \mathcal{L})$. The vector space structure of this graded algebra can be determined completely formally, and this we do first.

Indeed, recall some basic properties of Ext groups for schemes [8, section III.6]: (1) if $\mathcal{L}$ is locally free and $\mathcal{F}, \mathcal{G}$ are coherent, then $\mathrm{Ext}^*(\mathcal{L} \otimes \mathcal{F}, \mathcal{G}) = \mathrm{Ext}^*(\mathcal{F}, \mathcal{L}^\vee \otimes \mathcal{G})$, (2) $\mathrm{Ext}^*(\mathcal{O}_X, \mathcal{F}) = H^*(X, \mathcal{F})$ and (3) that Ext^* is bilinear. Here $\mathcal{L}^\vee$ denotes the dual sheaf. All three of these are easy consequences of the same statements for $\mathrm{Hom} = \mathrm{Ext}^0$ and universality of derived functors. Recall also Serre duality [8, section III.7]: if X is smooth projective of dimension n, then $H^k(X, \mathcal{L}) \simeq H^{n-k}(X, \mathcal{L}^\vee \otimes \omega_X)^\vee$.

We now get back to the case $X = E$ and $\mathcal{L} = \mathcal{O}_E \oplus \mathcal{O}_E(P)$. From now on we will also drop the subscript E. Using bilinearity, we can decompose the ext algebra into "matrix entries":

$$\mathrm{Ext}^*(\mathcal{L}, \mathcal{L}) = \begin{pmatrix} \mathrm{Ext}^*(\mathcal{O}, \mathcal{O}) & \mathrm{Ext}^*(\mathcal{O}(P), \mathcal{O}) \\ \mathrm{Ext}^*(\mathcal{O}, \mathcal{O}(P)) & \mathrm{Ext}^*(\mathcal{O}(P), \mathcal{O}(P)) \end{pmatrix}.$$

By this we mean that the left hand side decomposes, as a vector space, into a direct sum of the four subspaces on the right, and that multiplication on the left corresponds to matrix multiplication on the right together with the external multiplication on $\mathrm{Ext}^* = \mathrm{Hom}_{HD_\infty(E)}$. We can furthermore use properties (1) and (2) together with Serre duality to express this in terms of sheaf cohomology spaces.

Let us now specialise to the case of elliptic curves. We obtain the following.

Proposition 6.3. *Let E be an elliptic curve over an algebraically closed field k of characteristic zero, and $P \in E$ a closed point. Then*

$$\mathrm{Ext}^*(\mathcal{O} \oplus \mathcal{O}(P), \mathcal{O} \oplus \mathcal{O}(P)) \xrightarrow{\sim} Q,$$

where Q is the graded algebra from the previous section.

Proof. Almost all of the work has already been done. We have $H^0(E, \mathcal{O}) = k$ (i.e. there are no non-constant regular functions on E, this holds for any connected projective scheme), so in particular $H^0(E, \mathcal{O}(-P)) = 0$ (because $\mathcal{O}(-P)$ is the sheaf of regular functions vanishing at P). Also $H^0(E, \mathcal{O}(P)) = k$ because E cannot admit a rational function with a unique simple pole (otherwise it would be isomorphic to $\mathbb{P}^1$). Observing that the canonical bundle of elliptic curves is trivial, these computations suffice to show that our ext algebra is isomorphic to Q as a graded vector space.

But essentially all compositions in A are determined formally (from the action of identity elements of $\mathcal{O}$ and $\mathcal{O}(P)$, and degree reasons). Let h_0 denote a non-zero element of $\mathrm{Hom}(\mathcal{O}, \mathcal{O}(P))$ and h_1 a non-zero element of $\mathrm{Ext}^1(\mathcal{O}(P), \mathcal{O})$. To conclude the proof, we need only show that $h_0 h_1 \neq 0$ and $h_1 h_0 \neq 0$, since these are the only compositions which have not yet been determined.

We have an exact sequence $0 \to \mathcal{O} \to \mathcal{O}(P) \to \mathcal{O}/P \to 0$, where the first map is h_0. Applying $\mathrm{Ext}^*(\mathcal{O}(P), \bullet)$ we get $\mathrm{Ext}^1(\mathcal{O}(P), \mathcal{O}) \xrightarrow{\times h_0} \mathrm{Ext}^1(\mathcal{O}(P), \mathcal{O}(P)) \to \mathrm{Ext}^1(\mathcal{O}(P), \mathcal{O}/P)$. But $\mathrm{Ext}^1(\mathcal{O}(P), \mathcal{O}/P) = H^1(E, \mathcal{O}/P)$ (since $\mathcal{O}/P$ is a skyscraper sheaf), and this group is zero (skyscrapers being flasque). Thus right-multiplication by h_0 is surjective and $h_1 h_0 \neq 0$.

Applying instead $\mathrm{Ext}^*(\bullet, \mathcal{O})$ we get $\mathrm{Ext}^1(\mathcal{O}(P), \mathcal{O}) \xrightarrow{h_0 \times} \mathrm{Ext}^1(\mathcal{O}, \mathcal{O}) \to \mathrm{Ext}^2(\mathcal{O}/P, \mathcal{O})$. The same argument as before goes through as soon as we establish that $\mathrm{Ext}^2(\mathcal{O}/P, \mathcal{O}) = 0$. But this follows by applying $\mathrm{Ext}^*(\mathcal{O}, \bullet)$ to our exact sequence and observing that $H^2(X, \mathcal{F}) = 0$ for any line bundle $\mathcal{F}$, by Serre duality. □

We observe that the result is still true for singular cubics, if P is a smooth point, although one has to argue a bit more carefully for that.

Equivariant A_∞-structures and Weierstrass cubics

We would like to argue that the A_∞-structures on Q coming from elliptic curves essentially exhaust all possibilities (there are two exceptions, one formal and one not).

By this we mean the following. Fix a cubic curve E and a smooth point P, and set $\mathcal{Q}_{E,P} = \mathrm{Hom}_{D_\infty(E)}(\mathcal{O}_E \oplus \mathcal{O}_E(P), \mathcal{O}_E \oplus \mathcal{O}_E(P))$. This is a DGA, well-defined up to quasi-isomorphism, and by the proposition from the previous section, we have $H(\mathcal{Q}_{E,P}) \xrightarrow{\sim} Q$. Hence, by homological perturbation theory, we get a minimal A_∞-structure $(m_d^{E,P})$ on Q. We would like to say that in this way we obtain all isomorphism classes of A_∞-structures.

This is surprisingly tricky. Our method, following [11], will be to study how to turn all the $\mathcal{Q}_{E,P}$ *uniformly* into minimal A_∞-algebras.

More precisely, we put $R = k[a, b]$ and let $W = R(k)$. We then consider the family of curves $\mathcal{X} = \{X_0 Y^2 = X^3 + aX_0^2 X + bX_0^3\} \subset \mathbb{P}^2_R$, with its natural map to $Spec(R) = \mathbb{A}^2$. The multiplicative group $\mathbb{G}_m$ (with $\mathbb{G}_m(k) = k^\times$) acts on R by $a^u = u^4 a$ and $b^u = u^6 b$ for $u \in k^\times$. It acts on $\mathcal{X}$ by $((X_0 : X : Y), (a, b))^u = ((X_0 : u^2 X : u^3 Y), (u^4 a, u^6 b))$, covering the action on $Spec(R)$. For $w = (w_1, w_2) \in W$ we denote by $k(w) = R/(a - w_1, b - w_2)$ the residue field and by $\mathcal{X}_w = \mathcal{X} \times_R k(w)$ the fibre. The action of $\mathbb{G}_m$ on $\mathcal{X}$ induces isomorphisms $u : \mathcal{X}_w \to \mathcal{X}_{w^u}$.

We will construct a DGA $\mathcal{B}$, linear over R and with an action of $k^\times$. It will have the property that $\mathcal{B}_w := \mathcal{B} \otimes_R k(w)$ is quasi-isomorphic to $\mathcal{Q}_{\mathcal{X}_w, P_w}$. We will show that $H(\mathcal{B}) \otimes k(w) \to H(\mathcal{B}_w)$ is an isomorphism.

By an *equivariant splitting* of $\mathcal{B}$ we mean a decomposition $\mathcal{B} = H \oplus im(d) \oplus T$, where all the summands are $k^\times$-invariant, and $H \oplus im(d) = \ker(d)$. This yields an equivariant retraction $r : \mathcal{B} \to \mathcal{B}$ of degree -1 ($d : T \to im(d)$ is an isomorphism; let $r|_{im(d)} = d^{-1}$, $r = 0$ on the complement). If we denote by $i : H(\mathcal{B}) \to \mathcal{B}$ and $p : \mathcal{B} \to H(\mathcal{B}) = H$ the natural maps, then r is a homotopy between ip and $\mathbf{1}$. It follows that we can apply the homological perturbation lemma 2.1 to construct a minimal A_∞-structure on $H(\mathcal{B})$, linear over R, which we will denote by $\mathcal{A}$. Since $H = H(\mathcal{B})$ is $k^\times$-equivariant, $k^\times$ acts on $\mathcal{A}$. We will check that the action on degree i elements ($i = 0$ or $i = 1$ are the only possibilities) is pure of weight i, i.e. $h^u = u^{|h|} h$ for $h \in H$. For each $w \in W$ we get a minimal A_∞-structure $\mathcal{A}_w = \mathcal{A} \otimes_R k(w)$ on Q. We denote the structure maps by m_d^w, they are (vectors of) polynomials in the entries of w.

