The A_∞ Deformation Theory of a Point and the Derived Categories of Local Calabi-Yaus

Edward Paul Segal

Imperial College London
Department of Mathematics

Thesis presented for the degree of
Doctor of Philosophy

Abstract

Let A be an augmented algebra over a semi-simple algebra S. We show that the Ext algebra of S as an A-module, enriched with its natural A-infinity structure, can be used to reconstruct the completion of A at the augmentation ideal. We use this technical result to justify a calculation in the physics literature describing algebras that are derived equivalent to certain non-compact Calabi-Yau three-folds. Since the calculation produces superpotentials for these algebras we also include some discussion of superpotential algebras and their invariants.

Contents

1 Acknowledgements

Thanks to Kai Behrend, Tom Bridgeland, Lieven Le Bruyn, Tom Coates, Alessio Corti, Joel Fine, Dominic Joyce, Alistair King, Raphaël Rouquier, Jim Stasheff and Balázs Szendröi for many helpful conversations, comments and ideas. Particular thanks are due to Kevin Costello and Paul Seidel for their help and hospitality at the University of Chicago. Finally I owe an enormous debt to my supervisor Richard Thomas, for teaching me (amongst other things) that having a concrete example sometimes means more than just specifying the ground field to be $\mathbb{C}$.

For Debbie

2 Introduction

There are now various examples known of the phenomenon whereby a variety X can be derived equivalent to a non-commutative algebra A. The pioneering example is due to Beilinson [5] who proved that the derived category of $\mathbb{P}^n$ is generated by the line bundles $\mathcal{O}, ..., \mathcal{O}(n)$. This equivalent to saying that the functor

$$\mathrm{RHom}(\bigoplus_{i=0}^{n} \mathcal{O}(i), -) : D^b(\mathbb{P}^n) \to D^b(A)$$

is a derived equivalence between $\mathbb{P}^n$ and the non-commutative algebra

$$A := \mathrm{End}(\bigoplus_{i=0}^{n} \mathcal{O}(i))$$

In fact compact examples like this are rare, much more progress has been made for non-compact examples, in particular for local models of resolutions of singularities [10],[33].

The phenomenon is also well known in the physics literature. There the variety X should be a Calabi-Yau threefold, and we study the type II superstring compactification on X . Type B D-branes in the theory correspond to objects in $D^b(X)$. It has been known since the work of Douglas and Moore [13] that if a D-brane sits at the centre of a singularity the effective theory on its world-volume is a gauge theory whose content can be described by a quiver diagram. This is the same as the mathematical results - the quiver diagram is a presentation of an algebra A which is derived equivalent to a resolution of the singularity. Since then other physical approaches (e.g [16]) have been found that produce an effective quiver gauge theory from branes on X.

The example that we are interested is when $X = \omega$ is the canonical bundle of a del Pezzo surface Z. This a Calabi-Yau three-fold, and it is again 'local' in that we may think of it as the normal bundle to an embedded surface in a compact Calabi-Yau. A first step in describing $D^b(\omega)$ is to describe $D^b(Z)$, and we specified that Z should be a del Pezzo because in that case Beilinson's approach has been generalised. What we do is find a special collection of line bundles $\{T_i\}$ on Z that generate $D^b(Z)$, then as before Z is derived equivalent to

$$A := \mathrm{End}_Z(\oplus_i T_i)$$

As Bridgeland observed in [8], we have a similar description of the derived category of ω. If we pull up the T_i via the projection $\pi : \omega \to Z$ we find that they still generate the derived category, so ω is derived equivalent (under one further assumption) to

$$\tilde{A} := \mathrm{End}_{\omega}(\oplus_i \pi^* T_i)$$

Both A and $\tilde{A}$ can be presented as quiver algebras (with relations), where the nodes of the quiver correspond to the line bundles in the collection. Suppose we have such a presentation of A. What do we have to do to it to produce a presentation of $\tilde{A}$?

This was the question addressed, in rather more physical language, by Aspinwall and Fidkowski in [1]. This paper is a mathematical interpretation of their work, and of related physics papers ([2], [26], [6] etc.). For the remainder of this introduction we will discuss the answer to this question, leaving out many subtleties and technicalities.

Suppose we have a presentation of the algebra A as the path algebra of a quiver Q (with relations), where nodes of Q correspond to the line bundles T_i. Then an A-module is precisely a representation of the quiver that obeys the relations. We have some obvious one-dimensional modules S_i which are the representations with just a one-dimensional vector space at the ith node. The direct sum

$$\mathcal{S} = \bigoplus_i S_i$$

of these is a representation which is one-dimensional at each node and with all the arrows sent to zero maps.