By an *equivariant A_∞-structure* (m_d) on a graded vector space V which comes with a G-action, we mean an A_∞-structure such that $m_d(a_d^g, \dots, a_1^g) = m_d(a_d, \dots, a_1)^g$ for all $d > 0$, $a_i \in V$ and $g \in G$. We will check that $\mathcal{B}$ is a $k^\times$-equivariant DGA. The explicit formulas for the A_∞-structure on $\mathcal{A}$ then imply that $\mathcal{A}$ is a $k^\times$-equivariant A_∞-algebra. This, together with our computation of the $k^\times$-action on $\mathcal{A}$, will imply that $m_d^{w^u} = u^{d-2} m_d^w$. Since the m_d^w are also polynomials, an easy computation will show that $m_d^w = 0$ for $d < 6$ or $d = 7$, that $m_6^{(a,b)} = a$ and that $m_8^{(a,b)} = b$. It follows that the A_∞-structures $\mathcal{A}_w$ *automatically* satisfy $m_i = 0$ for $i = 3, 4, 5, 7$, and exhaust all possibilities, by the corollary to proposition 6.2, computing the Hochschild cohomology of Q.

Before checking the details of the above argument, we record the following theorem which follows from it:

Theorem 6.4. *Suppose $\mathcal{C}$ is a split-closed triangulated A_∞-category, generated by objects O and $O(P)$. Assume further that $\mathrm{Hom}_{H^*(\mathcal{C})}(O \oplus O(P), O \oplus O(P)) \simeq Q$, where Q is the graded algebra from section 6.1.*

Let O/P denote the cone on $g_0 : O \to O(P)$, let S denote the graded ring $\bigoplus_n \mathrm{Hom}_{H^0(\mathcal{C})}(O, T_{O/P}^{3n}(O))$, and put $E = Proj(S)$. Then E can be embedded in $\mathbb{P}^2$ as a cubic curve (i.e. S is generated in degree one by three elements satisfying a homogeneous cubic relation), and if E is smooth, then there exists a quasi-equivalence of A_∞-categories between $\mathcal{C}$ and the enhanced derived category $D_\infty(E)$ of E.

Proof. Let $\mathcal{Q}'$ denote the A_∞-algebra structure on Q coming from $\mathcal{C}$. By our above classification, $\mathcal{Q}' \simeq \mathcal{A}_w$ for some w. If $\mathcal{X}_w$ is smooth, its enhanced derived category $D(\mathcal{X}_w)$ is split-generated by $\mathcal{O} \oplus \mathcal{O}(P)$, and so $\mathcal{C}$ is quasi-equivalent to $D(\mathcal{X}_w)$ by theorem 2.4 on uniqueness of split-triangulated envelopes. Even if $\mathcal{X}_w$ is not smooth, $\mathcal{C}$ is quasi-equivalent to the split-closed triangulated envelope of $\mathcal{O} \oplus \mathcal{O}(P)$, a subcategory of $D(\mathcal{X}_w)$ (containing $\mathcal{O}(nP)$ for all n). Moreover, the ring computed in $\mathcal{C}$ is isomorphic to the ring computed in $D_\infty(\mathcal{X}_w)$, by proposition 2.6 on total endomorphism rings defined via twist functors.

Observe now that $\mathrm{Ext}^*(\mathcal{O}/P, \mathcal{O}(nP))$ is one-dimensional, generated in degree one. Hence $T_{n+1} := T_{\mathcal{O}/P}(\mathcal{O}(nP))$ fits into the exact triangle $\mathcal{O}/P[-1] \to \mathcal{O}(nP) \to T_{n+1} \to$. We have an exact sequence $0 \to \mathcal{O}(nP) \to \mathcal{O}((n+1)P) \to \mathcal{O}/P \to 0$, yielding the exact triangle $\mathcal{O}/P[-1] \to \mathcal{O}(nP) \to \mathcal{O}((n+1)P) \to$. Thus $T_{n+1} \simeq \mathcal{O}((n+1)P)$, since each is the cone on a non-zero morphism, unique up to rescaling. Hence $T_{\mathcal{O}/P}^{3n}(\mathcal{O}) = \mathcal{O}(3nP)$.

It hence remains to show the following: fix $w \in W$. Let $S = \bigoplus_n \mathrm{Hom}_{D^b(Coh(\mathcal{X}_w))}(\mathcal{O}, \mathcal{O}(3nP))$. Then $Proj(S)$ is isomorphic to $\mathcal{X}_w$. But this is obvious ($\mathcal{O}(3P)$ being very ample). □

$\mathcal{X}$ is covered by two affine opens $U = \{X_0 \neq 0\}$ and $V = \{X_2 \neq 0\}$. We write $\mathcal{U} = \{U, V\}$ for this cover. $\mathcal{X}$ has the smooth subvariety $P = \{X_0 = 0\}$. We set $\mathcal{B} = \check{C}^\bullet(\mathcal{U}, \mathcal{H}om(\mathcal{O}_\mathcal{X} \oplus \mathcal{O}_\mathcal{X}(P), \mathcal{O}_\mathcal{X} \oplus \mathcal{O}_\mathcal{X}(P))$. Here by $\check{C}^\bullet(\mathcal{U}, \mathcal{F})$ we denote the Čech complex with respect to the cover $\mathcal{U}$ and coefficients in the sheaf $\mathcal{F}$.

Lemma. *$\mathcal{B}$ has the desired properties. That is:*

1. *$\mathcal{B}$ is a DGA,*
2. *$k^\times$ acts on $\mathcal{B}$,*
3. *the differential and multiplication of $\mathcal{B}$ are $k^\times$-equivariant,*
4. *$\mathcal{B}$ admits a $k^\times$-equivariant splitting,*
5. *$H(\mathcal{B}) \otimes k(w) \to H(\mathcal{B}_w)$ is an isomorphism, and $\mathcal{B}_w$ is quasi-isomorphic to $\mathcal{Q}_{\mathcal{X}_w,P_w}$*
6. *the weights of the $k^\times$-action on $H(\mathcal{B}_w)$ coincide with the degree.*

Proof. We put $\mathcal{F} = \mathcal{H}om(\mathcal{O}_\mathcal{X} \oplus \mathcal{O}_\mathcal{X}(P), \mathcal{O}_\mathcal{X} \oplus \mathcal{O}_\mathcal{X}(P))$. We recall that $\mathcal{O}_\mathcal{X}(P)$ is the line bundle of functions with poles of order at most one along the smooth codimension one subvariety P. Explicitly, $\mathcal{O}_\mathcal{X}(P)(U) = \mathcal{O}_\mathcal{X}(U)$ and $\mathcal{O}_\mathcal{X}(P)(V) = x/y\mathcal{O}_\mathcal{X}(V)$, as subsets of the function field $K(\mathcal{X})$.

(1) There is a natural multiplication on $\mathcal{B}$ coming from the cup product $\check{C}^p(\mathcal{U},\mathcal{F}) \times \check{C}^q(\mathcal{U},\mathcal{F}) \to \check{C}^{p+q}(\mathcal{U},\mathcal{F})$ and the homomorphism composition map $\mathcal{F} \otimes \mathcal{F} \to \mathcal{F}$.

(2) $\mathbb{G}_m$ acts on $\mathcal{X}$ and hence $k^\times$ acts on $K(\mathcal{X})$. Since the $\mathbb{G}_m$-action preserves U, V and P it follows that $k^\times$ acts on $\mathcal{B}$.

(3) There is a canonical isomorphism $\mathcal{F} = \mathcal{O} \oplus \mathcal{O}(P) \oplus \mathcal{O}(-P) \oplus \mathcal{O}$ ("matrix entries"). The multiplication map $\mathcal{F} \times \mathcal{F} \to \mathcal{F}$ then corresponds to actual multiplication of elements in $K(\mathcal{X})$. Since the cup product is clearly $k^\times$-equivariant, it follows that the multiplication on $\mathcal{B}$ also is. The differential is just the Čech differential, which is also clearly equivariant.