If we pick projective resolutions of each S_i then we can form the dga

$$\mathrm{RHom}_A(\mathcal{S}, \mathcal{S})$$

and then, using the process of homological perturbation ([15] etc...) transfer the dga structure to an A_∞-structure on its homology $\mathrm{Ext}_A(\mathcal{S}, \mathcal{S})$. Of course since A is derived equivalent to Z we could also view the S_i as being objects in $D^b(Z)$ and compute this A_∞-algebra there.

Now we consider the algebra $\tilde{A}$ corresponding to ω. This is also a quiver algebra on the same number of nodes, so has a similar set of one-dimensional modules $\tilde{S}_i$. It is easy to show that under the derived equivalence these map to the objects

$$\iota_* S_i \in D^b(\omega)$$

where $\iota : Z \to \omega$ is the zero section. We can again form the sum

$$\tilde{\mathcal{S}} = \bigoplus_i \tilde{S}_i$$

and the A_∞-algebra

$$\mathrm{Ext}_{\tilde{A}}(\tilde{\mathcal{S}}, \tilde{\mathcal{S}}) = \mathrm{Ext}_\omega(\iota_*\mathcal{S}, \iota_*\mathcal{S})$$

This new A_∞ algebra has a straightforward relationship with the previous one. By resolving the structure sheaf of the zero section and using Serre duality on Z one easily shows

$$\mathrm{Ext}_\omega(\iota_*\mathcal{S}, \iota_*\mathcal{S}) = \mathrm{Ext}_Z(\mathcal{S}, \mathcal{S}) \oplus \mathrm{Ext}_Z(\mathcal{S}, \mathcal{S})[3]^\vee \tag{2.1}$$

The two summands are dual under the Calabi-Yau pairing on $D^b(\omega)$, and the A_∞ structure should be cyclic with respect to this pairing. In fact with a little more work one can show that the A_∞ structure is given by formally extending the A_∞ structure on $\mathrm{Ext}_Z(\mathcal{S}, \mathcal{S})$ to make it cyclic. We call this procedure *cyclic completion.*

Now we come to the key point:

Claim 2.1. The algebra A is determined by the A_∞-algebra $\mathrm{Ext}_A(\mathcal{S}, \mathcal{S})$, in that if $\{m_i\}$ are the A_∞-products on $\mathrm{Ext}_A(\mathcal{S}, \mathcal{S})$ then the map

$$(\oplus_i m_i)^\vee : \mathrm{Ext}^2_A(\mathcal{S}, \mathcal{S})^\vee \to T^\bullet \mathrm{Ext}^1_A(\mathcal{S}, \mathcal{S})^\vee$$

is a presentation of A. Similarly the A_∞-algebra

$$\mathrm{Ext}_{\tilde{A}}(\tilde{\mathcal{S}}, \tilde{\mathcal{S}})$$

gives rise to a presentation for $\tilde{A}$.

This says that generators for A are given by (the dual space to) $\mathrm{Ext}^1_A(\mathcal{S}, \mathcal{S})$ and relations are given by $\mathrm{Ext}^2_A(\mathcal{S}, \mathcal{S})$, with the form of the relations being determined by the A_∞ structure. If we split $\mathcal{S}$ into its summands we see that this presentation is actually of a quiver algebra: the generating arrows between nodes i and j are given by $\mathrm{Ext}^1_A(S_i, S_j)$, and the relations on paths between i and j are given by $\mathrm{Ext}^2_A(S_i, S_j)$.

This claim is the hard part of the argument, and Section 3 of this paper is devoted to the discussion and proof of it. However for now we put it to one side and return to the question of determining $\tilde{A}$.

Suppose that we have a presentation for A of the form given in Claim 2.1. What is the corresponding presentation of $\tilde{A}$? Using (2.1):

$$\mathrm{Ext}^1_{\tilde{A}}(\tilde{S}_i, \tilde{S}_j) = \mathrm{Ext}^1_A(S_i, S_j) \oplus \mathrm{Ext}^2_A(S_j, S_i)^\vee$$

and

$$\mathrm{Ext}^2_{\tilde{A}}(\tilde{S}_i, \tilde{S}_j) = \mathrm{Ext}^2_A(S_i, S_j) \oplus \mathrm{Ext}^1_A(S_j, S_i)^\vee$$

So the answer is that for each existing relation on paths from node j to node i we should insert a new generating arrow going from i to j. Then for each existing generator going from i to j we put on one extra relation on paths going from j to i. To understand what the form of the relations should be we need to unpack our definition of 'cyclic completion'. It is easier to express the result if we introduce the notion of a *superpotential*.