(4) It suffices to look at the matrix entries $\mathcal{O}_\mathcal{X}$ and $\mathcal{O}_\mathcal{X}(\pm P)$ separately. Begin with $\mathcal{O}_\mathcal{X}$. Notice that $\mathcal{O}_\mathcal{X}(U)$ is a free R-module with basis $x^i y^j$ for $i \geq 0$, $j = 0, 1$. Similarly $\mathcal{O}_\mathcal{X}(V)$ is free with basis $u^i v^j$ for $u = 1/y$, $v = x/y$ and $i \geq 0$, $v = 0, 1, 2$. The kernel of $\check{C}^0(\mathcal{U}, \mathcal{O}_\mathcal{X}) \to \check{C}^1(\mathcal{U}, \mathcal{O}_\mathcal{X})$ is given by the free R-module with basis $(1, 1)$, clearly this is equivariant and admits an equivariant complement. $\mathcal{O}_\mathcal{X}(U \cap V)$ is free with basis $x^i y^j$, $i \in \mathbb{Z}$, $j \leq 1$. The image of $\check{C}^0(\mathcal{U}, \mathcal{O}_\mathcal{X}) \to \check{C}^1(\mathcal{U}, \mathcal{O}_\mathcal{X})$ is free with basis $x^i y^j$, $(i, j) \neq (1, -2)$. Again this is clearly equivariant with an obvious equivariant complement. The cases $\mathcal{O}_\mathcal{X}(\pm P)$ can be handled by a similar argument.

(5) $\mathcal{B}_w$ is naturally isomorphic to $\check{C}^\bullet(\{U_w, V_w\}, \mathcal{H}om(\mathcal{O}_{\mathcal{X}_w} \oplus \mathcal{O}_{\mathcal{X}_w}(P_w), \mathcal{O}_{\mathcal{X}_w} \oplus \mathcal{O}_{\mathcal{X}_w}(P_w))$. To see that this is quasi-isomorphic to $\mathcal{Q}_{\mathcal{X}_w,P_w} = \mathrm{Hom}_{D_\infty(\mathcal{X}_w)}(\mathcal{O}_{\mathcal{X}_w} \oplus \mathcal{O}_{\mathcal{X}_w}(P_w), \mathcal{O}_{\mathcal{X}_w} \oplus \mathcal{O}_{\mathcal{X}_w}(P_w))$, it suffices to observe that, for any locally free sheaf $\mathcal{L}$ with an injective resolution $I^\bullet$, we have a diagram of quasi-isomorphisms

$$\begin{array}{ccccc} \check{C}^\bullet\mathcal{H}om(\mathcal{L},\mathcal{L}) & \xrightarrow{a} & \mathrm{Tot}\,\check{C}^\bullet\mathcal{H}om^\bullet(\mathcal{L}, I) & \xleftarrow{b} & \mathrm{Tot}\,\check{C}^\bullet\mathcal{H}om^\bullet(I, I) \\ & & & & \uparrow{\scriptstyle c} \\ & & & & \mathrm{Hom}^\bullet(I, I). \end{array}$$

Indeed, $\mathrm{Tot}\,\check{C}^\bullet$ computes hypercohomology, which factors through quasi-isomorphisms, and the map of complexes of sheaves underlying b is a quasi-isomorphism by proposition

3.1. (For basic facts about hypercohomology, see [23, section 5.7].) For any injective sheaf J and coherent sheaf $\mathcal{F}$, $\mathcal{H}om(\mathcal{F}, J)$ is flasque and hence acyclic, and so the hypercohomology of $\mathcal{H}om^\bullet(I, I)$ is just the cohomology of its global sections, whence c is a quasi-isomorphism. But we already know that $\mathrm{Ext}^*(\mathcal{L}, \mathcal{L}) = H^*(\mathcal{L}^\vee \otimes \mathcal{L}) = H^*(\mathcal{H}om(\mathcal{L}, \mathcal{L})) = \check{C}^*(\mathcal{H}om(\mathcal{L}, \mathcal{L}))$, so a is a quasi-isomorphism as well.

Hence $H(\mathcal{B}) \otimes k(w) \to H(\mathcal{B}_w)$ is an isomorphism (we have computed both sides!).

(6) $H^0(\mathcal{X}, \mathcal{O}_\mathcal{X}) = k$ has weight zero $k^\times$-action, similarly for $\mathcal{O}_\mathcal{X}(\pm P)$. $H^1(\mathcal{X}, \mathcal{O}_\mathcal{X})$ is represented by the Čech cocycle x^2/y which has weight one. Again $\mathcal{O}_\mathcal{X}(\pm P))$ can be handled by a similar argument. □

We recall that (5) allows us to construct a $k^\times$-equivariant minimal model $\mathcal{A}$ for $\mathcal{B}$. We now exploit (6).

Lemma. *1.* $m_d^{w^u} = u^{d-2} m_d^w$ *for* $u \in k^\times$,

2. $m_d^w = 0$ *for* $d = 3, 4, 5$ *or* $d = 7$.

3. Identify $HH^{6,-4}(A)$ *and* $HH^{8,-6}(A)$ *with* k. *Then the map* $w \mapsto ([m_6^w], [m_8^w]) \in HH^{6,-4}(A) \oplus HH^{8,-6}(A) \simeq k^2$ *is given by* $(a, b) \mapsto (Ca, Db)$ *for non-zero constants* C *and* D.

Proof. We recall that $m_d(a_d^u, \dots, a_1^u) = m_d(a_d, \dots, a_1)^u$. The left hand side is the same as $m_d^{(a,b)}([a_d]^u, \dots, [a_1]^u)$, whereas the right hand side is $m_d^{(a,b)^u}([a_d], \dots, [a_1])^u$. Using (6), this implies that $m_d^w([a_d], \dots, [a_1]) u^{|a_d| + \dots + |a_1|} = m_d^{w^u}([a_d], \dots, [a_1]) u^{|a_d| + \dots + |a_1| + 2 - d}$. (The extra factor of u^{2-d} comes about because m_d has degree $2 - d$.) This immediately implies (1).

We now know that m_d^w is a function from W to some vector space (namely $\mathrm{Hom}(Q^{\otimes d}, Q(2-d)))$ which, in our preferred bases, has polynomial entries. Let $f(a, b)$ be one such entry. We know that $f(u^4 a, u^6 b) = u^{d-2} f(a, b)$. Differentiating with respect to u and setting $u = 1$ yields

$$4a\frac{\partial f}{\partial a} + 6b\frac{\partial f}{\partial b} = (d-2)f. \tag{4}$$

This is going to apply to all homogeneous components of f separately. Suppose f is homogeneous of degree r. Then $a\frac{\partial f}{\partial a} + b\frac{\partial f}{\partial b} = rf$, and so we get $b\frac{\partial f}{\partial b} = \alpha f$ for some constant α. This is only possible if f is a pure power (i.e. $f = a^r$ or $f = b^r$). But then equation (4) implies that $4r = d - 2$ (if $f = x^r$) or $6r = d - 2$ (if $f = y^r$). Hence (2).

Finally, everything in (3) is clear, except that C and D might vanish. However, if either of them vanishes, then we obtain at most two isomorphism classes of A_∞-structures on Q in this way, by the arguments of the corollary to proposition 6.2. But non-isomorphic $\mathcal{X}_w$ yield non-isomorphic A_∞-structures, by the argument of theorem 6.4, which allows us to recover the isomorphism class of $\mathcal{X}_w$ from $\mathcal{A}_w$. This is a contradiction, since there are infinitely many isomorphism classes of elliptic curves. □

6.3. The Fukaya category of the torus

We now turn to the symplectic side of the homological mirror symmetry conjecture for elliptic curves. We thus let T denote the standard torus $\mathbb{R}^2/\mathbb{Z}^2$, with its standard symplectic and complex structure. We choose the Maslov cover corresponding to the holomorphic trivialisation of Δ, whence all non-contractible embedded curves and all symplectomorphisms can be graded (refer to the discussion in section 4.3.2). We let throughout Fuk denote the Fukaya category of the torus with respect to these choices.

Generation

We first need to show that Fuk is generated by two objects, corresponding to $\mathcal{O}$ and $\mathcal{O}/P$. Let A and B be two meridians (so they in particular freely generate $H_1(T, \mathbb{Z})$). For definiteness, we assume that the basis A, B at their intersection is positively oriented – even more concretely, we can just assume that A corresponds to the line $(t, 0)$ and B to the line $(0, t)$. We claim these serve as generators. We aim to use theorem 2.7 on split-generation criteria using twists, by showing that for $X \in \text{Fuk}$, there exists $n > 0$ such that $\text{Hom}_{H^0(\text{Fuk})}(X, (T_A T_B)^n(X)) = 0$. This is clearly sufficient.

Definition. *Let $(M, \omega, \mathcal{L})$ be a graded manifold and ϕ a graded automorphism. Choose a compatible almost-complex structure I and realise $\mathcal{L}$ as $\mathcal{L}_{\Omega^2}$. We say that ϕ is* of constant sign *(with respect to these choices) if the lift $\tilde{s}_\phi : LGr(M) \to \mathbb{R}$ (compare section 4.3.1) is never zero.*

Proposition 6.5 ([19], Lemma 9.2)**.** *Let (M, ω, I, Ω^2) be a graded manifold as above and ϕ a graded automorphism such that ϕ^d is of constant sign for some $d > 0$. Assume additionally that M is compact and connected. Let $X_1, \ldots, X_n \in \text{Fuk}(M)$ be Lagrangian spheres such that $\tau_{X_1} \ldots \tau_{X_n}$ is Hamiltonian isotopic to ϕ.*

Then $X_1, \ldots, X_n$ split-generate $\text{Fuk}(M)$.

Proof. We know that the Dehn twists τ_{X_i} correspond to the twist functors T_{X_i}, and that ϕ and $\tau_{X_1} \ldots \tau_{X_n}$ act quasi-isomorphically on $\text{Fuk}(M)$ (by standard properties of the Fukaya category, i.e. theorem 4.2).

It hence suffices to show, by theorem 2.7, that for $X \in \text{Fuk}(M)$ and $n > 0$ sufficiently large, $\text{Hom}_{H^0(\text{Fuk}(M))}(X, \phi^n(X)) = 0$. We may thus replace ϕ by ϕ^d and so assume that ϕ is of constant sign.

Since both $LGr(2n)$ and M are compact and connected (the former being a quotient of $U(n)$), it follows that $LGr(M)$ is compact and connected. In particular $\tilde{s}_\phi$ is either always positive or always negative. Let us assume that $\tilde{s}_\phi$ is always positive (the argument being essentially the same if it is always negative). Then by compactness, there exists $\epsilon > 0$ such that $\tilde{s}_\phi > \epsilon$. Then $\tilde{s}_{\phi^n} > n\epsilon$.

Let $\tilde{s}_X$ denote the grading of X. By compactness of X, this function is bounded, say $|\tilde{s}_X| < r$. It follows from the formula $\tilde{s}_{\phi(L)} = \tilde{s}_L \circ \phi^{-1} + \tilde{s}_\phi \circ s_L \circ \phi^{-1}$ that for n suitably large, $\tilde{s}_{\phi^n(X)} > r + 2$. Then the same inequality with $r+1$ instead of $r+2$ will hold for a sufficiently small Hamiltonian perturbation of $\phi^n(X)$, which we may assume transverse to X. It follows that all intersection points have Maslov index ≥ 1. This follows from proposition 4.1 (which shows how to compute Maslov indices using lifts) for the 2D case, which is all that we will use in the sequel. For the general case, see the reference.

This concludes the proof. □

It remains to apply this to our case. A symplectomorphism ϕ of T induces an action on $\pi_1(T) = \mathbb{Z}^2$, i.e. an element M_ϕ of $SL_2(\mathbb{Z})$. Then M_ϕ itself acts on T, and it is well known that ϕ and M_ϕ are homotopic. In the case of Dehn twists, $M_{\tau_A} = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}$ and $M_{\tau_B} = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. In this case the symplectomorphisms and their linearisations are not only homotopic, but Hamiltonian isotopic. One way to see this is to observe that M_{τ_A} is actually a model Dehn twist corresponding to a piecewise linear but non-smooth choice of f in the notation of section 4.2.1, implanted via a Weinstein neighbourhood which actually wraps once around the whole torus. Nonetheless it is easy to see that the arguments from section 4.2.1 generalize to show our claim.

It follows that $\tau_A\tau_B$ is Hamiltonian isotopic to $M_{\tau_A}M_{\tau_B} = \begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}$. This rotates the first basis vector by $\pi/2$ and the second by $\pi/4$. From this one easily sees it is of constant sign (since it is linear). Its value varies by $1/4$ without changing sign (for some choice of grading), and hence *any* grading will be of constant sign.

Counting polygons

As explained at the beginning of this sections, for technical reasons (i.e. having studied the derived category of E in terms of $\mathcal{O}$ and $\mathcal{O}(P)$ instead of $\mathcal{O}$ and $\mathcal{O}/P$), we should really be using the generating set A and $B' = \tau_B(A)$. However, as also explained there, this will yield the same relations algebra (up to quasi-isomorphism) as using the more standard generators A and B, so we will do this here.

We now determine $H\mathcal{Q} = \mathrm{Hom}_{HD^\pi\mathrm{Fuk}}(A \oplus B, A \oplus B)$ geometrically, using the methods explained in section 4.4. This is straightforward, after using bilinearity of Hom to break this up into $\mathrm{Hom}(A, A)$, $\mathrm{Hom}(A, B)$ etc. Indeed, we can choose the evident gradings for A and B, corresponding to constant functions with values in $[0, 1)$. Even more explicitly, we may put $\tilde{s}_A \equiv 0$ and $\tilde{s}_B \equiv \frac{1}{2}$. This induces orientations for A and B. We will also need to consider small Hamiltonian perturbations A' and B'.

We have assembled sketches of the relevant situations in figure 6.3. We begin with determining $\mathrm{Hom}(A, B)$ and $\mathrm{Hom}(B, A)$, so consider subfigure 2(a). Since A and B have a unique intersection point, it follows immediately that both spaces are one-dimensional, and we need only determine the grading. This is given by the Maslov indices $\mu_p(A, B)$ and $\mu_p(B, A)$ respectively. We have, by theorem 4.1 determining the Maslov index in terms of the lifts, that $\mu_p(A, B) = -[\tilde{s}_B(p) - \tilde{s}_A(p)] = -\left[\frac{1}{2} - 0\right] = 0$, and that $\mu_p(B, A) = -\left[0 - \frac{1}{2}\right] = 1$.

We next consider $\mathrm{Hom}(A, A)$. We already know by theorem 4.2 that this must have cohomology algebra isomorphic to the singular cohomology $H^*(A, \Lambda_t)$, which is $\Lambda_t \oplus \Lambda_t(-1)$ (since A is a circle). It seems enlightening to compute this, by perturbing A

slightly to yield a Hamiltonian isotopic curve A'. One such perturbation is shown in subfigure 2(b). We see that $\mathrm{Hom}(A, A')$ is two-dimensional, generated by p and q. Also note that $\tilde{s}_{A'}(p) = -\epsilon < 0$, and so $\mu_p(A, A') = -[-\epsilon - 0] = 1$. In contrast, $\tilde{s}_{A'}(q) = \epsilon > 0$, and so $\mu_q(A, A') = 0$. It follows that $\mathrm{Hom}(A, A')$ might carry a non-zero differential $d(q) = \lambda p$. From the explanations in section 4.4 describing the Fukaya category, we determine the differential by looking for immersed bigons with sides A and A', up to equivalence. For reference, the form of such bigons is sketched in the figure. One sees that there are precisely two such bigons, labelled F and G in the figure. We thus find $dq = (t^F \pm t^G)p$. Here the sign depends on conventions which we have not explained. But observe that we must have $\mathrm{Hom}_{H\mathrm{Fuk}}(A, A') \neq 0$, since one element of this set must represent $\mathbf{1}_A$. This can happen only if $dq = 0$, and then indeed $\mathrm{Hom}_{H\mathrm{Fuk}}(A, A') \simeq H^*(A, \Lambda_t)$, as expected. Note also that $dq = 0$ is indeed possible: the perturbation is Hamiltonian, so F and G have the same area, and so $dq = 0$ if the negative sign is chosen above.

A similar argument applies for $\mathrm{Hom}(B, B)$. We thus find that, as a vector space, $H(\mathcal{Q}) \simeq Q$ (where Q is the algebra from section 6.1 again). As when we investigated the ext algebra of an elliptic curve, again almost all compositions are determined formally, and we can show that $H(\mathcal{Q}) \simeq Q$ *as algebras* provided that we show $\mathrm{Hom}_{H\mathrm{Fuk}}(A, B)\,\mathrm{Hom}_{H\mathrm{Fuk}}(B, A) \neq 0$ and $\mathrm{Hom}_{H\mathrm{Fuk}}(B, A)\,\mathrm{Hom}_{H\mathrm{Fuk}}(A, B) \neq 0$.

That is to say, we need to determine some m_2, i.e. count some triangles. Consider now subfigure 2(c). We will show that $\mathrm{Hom}_{H\mathrm{Fuk}}(B, A') \otimes \mathrm{Hom}_{H\mathrm{Fuk}}(A, B) \to \mathrm{Hom}_{H\mathrm{Fuk}}(A, A') \simeq \mathrm{Hom}_{H\mathrm{Fuk}}(A, A)$ is not the zero map. In the figure, we have the two meridians A and B, and the perturbation A' of A. For illustrative purposes we have drawn two copies of the square representing the torus, which we think of as two fundamental domains of the lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$. On T, as before A and B have a unique intersection point, denoted q. The point q' is equivalent to it. Similarly A' and B intersect in a unique point p, and there is another equivalent intersection point p' drawn. Also A' and A intersect in r and s. As before, we determine $|q| = 0$, $|p| = 1$, $|r| = 1$ and $|s| = 0$. We suspect $m_2(p, q) = \lambda r$, for some $\lambda \neq 0$. (Indeed r is the only intersection point of the right degree of B and A'.) As explained in section 4.4, to compute this product, we need to count triangles. The type of triangle we are looking for is sketched in the figure: triangles with vertices (in clockwise order) q, p and x, connected along segments of curves as indicated. Also in this figure we shaded one such triangle.