In fact from the physics perspective, working out the superpotential is the primary goal, as it specifies the quiver gauge theory coming from ω. For the moment however we shall treat it just as the following little trick from linear algebra. The spaces $\mathrm{Ext}^1_{\tilde{A}}(\tilde{\mathcal{S}}, \tilde{\mathcal{S}})$ and $\mathrm{Ext}^2_{\tilde{A}}(\tilde{\mathcal{S}}, \tilde{\mathcal{S}})$ are dual under the Calabi-Yau pairing. Therefore the presentation

$$\mathrm{Ext}^2_{\tilde{A}}(\tilde{\mathcal{S}}, \tilde{\mathcal{S}})^\vee \to T^\bullet \mathrm{Ext}^1_{\tilde{A}}(\tilde{\mathcal{S}}, \tilde{\mathcal{S}})^\vee$$

alluded to in Claim 2.1 is given by an element

$$W \in T^\bullet \mathrm{Ext}^1_{\tilde{A}}(\tilde{\mathcal{S}}, \tilde{\mathcal{S}})^\vee$$

This is the superpotential for $\tilde{A}$. It is a formal non-commutative polynomial in the generators, and taking partial derivatives of it one recovers the relations. It is moreover cyclicly symmetric since the A_∞ structure is cyclic.

Now we can state the result. Suppose A is given by generators $\{x_1, ..., x_i\}$ and relations $\{\rho_1, ..., \rho_j\}$ (which are formal expressions in the x_i). Then the algebra $\tilde{A}$ is generated by the set $\{x_1, .., x_i, y_1, ..., y_j\}$ with relations coming from the superpotential

$$W = \sum_{\substack{\text{cyclic} \\ \text{permutations}}} \sum_j y_j \otimes \rho_j$$

2.1 An example

We illustrate the procedure with the prototypical example of $\mathbb{P}^2$, with the line bundles $T_i = \mathcal{O}(i)$, $i = 0, 1, 2$. The endomorphism algebra A of this collection is given by the Beilinson quiver

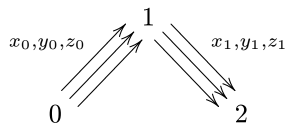

subject to the relations

$$x_0y_1 - y_0x_1 = 0$$
$$y_0z_1 - z_0y_1 = 0$$
$$z_0x_1 - x_0z_1 = 0$$

Now we pass to the local Calabi-Yau $\omega = \mathcal{O}(\mathbb{P}^2, -3)$, and pull up the line bundles. This corresponds to cyclically completing the quiver algebra. Firstly we insert extra arrows, dual to the relations. We have three relations, each of which applies to paths from T_0 to T_2. Hence we should insert three dual arrows from T_2 to T_0, so $\tilde{A}$ is generated by the quiver

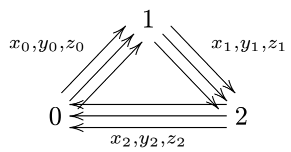

The superpotential is given by multiplying these new arrows by their corresponding relations, so it is

$$\begin{aligned} W &= \sum_{\substack{\text{cyclic} \\ \text{permutations}}} (x_0y_1 - y_0x_1)z_2 + (y_0z_1 - z_0y_1)x_2 + (z_0x_1 - x_0z_1)y_2 \\ &= \sum \epsilon^{ijk} x_i y_j z_k \end{aligned}$$

Now we compute the relations in $\tilde{A}$, which are given by taking formal partial derivatives of W. Taking derivatives with respect to the new generators just gives back the original three relations. Taking derivatives with respect to the original generators gives six new relations, each of which is a commutativity relation of the form of the one of the original relations but lying between a different pair of nodes.

According to our prescription the resulting algebra $\tilde{A}$ should be

$$\mathrm{End}_\omega(\pi^*\mathcal{O} \oplus \pi^*\mathcal{O}(1) \oplus \pi^*\mathcal{O}(2))$$

This is easily seen to be correct, since the latter is given by

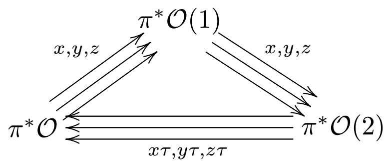

where τ is the tautological section of $\pi^*\mathcal{O}(-3)$.