The main observation is now that this is the *only* such triangle, up to equivalence. Namely it is the only such triangle with one vertex given by precisely q (and not an equivalent point). Indeed, the side $q \to p$ of such a triangle is the *only* side along which the y coordinate can change by more than one, and even along the other two sides together it can never change by one (since the other sides lie on A or A'). But going once round the whole triangle must leave the y coordinate unchanged, so if it changes by more than one once it must do so twice, which is impossible. It follows that the second vertex of the triangle must lie in the same fundamental domain as we started, i.e. it must be p. The rest is now easy, since the map of the triangle into T must be orientation-preserving and an embedding away from the corners. Hence the second side of the triangle must "start out to the left" (orientation), and the third vertex must be r

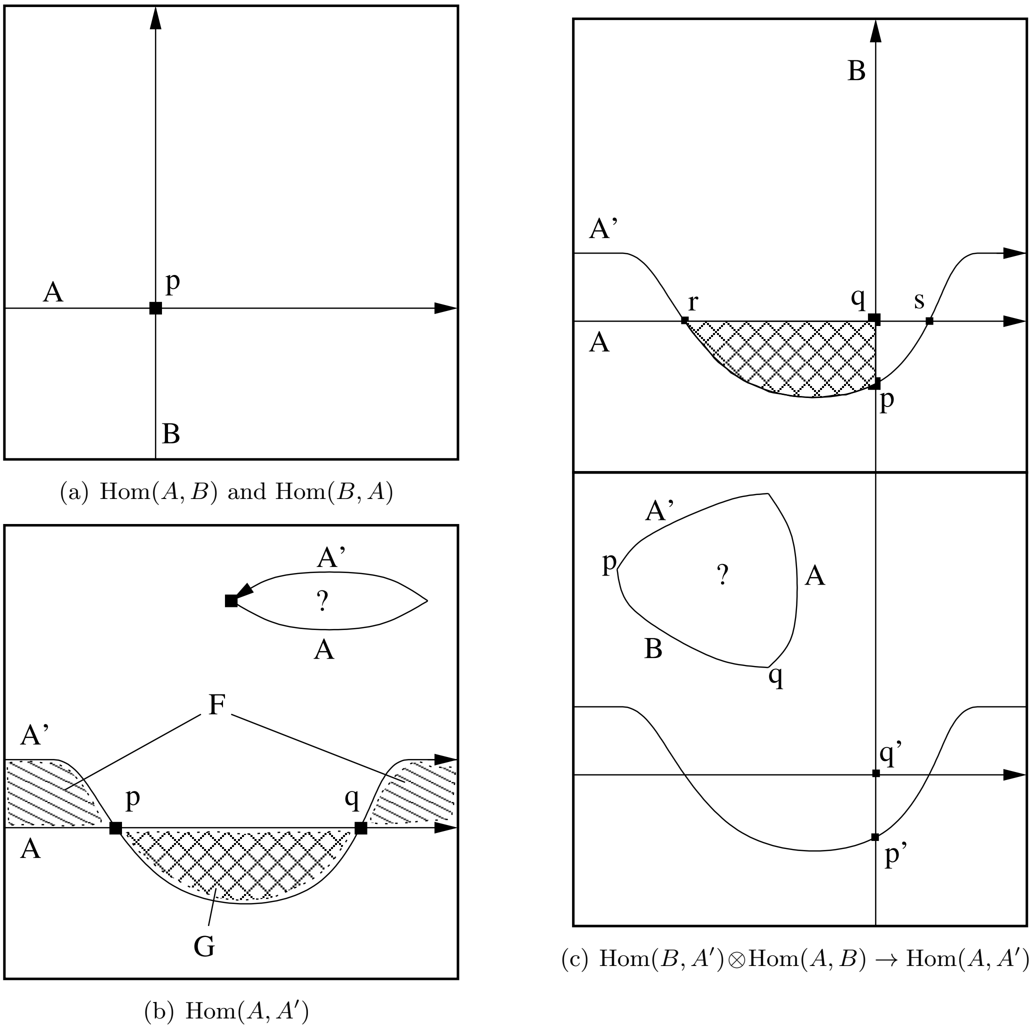

(a) $\mathrm{Hom}(A, B)$ and $\mathrm{Hom}(B, A)$

(b) $\mathrm{Hom}(A, A')$

(c) $\mathrm{Hom}(B, A') \otimes \mathrm{Hom}(A, B) \to \mathrm{Hom}(A, A')$

Figure 2: Computing Floer chain complexes.

(and not one of its translates), since otherwise we cannot get an embedding.

The other composition is treated entirely similarly, and so the proof of abstract mirror symmetry is complete.

7. Illustrations of mirror symmetry for elliptic curves

In this section we intend to reap in the benefits of all the hard work that went into proving theorems 6.1 and 6.4. Recall that we have proved there that there exists a quasi-equivalence $\psi : D^\pi \mathrm{Fuk}(T) \to D_\infty(E)$ for some cubic curve E. It is determined by $\psi(A) = \mathcal{O}$ and $\psi(B) = \mathcal{O}/P$, where A and B are the standard meridians.

However, we do not know what curve E is. We do not even know if it is smooth. We remedy this problem in subsection 7.1. Indeed theorem 6.4 also includes a recipe on how to determine E: start with the two meridians A and B corresponding to $\mathcal{O}$ and $\mathcal{O}/P$. Consider the ring $S = \bigoplus_{n\geq 0} \mathrm{Hom}_{H^0 D^\pi \mathrm{Fuk}(T)}(A, T_B^{3n}(A))$. This is the projective coordinate ring of E. Using some explicit polygon counts, we can find a defining equation for E, check it is smooth, and compute its j-invariant.

Once we have done that, we investigate in somewhat more detail the mirror correspondence ψ. In subsection 7.2, we define the rank and slope of a vector bundle, and explain how the slope m/n of a geodesic circle C in T corresponds to the rank and slope of the vector bundle it represents in $D_\infty(E)$ under the mirror correspondence.

Observe that the group $SL_2(\mathbb{Z})$ acts on T, and hence it (or rather its central extension) acts on $\mathrm{Fuk}(T)$. By the mirror correspondence, there must be a corresponding action on $D_\infty(E)$. In subsection 7.3 we construct this action, and use it to prove the sketched correspondence of rank and slope from the previous subsection.

7.1. Determination of the mirror curve

We now carry out the recipe sketched above. To do this, we work in the universal cover $\mathbb{R}^2$ of T. We note that $\rho(x,y) = (x, y+3x)$ is a model for τ_B^3. We put $L_i = \rho^i(A)$.

We must find (three) generators of $S_1 = \mathrm{Hom}(A, L_1)$ and a cubic relation they satisfy in $S_3 = \mathrm{Hom}(A, L_3)$. For this, we put $X_i = (i/3, 0)$ (for $i = 0,1,2$), $Y_i = (i/6, 0)$ (for $i = 0,\ldots,5$) and $Z_i = (i/9, 0)$ (for $i = 0, \ldots, 8$). Here and elsewhere, indices will be understood " (mod n)" for appropriate n (e.g. $n = 3$ for the X_i). We note that $\mathrm{Hom}(A, L_i)$ is a graded vector space concentrated in degree zero, of dimension $3i$. Indeed $X_i \in \mathrm{Hom}(A, L_1)$, $Y_i \in \mathrm{Hom}(A, L_2)$ and $Z_i \in \mathrm{Hom}(A, L_3)$ are the natural generators of these vector spaces (i.e. intersection points).

What we need to do, then, is to write ten cubic monomials in the X_i in terms of the nine basis elements Z_j of $\mathrm{Hom}(A, L_2)$. Here we have to recall that the ring structure on S is produced "via ρ". So for example by $X_1 X_2$ we mean $m_2(\rho(X_1), X_2) \in \mathrm{Hom}(A, L_2)$, where we consider $\rho(X_1) \in \mathrm{Hom}(\rho(A), \rho(L_1)) = \mathrm{Hom}(L_1, L_2)$.

So what we really want to do is to first write the six quadratic monomials in the X_i in terms of the Y_j, and then work out the cubic monomials from there. Here we are using the fact that we *know* that S must be commutative and associative.

Let us consider an example. Suppose we want to express X_0^2 in terms of the Y_j. We are thus looking for triangles with vertices X_0, $\rho(X_0) = X_0$ and some Y_j. Consider figure 7.1. This shows the universal cover, in which all points and curves (lines) are replicated infinitely many times. Shaded are two triangles F and G of the form we are looking for. It is easy to see that all other triangles contributing to X_0^2 are translates and scalings of these two triangles. It follows that $X_0^2 = A_0Y_0 + A_3Y_3$. We determine A_3. This means we have to count scalings of F. One easily sees that these are indexed by integers $n \in \mathbb{Z}$, and have vertices $(0,0)$, $(2n+1, 6n+3)$ and $(0, n+1/2)$. The triangle F shown corresponds to $n = 0$. The triangle corresponding to n has area $1/2.(n+1/2).(6n+3) = 3(n+1/2)^2$. Hence

$$A_3 = \sum_{n\in\mathbb{Z}} t^{3\left(n+\frac{1}{2}\right)^2}.$$

This is a classical theta series. It turns out that all the triangle counts we have to do yield these. Zaslow [24] has carried out the counts in detail. His results can be summarised as follows. Write

$$A_i = \sum_{n\in\mathbb{Z}} t^{3\left(n+\frac{i}{6}\right)^2}$$
$$B_i = \sum_{n\in\mathbb{Z}} t^{9\left(n+\frac{i}{18}\right)^2}.$$

Then

$$X_iX_j = A_{i-j}Y_{i+j} + A_{i-j+3}Y_{i+j+3}$$
$$Y_iX_j = B_{2j-i}Z_{i+j} + B_{2j-i+6}Z_{i+j+3} + B_{2j-i+12}Z_{i+j+6}.$$

From this, working out the linear relation among the cubic monomials is entirely routine, if rather cumbersome. From it the j-invariant of the mirror can be computed. The result is as follows.