2.2 The physical argument

It is instructive to look at the physical arguments involved in justifying Claim 2.1. The set-up is type II superstring theory on the ten-dimensional space $\omega \times \mathbb{R}^{3,1}$. We have a D3-brane, which is a $(3+1)$-dimensional object, extending in the flat directions, so from the point of view of ω it is just a point p. The effective (i.e. low-energy limit) theory on the world-volume of this brane is a gauge theory on $\mathbb{R}^{3,1}$. The quiver diagram for $\tilde{A}$ specifies this gauge theory - the nodes are $U(1)$ gauge groups, the arrows are fields, and the relations are constraints on the fields.

In terms of the derived category this D3-brane is the skyscraper sheaf $\mathcal{O}_p$. Under the derived equivalence between Z and $\tilde{A}$ this gets mapped to the $\tilde{A}$-module

$$\mathrm{RHom}_\omega(\bigoplus_i \pi^* T_i, \mathcal{O}_p) = \bigoplus_i (\pi^* T_i^\vee)|_p$$

This is a quiver representation that is one-dimensional at each node.

The moduli space of p is obviously just ω. However, on the other side of the derived equivalence it is also a moduli space $\mathcal{M}$ of quiver representations that are one-dimensional at each node, physically this is the vacuum moduli space of the quiver gauge theory. We can construct this space as follows. Suppose we have a presentation of $\tilde{A}$ as a quiver algebra with generating arrows V and some relations. Then a $(1, ..., 1)$-dimensional representation is just an assignment of a complex number to each generating arrow, such that the relations hold. This means that the space of such representations is a subvariety of $V^\vee$ cut out by the relations. Finally we must quotient this space by the gauge action of $\mathbb{C}^* \times ... \times \mathbb{C}^*$

given by changing the bases of the vector spaces at each node.

We now pick a Kähler metric on ω, which gives us a notion of stability for branes, and then deform the Kähler class to the limit where the metric collapses the zero section Z to a point. If our D3-brane was sitting at a point $p \in Z$ then it becomes unstable in this limit, and decays into a collection of so-called *fractional* branes, the $\tilde{S}_i$. We can see this mathematically in the construction of $\mathcal{M}$. When we take the gauge group quotient we should really pick a character χ of the gauge group and form the GIT quotient $\mathcal{M}^\chi$. For appropriate characters this should make the stable representations correspond precisely to points $p \in \omega$, and thus $\mathcal{M}^\chi = \omega$. But if we set $\chi = 0$ then all representations corresponding to points in Z become semi-stable and S-equivalent to the origin in $V^\vee$, which is the representation $\bigoplus_i \tilde{S}_i$. The moduli space $\mathcal{M}^\chi$ is then the singularity obtained by collapsing the zero section in ω.

Now the physics of the D3-brane is encoded in the superpotential W for the quiver gauge theory. This means that the equations of motion for p are the partial derivatives ∂W. However from the construction of $\mathcal{M}$ we know that the equations restricting p are precisely the relations in $\tilde{A}$, so in fact W is a superpotential in the mathematical sense for the algebra $\tilde{A}$.

On the other hand we can also see the behaviour of p by deforming $\oplus \tilde{S}_i$, since the deformation space is just $\mathcal{M}$. These deformations will be governed by the A_∞-algebra

$$\mathrm{Ext}_{\tilde{A}}(\oplus \tilde{S}_i, \oplus \tilde{S}_i)$$

in the sense that if $W' \in T^\bullet(\mathrm{Ext}^1)^\vee$ encodes the A_∞ structure then the critical locus of W' is the deformation space of $\oplus_i \tilde{S}_i$. Thus $W' = W$ is the superpotential for the quiver gauge theory, and hence for the algebra $\tilde{A}$.

2.3 Notation and basics

We will work over the ground field $\mathbb{C}$, although Section 3 works over an arbitrary ground field, and Section 4 works over any field of characteristic zero. $\mathbf{Alg}_\mathbb{C}$ is the category of associative unital $\mathbb{C}$-algebras. Undecorated tensor products will be over $\mathbb{C}$.

We will also need the category $\mathbb{C}^r$**-bimod** of bimodules over the semi-simple ring $\mathbb{C}^r$. We denote the obvious idempotents in $\mathbb{C}^r$ by $1_1, ..., 1_r$, then any

$V \in \mathbb{C}^r$-**bimod** is a direct sum of the subspaces

$$V_{ij} := 1_i.V.1_j$$

We may think of V as a 'categorified' vector space.

Let $\mathbf{Alg}^r_{\mathbb{C}}$ be the category of $\mathbb{C}^r$-algebras, i.e. associative unital algebra objects in $\mathbb{C}^r$-**bimod**. Equivalently this is the category of $\mathbb{C}$-linear categories whose objects form an ordered set of size r, if we only allow functors that preserve the ordering on the objects.