Proposition 7.1 (Zaslow). *Put $u = A_2B_0 + A_1B_9$, $p = A_0B_0 + A_3B_9$, $q = A_0B_6 + A_3B_3$ and $z = \frac{2q+p}{3u}$. Then the cubic relation satisfied by the X_i is*

$$X_0^3 + X_1^3 + X_2^3 - 3zX_0X_1X_2 = 0.$$

This is a smooth cubic, with j-invariant given by

$$j_{\mathrm{Fuk}} = -27\frac{z^3(z^3+8)^3}{(1-z^3)^3}.$$

This answers completely the question posed for this subsection. Moreover, Zaslow observed that, in the way we have presented mirror symmetry, a small miracle happens. Consider, for τ in the upper half plane, the complex torus $E_\tau = \mathbb{C}/\mathbb{Z} \oplus \tau\mathbb{Z}$. This is an elliptic curve, and its j-invariant is a classical function. It is usually written in terms of $q = e^{2\pi i\tau}$, because it then has a Laurent series expansion with integral coefficients [13]

$$j(q) = \frac{1}{q} + 744 + 196884q + \ldots$$

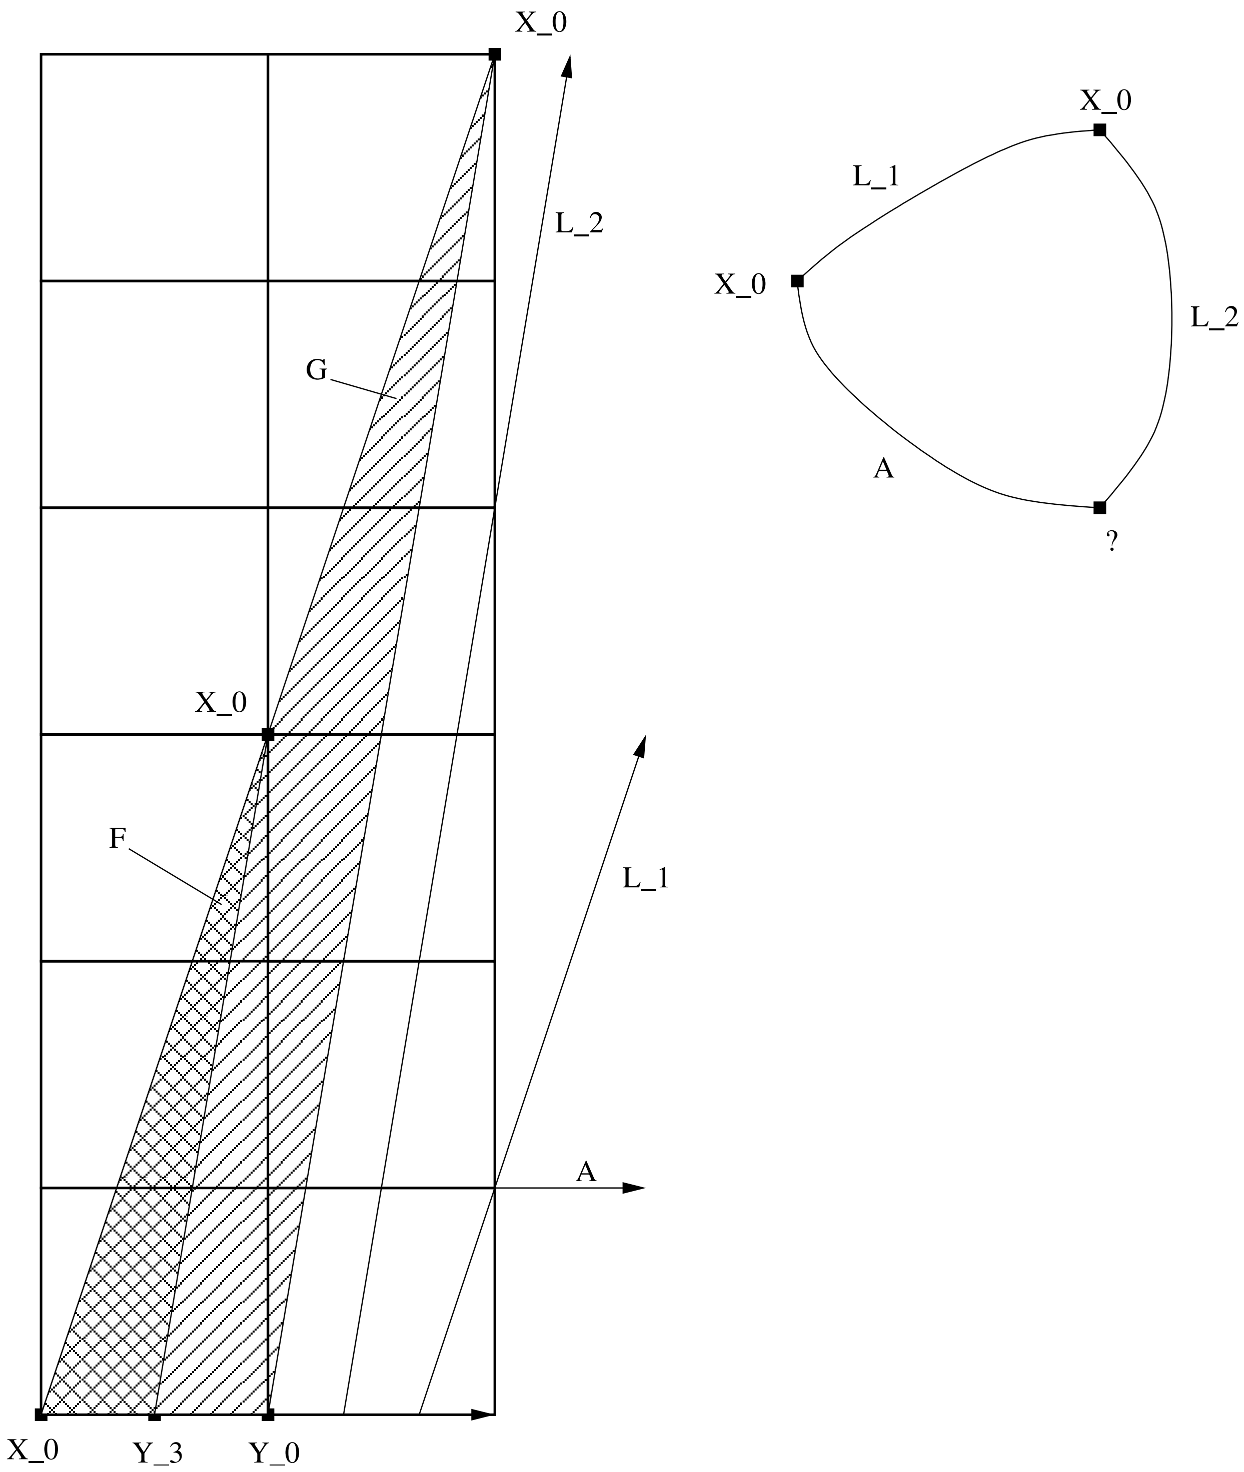

Figure 3: Some of the triangles contributing to X_0^2.

Theorem 7.2. *The mirror curve to T is the smooth plane cubic over Λ_t with j-invariant j given by the Laurent expansion of the classical $j(t)$-function.*

In principle, this can be proven by showing that $j_{\mathrm{Fuk}}(t)$ is a modular function and then computing the first few terms of the series expansion (since spaces of modular functions are finite-dimensional). However, this is fairly non-trivial. We instead contend ourselves with verifying the claim order-by-order, for as many terms as one wishes (which is of course not a proof, but still quite convincing): the python script `series.py` in figure 4 in the appendix can be used to compute the series expansion of $j_{\mathrm{Fuk}}(t)$, to any order desired. It proceeds by brute force, i.e. computing the series of the theta functions appearing in $j_{\mathrm{Fuk}}(t)$ to sufficiently high order, and then inverting them and multiplying them together. We have used it to verify the claim up to $O(q^4)$.

Mirror symmetry over $\mathbb{C}$ An interesting observation is that, given a complex number q in the upper half plane, one may form an A_∞-category $\mathrm{Fuk}_q(T)$, which is obtained from $\mathrm{Fuk}(T)$ by "substituting" q for t. That is to say, the appropriate series in the definition actually converge. The above results then imply that if E is an elliptic curve over $\mathbb{C}$ with j-invariant $j(q)$ (for some particular $q \in \mathbb{C}$), then $D_\infty(X)$ is quasi-equivalent to $D^\pi \mathrm{Fuk}_q(T)$. Hence, over $\mathbb{C}$, we have found mirrors for all elliptic curves.

7.2. Pictorial mirror symmetry: rank and slope of vector bundles

Our proof of mirror symmetry is rather abstract. In particular, determining the image $\psi(C)$ of an embedded circle $C \in \mathrm{Fuk}(T)$ under the mirror correspondence is not at all trivial. Here we intend to describe at least what kind of object $\psi(C)$ is. In fact, one can show that it is always (a shift of) a locally free sheaf or a skyscraper sheaf, and so one very natural question is what rank and first Chern class $\psi(C)$ has. This is the problem we tackle in this subsection.

Suppose $\mathcal{A}$ is a triangulated A_∞-category and B is an abelian group. A map $e : Ob(\mathcal{A}) \to B$ is called an *Euler characteristic* (or *additive*) if for any morphism $f : X \to Y$ in $\mathcal{A}$, we have $0 = e(X) - e(Y) + e(cone(f))$. It is clear that e is determined by its value on any set of generators (not split-generators) of $\mathcal{A}$. For example, if $\mathcal{A} = D_\infty(X)$, for X a smooth projective variety, then e is determined by its values on vector bundles (since these are generators, by theorem 3.2).

Suppose $\mathcal{A}$ is a triangulated A_∞-category over k and $X \in A$ is of finite cohomological dimension, by which we mean that for all $Y \in \mathcal{A}$ there exists n such that $H^r \operatorname{Hom}_{\mathcal{A}}(X, Y) = 0$ for all $|r| > n$, and also $\dim_k H^r \operatorname{Hom}_{\mathcal{A}}(X, Y) < \infty$ for all r. In this case $e(Y) = \sum_i (-1)^i \dim_k H^i \operatorname{Hom}_{\mathcal{A}}(X, Y) =: \chi(H^* \operatorname{Hom}(X, Y))$ is an Euler characteristic on $\mathcal{A}$. This follows from the fact that $H^0\mathcal{A}$ is a triangulated category. A similar result holds for $\operatorname{Hom}(\bullet, X)$.

Suppose now X is a smooth projective variety and $P \in X$ is a closed point. We define two Euler chararacteristics on $D_\infty(X)$, called *rank* and *slope*, by $rk(\mathcal{F}) = \chi(\mathrm{Ext}^*(\mathcal{F}, \mathcal{O}_X/P))$ and $sl(\mathcal{F}) = \chi(\mathrm{Ext}^*(\mathcal{O}, \mathcal{F})) = \chi(H^*(X, \mathcal{F}))$. It is easy to see that for locally free sheaves, this definition of rank agrees with the conventional one, and so is in particular independent of the choice of P. Also if X is an elliptic curve and $\mathcal{L}$ is a line bundle, then, by the Riemann-Roch theorem and Serre duality, the slope of $\mathcal{L}$ equals its degree (or first Chern class).[2]

Of course, we can combine these two Euler characteristics and form a map $k : D_\infty(E) \to \mathbb{Z}^2, \mathcal{F} \mapsto (sl(\mathcal{F}), rk(\mathcal{F}))$. This corresponds under mirror symmetry to a map $k' : \mathrm{Fuk}(T) \to \mathbb{Z}^2$, defined by the same formulas.

There is another geometrically inspired similar-looking map from $\mathrm{Fuk}(T)$. Namely, recall that objects of $\mathrm{Fuk}(T)$ (not of $D^\pi\mathrm{Fuk}(T)$) are represented by embedded circles with extra data. We thus have a map $h : \mathrm{Fuk}(T) \to H_2(T) \simeq \mathbb{Z}^2$. Explicitly, if C is a curve represented in the universal cover of T by a line of slope m/n (and oriented "to the right"), where the fraction is in lowest terms and $n > 0$, then $h(C) = (m, n)$. This explains the term "slope".

We claim that $h = k'$, where defined. This is not particularly hard to check by hand. We verify here that $h(C_x) = k'(C_x)$, where C_x is the curve corresponding to the line $\{(t, x)|t \in [0, 1]\}$ parallel to A. The general case will be established in the next subsection.

The claim is clear if $x = 0$, since $C_0 = A$, and $\psi(A) = \mathcal{O}$ has rank one and slope (i.e. degree) zero Moreover, if $x \neq 0$, then $\mathrm{Hom}_{\mathrm{Fuk}(T)}(A, C_x) = 0$ (since there are no intersection points) and so $sl'(C_x) = \chi(H^* \mathrm{Hom}_{\mathrm{Fuk}(T)}(A, C_x) = 0$. On the other hand $\mathrm{Hom}_{\mathrm{Fuk}(T)}(C_x, B)$ is one-dimensional, generated in degree zero, by precisely the same picture as for $t = 0$, and so $rk'(C_x) = 1$.

7.3. The braid action

The group $SL_2(\mathbb{Z})$ acts on T in a natural way. It clearly preserves the symplectic form (since determinant one matrices preserve volume), and so by the discussion in subsection 4.3.2, all of its elements are gradable. In fact, it is well known that $SL_2(\mathbb{Z})$ is generated by the standard linear Dehn twists M_A and M_B (compare section 6.3) with respect to the two relations $(M_A M_B)^6 = \mathbf{1}$ and $M_A M_B M_A = M_B M_A M_B$. Denote by $s_1 = \tilde{\tau}_A$ and $s_2 = \tilde{\tau}_B$ the standard gradings of the Dehn twists. It follows that $(s_1 s_2)^6 = \mathbf{1}[k_1]$ and $s_1 s_2 s_1 = s_2 s_1 s_2[k_2]$. It is easy to check that M_A, M_B and $M_A M_B$ all rotate their arguments by small positive amounts. From this observation it is easy to see that $k_1 \neq 0$ and $k_2 = 0$. (One may in fact easily check that $k_1 = 2$, but we do not really need this.) It follows that the subgroup of $Sym^{gr}(T)$ generated by s_1 and s_2 is isomorphic to the braid group

$$B_3 = \langle s_1, s_2 | s_1 s_2 s_1 = s_2 s_1 s_2 \rangle.$$

We have hence found a *braid action* on $\mathrm{Fuk}(T)$. Via the mirror correspondence ψ, this yields a braid action on $D_\infty(E)$ as well (recall that quasi-equivalences of A_∞-categories admit quasi-inverses).

We can construct this action on $D_\infty(E)$ directly. Indeed, by theorem 4.2 summarizing the properties of the Fukaya category, s_1 acts by T_A on $\mathrm{Fuk}(T)$, and similarly s_2 acts by T_B. It then follows from theorem 2.5 (which states that twist functors are preserved

[2] One may show that the first Chern class c_1 is also an Euler characteristic. If X is not an elliptic curve, this is probably a better definition for the slope than ours. Since we only deal with elliptic curves, we will ignore this problem.

by A_∞-functors) that under ψ, the action of s_1 is via $T_{\psi(A)} = T_{\mathcal{O}}$ and the action of s_2 is via $T_{\mathcal{O}/P}$. This by itself is already quite an interesting statement, since it is a priori not at all clear that these twists are invertible (although this follows from the theory of "spherical objects").

Comparison of h and k

We now exploit this braid action to show that $h(C) = k(\psi(C))$ for and $C \in \mathrm{Fuk}(T)$.

Lemma. *For $C \in \mathrm{Fuk}(T)$, $\mathcal{F} \in D_\infty(E)$ and $s \in B_3$, we have*

$$h(sC) = sh(C)$$
$$k(s\mathcal{F}) = sk(\mathcal{F}),$$

where the action on $\mathbb{Z}^2$ is via the natural quotient $SL_2(\mathbb{Z})$.

Proof. It suffices to prove this for $s = s_1$ and $s = s_2$ (since these generate B_3), which act by M_{τ_A} and M_{τ_B} on $\mathbb{Z}^2$ respectively. On the symplectic side, they correspond respectively to $\tilde{\tau}_A$ and $\tilde{\tau}_B$, whereas on the algebraic side they correspond to $T_{\mathcal{O}}$ and $T_{\mathcal{O}/P}$. We thus need to show $h(\tilde{\tau}_A C) = M_{\tau_A} h(C)$, for $C \in \mathrm{Fuk}$, and so on.