Any algebra $A \in \mathbf{Alg}^r_{\mathbb{C}}$ may be pictured as a quiver algebra (with relations) - just pick a basis for each A_{ij}, then A is a quotient of the path-algebra of the quiver with r nodes and arrows given by the basis elements. We may also consider A as an object of $\mathbf{Alg}_{\mathbb{C}}$ equipped with an ordered complete set of orthogonal idempotents $\{1_1, ..., 1_r\}$.

If V is any $\mathbb{C}^r$-bimodule then it generates a free $\mathbb{C}^r$-algebra

$$TV := \bigoplus_n V^{\otimes_{\mathbb{C}^r} n}$$

and a completed algebra

$$\hat{T}V := \prod_n V^{\otimes_{\mathbb{C}^r} n}$$

$\mathbf{Alg}^r_{\mathbb{C}}$ admits a symmetric monoidal product $\underline{\otimes}$ given by

$$(A \underline{\otimes} B)_{ij} = A_{ij} \otimes B_{ij}$$

Note that this is certainly not $A \otimes_{\mathbb{C}^r} B$, in general $A \otimes_{\mathbb{C}^r} B$ does not have an algebra structure.

An *augmentation* of an algebra $A \in \mathbf{Alg}^r_{\mathbb{C}}$ is a splitting $p : A \to \mathbb{C}^r$ of the inclusion $\mathbb{C}^r \hookrightarrow A$ of the identity arrows, or equivalently a choice of a two-sided ideal $\bar{A} \subset A$ such that $A/\bar{A} = \mathbb{C}^r$. Alternatively we may think of A as an algebra in $\mathbf{Alg}_{\mathbb{C}}$ for which we have chosen r $\mathbb{C}$-points $p : A \to \mathbb{C}^r$ and then split p. We denote the category of augmented algebras by $\mathbf{Alg}^{\star r}_{\mathbb{C}}$. Morphisms must respect the augmentations.

A module always means a left module. If we are picturing $A \in \mathbf{Alg}^r_{\mathbb{C}}$ as a quiver algebra then a module over A is precisely a representation of the quiver (that respects the relations). It is also the same as a functor $A \to \mathbf{Vect}$.

3 The A_∞ Deformation Theory of a Point

In this section we address the following claim, which we made in the introduction: suppose we have an appropriate set $\{S_i\}$ of one-dimensional modules for some algebra A. Then we can reconstruct A from the A_∞-algebra $\mathrm{Ext}_A(\oplus S_i, \oplus S_i)$.

In fact if we assume that A is graded, and that $A_0 = \oplus S_i$, then this statement has been part of the mathematical folklore for some time. The result is claimed (although not proven) by Keller [19] for a particular class of graded algebras, and his statement is closely related to a result of Laudal [25], who uses the terminology of Massey products. The fullest investigation to date appears to be the work of Lu, Palmieri, Wu and Zhang [28].

Let us start by explaining the statement a little. Let A be an $\mathbb{N}$-graded algebra over $\mathbb{C}$, and for simplicity let $A_0 = \mathbb{C}$. Now suppose we are given a presentation

$$A = TV/(\iota R)$$

so A is generated by a vector space V, modulo the two-sided ideal generated by a space of relations R under an inclusion

$$\iota : R \to TV$$

Assume that the presentation is minimal, in the sense that V and R are of minimal dimension. Then using the free resolution

$$... \to A \otimes R \to A \otimes V \to A \to \mathbb{C} \to 0$$

of $\mathbb{C} = A_0$ it is elementary to show that V must be dual to $\mathrm{Ext}^1_A(\mathbb{C}, \mathbb{C})$, and R must be dual to $\mathrm{Ext}^2_A(\mathbb{C}, \mathbb{C})$. Hence we might ask: if we are just given $\mathrm{Ext}^\bullet_A(\mathbb{C}, \mathbb{C})$, can we recover A?

We know immediately that A is generated by the space $V := (\mathrm{Ext}^1_A(\mathbb{C}, \mathbb{C}))^\vee$, and that relations are counted by the space $R := (\mathrm{Ext}^2_A(\mathbb{C}, \mathbb{C}))^\vee$, but we still need to know what form these relations take, i.e. we need the map

$$\iota : R \to TV$$

or dually, a map

$$\iota^\vee : \hat{T}\mathrm{Ext}^1_A(\mathbb{C},\mathbb{C}) \to \mathrm{Ext}^2_A(\mathbb{C},\mathbb{C})$$

We certainly have something that might be a part of this map, namely the usual Yoneda (wedge) product, which is a map

$$\mathrm{Ext}^1_A(\mathbb{C},\mathbb{C})^{\otimes 2} \to \mathrm{Ext}^2_A(\mathbb{C},\mathbb{C})$$

If we knew that our relations were purely quadratic then we might reasonably conjecture that this dualising this map gave a presentation of A. In fact although this is true for many algebras it is false in general - the study of those algebras for which it works is the subject of classical Koszul duality. What happens when our relations are definitely not just quadratic? Then we would need, in addition to the bilinear Yoneda product, some 'higher' multi-linear products

$$m_i : \mathrm{Ext}^1_A(\mathbb{C},\mathbb{C})^{\otimes i} \to \mathrm{Ext}^2_A(\mathbb{C},\mathbb{C})$$

Fortunately these higher products do exist (though not quite canonically), they form an A_∞-structure on $\mathrm{Ext}^\bullet_A(\mathbb{C},\mathbb{C})$ which measures the failure of the dga $\mathrm{RHom}_A(\mathbb{C},\mathbb{C})$ to be formal. Furthermore when you dualize they do indeed give a presentation of A. It is this result (essentially our Theorem 3.16) that is proven by [28].

One of our original aims was to prove this result for the case $A_0 = \mathbb{C}^r$, i.e. when A is a graded quiver algebra (with relations) on r vertices. However, given the proof in [28] this is easy - you simply change your ground category from vector spaces to the category of $\mathbb{C}^r$-bimodules (which one may picture as vector spaces strung between r vertices) and the same proof works. Instead we take a different tack which we feel is a bit more conceptual.

It seemed to us that the graded hypothesis was a little unnatural. We instead ask what happens if we take an arbitrary algebra A with a one-dimensional module S and perform the same construction, i.e. take the Yoneda algebra $\mathrm{Ext}_A(S,S)$ equipped with A_∞ products $\{m_i\}$, then dualize the map

$$m = \oplus_i m_i : T\mathrm{Ext}^1_A(S,S) \to \mathrm{Ext}^2_A(S,S)$$

to get the presentation of a new algebra

$$E := \frac{\hat{T}\mathrm{Ext}^1_A(S,S)^\vee}{(m^\vee \mathrm{Ext}^2_A(S,S)^\vee)}$$

What is this new algebra? Firstly note that a one-dimensional module is just a map $p : A \to \mathbb{C}$. Hence if A is commutative then this is simply a closed point of the affine scheme $\operatorname{Spec}(A)$, and the module is its sky-scraper sheaf. It is then geometrically obvious that the algebra E can only depend on a formal neighbourhood of the point p. In fact the result is that E is precisely the formal neighbourhood of p, i.e. it is the completion of A at the kernel of p. We explain this result (which contains nothing new) informally in Section 3.1, the key point is that the A_∞-algebra $\operatorname{Ext}_A(S, S)$ controls the deformations of the module S and hence those of p.

This is of course the case $r = 1$, in general we wish to pick r points

$$p = \oplus p_i : A \to \mathbb{C}^r$$

so that $\mathbb{C}^r$ becomes an A-module (strictly speaking we must also choose a splitting of the the map p, so this is more like choosing a single point of a $\mathbb{C}^r$-algebra). Then our main result (Theorem 3.14) is that performing the above construction on $\operatorname{Ext}_A(\mathbb{C}^r, \mathbb{C}^r)$ again produces the completion of A at the kernel of p.

If we stick with a commutative A then this generalization is trivial, since $\operatorname{Ext}_A(\mathbb{C}^r, \mathbb{C}^r)$ splits as a direct product over the different points. If A is non-commutative however this is no longer true and the proof becomes rather more difficult. In particular it is not correct to study deformations of $\mathbb{C}^r$ as an A-module, instead one should follow [25] and study non-commutative deformations of the category whose objects are these r one-dimensional modules. The technical challenge of this paper is checking that the sketch proof given for the $r = 1$ commutative case continues to work in the general setting, which means firstly checking that the non-commutative deformations of a set of modules are governed by the A_∞-category of their Ext groups (Section 3.2) and secondly relating deformations of a set of one-dimensional modules to deformations of the corresponding set of points (Section 3.3).

Having understood this more general situation it is then straightforward to deduce the required result for graded algebras, which we do in Section 3.4. This is because the completion of an $\mathbb{N}$-graded algebra at its positively-graded ideal contains the original algebra in a natural way.

3.1 A geometric sketch

Let $X = \operatorname{Spec} A$ be an affine scheme over $\mathbb{C}$, and let $p : A \to \mathbb{C}$ be a point.