This is obvious on the symplectic side, i.e. for h. On the algebraic side, we just observe the exact triangles (from the definition of twist functors)

$$\mathcal{O} \otimes \mathrm{Ext}^*(\mathcal{O}, \mathcal{F}) \to \mathcal{F} \to T_{\mathcal{O}}(\mathcal{F}) \to$$
$$\mathcal{O}/P \otimes \mathrm{Ext}^*(\mathcal{O}/P, \mathcal{F}) \to \mathcal{F} \to T_{\mathcal{O}/P}(\mathcal{F}) \to .$$

Since sl and rk are Euler characteristics, it follows that

$$\begin{aligned}
sl(T_{\mathcal{O}}(\mathcal{F})) &= sl(\mathcal{F}) - sl(\mathcal{O} \otimes \mathrm{Ext}^*(\mathcal{O}, \mathcal{F})) = sl(\mathcal{F}) \\
rk(T_{\mathcal{O}}(\mathcal{F})) &= rk(\mathcal{F}) - rk(\mathcal{O} \otimes \mathrm{Ext}^*(\mathcal{O}, \mathcal{F})) = rk(\mathcal{F}) - \chi(\mathrm{Ext}^*(\mathcal{O}, \mathcal{F})) = rk(\mathcal{F}) - sl(\mathcal{F}) \\
sl(T_{\mathcal{O}/P}(\mathcal{F})) &= sl(\mathcal{F}) - sl(\mathcal{O}/P \otimes \mathrm{Ext}^*(\mathcal{O}/P, \mathcal{F})) = sl(\mathcal{F}) - \chi(\mathrm{Ext}^*(\mathcal{O}/P, \mathcal{F})) \\
rk(T_{\mathcal{O}/P}(\mathcal{F})) &= rk(\mathcal{F}) - rk(\mathcal{O}/P \otimes \mathrm{Ext}^*(\mathcal{O}/P, \mathcal{F})) = rk(\mathcal{F}).
\end{aligned}$$

Here we have used that Euler characteristics are additive, change sign under shifts, and that $k(\mathcal{O}) = (0, 1)$, $k(\mathcal{O}/P) = (1, 0)$.

It remains to show that $\chi(\mathrm{Ext}^*(\mathcal{O}/P, \mathcal{F})) = -rk(\mathcal{F})$, and as usual it suffices to show this for vector bundles. One simple way to see this is to consider the exact triangle $\mathcal{O} \to \mathcal{O}(P) \to \mathcal{O}/P$. If we write $e(\mathcal{F}, \mathcal{G}) = \chi(\mathrm{Ext}^*(\mathcal{F}, \mathcal{G}))$, then we obtain

$$\begin{aligned}
e(\mathcal{O}/P, \mathcal{F}) &= e(\mathcal{O}(P), \mathcal{F}) - e(\mathcal{O}, \mathcal{F}) \\
&= e(\mathcal{F}^\vee, \mathcal{O}(-P)) - e(\mathcal{F}^\vee, \mathcal{O}) \\
&= -e(\mathcal{F}^\vee, \mathcal{O}/P)
\end{aligned}$$

(where in the second line we have used that $\mathrm{Ext}^*(\mathcal{F} \otimes \mathcal{L}, \mathcal{G}) = \mathrm{Ext}^*(\mathcal{F}, \mathcal{G} \otimes \mathcal{L}^\vee)$ twice). But this last expression is just $-rk(\mathcal{F}^\vee) = -rk(\mathcal{F})$. This concludes the proof. $\square$

Theorem 7.3. *For any* $C \in \mathrm{Fuk}(T)$, *we have* $h(C) = k(\psi(C))$.

Proof. It is well known that $SL_2(\mathbb{Z})$ acts transitively on subsets of $\mathbb{Z}^2$ of the form $Z_r = \{(m,n) |\, \gcd(m,n) = r\}$ (this is essentially a restatement of Euclid's algorithm). Hence it follows from the lemma that we may restrict to curves C with $h(C) = (r, 0)$. Recalling that contractible curves are not gradable (compare section 4.3.2), we know $r > 0$. We will in fact show that $r = 1$ is the only possibility.

An argument using the theory of mean curvature flow [6], which we need to omit for space reasons, shows that if $r = 1$, any such curve C is Hamiltonian isotopic to a geodesic C'. The geodesics of T correspond to straight lines (of rational slope) in the universal cover, and we have described the effect of h on such curves before. It follows that $C' = C_x$, for some x, in the notation of subsection 7.2. But this is the case we already verified there.

If $r > 1$, then we may still apply mean curvature flow to find a Hamiltonian isotopic curve C' of arbitrarily small curvature. But such a curve (with small curvature) must intersect itself. Since mean curvature flow does not create self-intersections, we conclude that C already had self-intersections, which is not allowed (we require our curves to be embedded). □

References

[1] Mohammed Abouzaid. On the fukaya categories of higher genus surfaces. *Advances in Mathematics*, 217(3):1192 – 1235, 2008.

[2] Sylvain E. Cappell, Ronnie Lee, and Edward Y. Miller. On the maslov index. *Communications on Pure and Applied Mathematics*, 47(2):121–186, 1994.

[3] David A Cox and Sheldon Katz. *Mirror symmetry and algebraic geometry.* Math. Surveys Mono. AMS, Rhode Island, NJ, 1999.

[4] Ana Cannas da Silva. Symplectic geometry. In F.J.E. Dillen and L.C.A. Verstraelen, editors, *Handbook of Differential Geometry.* Elsevier, 2005.

[5] Alice Fialowski and Michael Penkava. Deformation theory of infinity algebras. *Journal of Algebra*, 255(1):59 – 88, 2002.

[6] Matthew A. Grayson. Shortening embedded curves. *Annals of Mathematics*, 129(1):pp. 71–111, 1989.

[7] Dieter Happel. Hochschild cohomology of finite—dimensional algebras. In Marie-Paule Malliavin, editor, *Séminaire d'Algèbre Paul Dubreil et Marie-Paul Malliavin*, volume 1404 of *Lecture Notes in Mathematics*, pages 108–126. Springer Berlin Heidelberg, 1989.

[8] Robin Hartshorne. *Algebraic geometry.* Springer-Verlag, New York, 1977. Graduate Texts in Mathematics, No. 52.

[9] Bernhard Keller. Introduction to A-infinity algebras and modules., 2001.

[10] Maxim Kontsevich. Homological algebra of mirror symmetry. In *Proceedings of the International Congress of Mathematicians*, pages 120–139. Birkhäuser, 1994.

[11] Yanki Lekili and Timothy Perutz. Arithmetic mirror symmetry for the 2-torus, 2012. Preprint.

[12] Yankı Lekili and Timothy Perutz. Fukaya categories of the torus and dehn surgery. *Proceedings of the National Academy of Sciences*, 108(20):8106–8113, 2011.

[13] James S. Milne. *Modular Functions and Modular Forms.* 2009. Available at www.jmilne.org/math/.

[14] Alexander Polishchuk. Massey and fukaya products on elliptic curves, 1998.

[15] Alexander Polishchuk. Homological mirror symmetry with higher products, 1999.

[16] Alexander Polishchuk. A-infinity structures on an elliptic curve. Commun.Math.Phys. 247 (2004) 527-551, 2004.

[17] Alexander Polishchuk and Eric Zaslow. Categorical mirror symmetry: The elliptic curve. Adv.Theor.Math.Phys.2:443-470,1998, 1998.

[18] Paul Seidel. Graded lagrangian submanifolds. *Bulletin de la Société Mathématique de France*, 128(1):103–149, 2000.

[19] Paul Seidel. Homological mirror symmetry for the quartic surface, 2003.

[20] Paul Seidel. *Fukaya categories and Picard-Lefschetz theory.* Zurich lectures in advanced mathematics. European Mathematical Society, Zürich, 2008.

[21] Nicholas Sheridan. *Homological Mirror Symmetry for Calabi-Yau hypersurfaces in projective space.* PhD thesis, Massachusetts Institute of Technology, 2012.

[22] Richard Vale. 7350: Topics in finite-dimensional algebras, 2009. Lecture notes on a course at Cornell.

[23] Charles Weibel. *An Introduction to Homological Algebra.* Cambridge University Press, 1994.

[24] Eric Zaslow. Seidel's mirror map for the torus. Adv.Theor.Math.Phys.9:999-1006,2005, 2005.

A. Computing the mirror map

The following Python code (see figure 4) can be used to compute the j-invariant $j(\tau)$ of the mirror curve to arbitrary order, by expanding the relevant theta function formula. See section 7.1 for more details.

This code uses the free software symbolic computer algebra package Sympy[3] for manipulating power series. One convenient way to run it is to execute sympy in its build directory, and then load the script using `%run series.py`. This will automatically compute $j(x)$ to order x. More terms can be computed now using e.g. `j.series(x,n=3)` (but this gets slow pretty quickly).

```
from sympy import *

var('x')

class theta(Function):

    nargs = 3

    def _eval_nseries(self, x, n, logx):
        if x != self.args[2]:
            raise ValueError('series expansion not supported')
        [a, b, x] = self.args
        return sum(x**(a*(k + b)**2)
                   for k in range(-n+1, n)) + \
               O(x**(a*(n + b)**2), x) + O(x**(a*(-n + b)**2),
               x)

def b(i): return theta(9, i/S(18), x)
def a(i): return theta(3, i/S(6), x)
p = a(0)*b(0) + a(3)*b(9)
q = a(0)*b(6) + a(3)*b(3)
u = a(2)*b(0) + a(1)*b(9)
z = (2*q+p)/(3*u)
j = -27*z**3*(z**3+8)**3*(1-z**3)**(-3)
pprint(j.series(x, n=2))
```

Figure 4: Python script `series.py` to compute the Laurent expansion of $j(x)$.

[3]See `http://sympy.org`.