Obviously the formal deformations of p see precisely a formal neighbourhood of p in X, the algebraic way to say this is that the formal deformation functor of p is pro-represented by the completion $\hat{A}_p$ of A at the kernel of p.

It is easy to show that the deformation theory of p is precisely the same as the deformation theory of the associated 'sky-scraper' sheaf $\mathcal{O}_p$, i.e. the 1-dimensional A-module given by p. In accordance with the philosophy of dga (or dgla, see Remark 3.1) deformation theory, the deformations of $\mathcal{O}_p$ are governed by the differential graded algebra $\mathrm{RHom}_X(\mathcal{O}_p, \mathcal{O}_p)$, by which we mean that formal deformations of $\mathcal{O}_p$ are formal solutions of the *Maurer-Cartan* equation

$$MC : \mathrm{RHom}^1_X(\mathcal{O}_p, \mathcal{O}_p) \to \mathrm{RHom}^2_X(\mathcal{O}_p, \mathcal{O}_p)$$

$$MC(a) := da + a^2 = 0$$

taken modulo the 'infinitesimal gauge action' of $\mathrm{RHom}^0_X(\mathcal{O}_p, \mathcal{O}_p)$. This resulting 'formal deformation space' is, as we just said, simply the formal scheme $\hat{A}_p$.

According to Kontsevich [21] the formal deformation theory attached to a dga is a homotopy invariant, so we may replace our dga by any quasi-isomorphic A_∞-algebra and compute the deformations there instead. The Maurer-Cartan equation picks up higher terms from the A_∞ structure and becomes the *Homotopy Maurer-Cartan equation*:

$$HMC(a) := \sum_i m_i(a^{\otimes i}) = 0$$

and there are similar homotopy corrections to the gauge action. In particular, using the process of homological perturbation, we may replace $\mathrm{RHom}_X(\mathcal{O}_p, \mathcal{O}_p)$ by its homology $\mathrm{Ext}_X(\mathcal{O}_p, \mathcal{O}_p)$ equipped with an appropriate A_∞-structure.

Since $\mathrm{Ext}^0_X(\mathcal{O}_p, \mathcal{O}_p) = \mathbb{C}$, the gauge action is now trivial, so the formal deformation space is the formal zero locus of

$$HMC : \mathrm{Ext}^1_X(\mathcal{O}_p, \mathcal{O}_p) \to \mathrm{Ext}^2_X(\mathcal{O}_p, \mathcal{O}_p)$$

The algebra of functions on this formal scheme is the formal power series ring on Ext^1 modulo the ideal generated by the $\mathbb{C}$-linear dual of HMC, so we have shown

$$\frac{\mathbb{C}[[\ \mathrm{Ext}^1_X(\mathcal{O}_p, \mathcal{O}_p))^\vee\ ]]}{\big(\ HMC^\vee(\mathrm{Ext}^2_X(\mathcal{O}_p, \mathcal{O}_p)^\vee)\ \big)} = \hat{A}_p$$

This says that formally around p, X is cut out of T_pX by the HMC equation, with $\mathrm{Ext}^2_X(\mathcal{O}_p, \mathcal{O}_p)$ being a canonical space of obstructions. Of course if p is a smooth point the denominator of this expression should vanish, since there are no (commutative) obstructions. However in that case Ext^2 is precisely the commutativity relations, which tells us that we are really measuring non-commutative obstructions, and that the numerator should really be non-commutative power series.

Remark 3.1. It is more traditional to control deformations with dg-Lie (or L_∞) algebras, but in this paper we will always in fact have a dg (or A_∞) algebra. We should really take the associated commutator algebra, as this is all that the deformation theory depends on, but we shall not bother to do so.

3.2 Deformation theory of sets of modules

Let $A \in \mathbf{Alg}_\mathbb{C}$ and let $\mathcal{M} = \{M_1, ..., M_r\}$ be a set of A-modules. We show how the non-commutative deformation theory of $\mathcal{M}$ as developed by Laudal [25] may be viewed as a dga deformation problem.

If we wanted to deform a single module M then we would just deform the module map

$$\mu : A \to \mathrm{End}_\mathbb{C}(M)$$

When we have a set of modules we can form the endomorphism algebra

$$\mathrm{End}_\mathbb{C}(\mathcal{M}) := \mathrm{End}_\mathbb{C}(\oplus_i M_i)$$

and we could deform the map

$$\mu = \oplus_i \mu_i : A \to \mathrm{End}_\mathbb{C}(\mathcal{M})$$

If we just treat this as a map of $\mathbb{C}$-algebras and deform it then we are just studying deformations of $\oplus_i M_i$ as an A-module. We wish to do something slightly different, and use the fact that $\mathrm{End}_\mathbb{C}(\mathcal{M})$ is actually a $\mathbb{C}^r$-algebra.

Recall that $\mathbf{Alg}^{\star r}_\mathbb{C}$ is the category of augmented algebras over $\mathbb{C}^r$. Let

$$\mathbf{Art}^r_\mathbb{C} \subset \mathbf{Alg}^{\star r}_\mathbb{C}$$

be the subcategory of consisting of algebras $(R, \mathfrak{m})$ for which the augmentation ideal $\mathfrak{m}$ is nilpotent. Of course $\mathbf{Art}^1_\mathbb{C}$ is just the category of Artinian local (non-commutative) $\mathbb{C}$-algebras. Recall also the product of two $\mathbb{C}^r$-algebras that we defined by

$$(A\underline{\otimes}B)_{ij} = A_{ij} \otimes B_{ij}$$

Definition 3.2. [25] For $(R, \mathfrak{m}) \in \mathbf{Art}^{\mathbf{r}}_{\mathbb{C}}$, an R-deformation of $\mathcal{M}$ is a map of $\mathbb{C}$-algebras

$$\mu_R : A \to \mathrm{End}_{\mathbb{C}}(\mathcal{M})\underline{\otimes}R$$

which reduces modulo $\mathfrak{m}$ to the given module maps

$$\bigoplus_i \mu_i : A \to \bigoplus_i \mathrm{Hom}_{\mathbb{C}}(M_i, M_i)$$

Two R-deformations are equivalent if they differ by an inner automorphism of $\mathbb{C}$-algebras in $\mathrm{End}_{\mathbb{C}}(\mathcal{M})\underline{\otimes}R$.

Let $\mathcal{D}ef_{\mathcal{M}} : \mathbf{Art}^r_{\mathbb{C}} \to \mathbf{Set}$ be the resulting deformation functor.

We are going to present this as a dga (and later A_∞) deformation problem. It is well known that the deformation functor of a single module M can be seen as a dga deformation problem - it is controlled by the dga $\mathrm{RHom}_A(M, M)$. We now show that a similar statement is true for our deformation functor $\mathcal{D}ef_{\mathcal{M}}$, but since we are deforming $\mathbb{C}^r$-algebras we look not for a dga but for a dg-category with r objects.

Definition 3.3. Let $(\mathcal{A}^\bullet, d, m)$ be a dga over $\mathbb{C}^r$. The *deformation functor* associated to $\mathcal{A}$ is the functor

$$\mathcal{D}ef_{\mathcal{A}} : \mathbf{Art}^r_{\mathbb{C}} \to \mathbf{Set}$$

which sends $(R, \mathfrak{m})$ to the set

$$\left\{a \in \mathcal{A}^1\underline{\otimes}\mathfrak{m};\;\; da + m(a \otimes a) = 0\right\} / \sim$$

where the equivalence relation $\sim$ is given by taking the exponential of the following action of the commutator Lie algebra of $\mathcal{A}^0\underline{\otimes}\mathfrak{m}$

$$b : a \to a + db - [b, a]$$

For our deformation problem the obvious choice of dg-$\mathbb{C}^r$-algebra is

$$\mathrm{RHom}_A(\mathcal{M}, \mathcal{M}) = \bigoplus_{i,j} \mathrm{RHom}_A(M_i, M_j)$$

This is only defined up to quasi-isomorphism. To produce models for it we need to resolve each M_i, which we may do using the following standard construction:

Definition 3.4. For an A-module M the *bar resolution* of M to be the complex of free A-modules (concentrated in non-positive degrees)

$$B(A,M)^{-t} := A^{\otimes t+1} \otimes M$$

with differential given by

$$d(a_1 \otimes ... \otimes a_t \otimes m) = \sum_{s=2}^{t} (-1)^s a_1 \otimes ... \otimes a_{s-1} a_s \otimes ... \otimes a_t \otimes m \;\; - \;\; (-1)^t a_1 \otimes ... \otimes a_t m$$

Lemma 3.5. *The module map $\mu : A \otimes M \to M$ induces a quasi-isomorphism*

$$\mu : B(A,M) \to M$$

Proof. Since A is unital μ is a surjection, and $B(A,M)$ is acyclic in all negative degrees since

$$d(1_A \otimes \mathbf{b}) = \mathbf{b} - 1_A \otimes d(\mathbf{b})$$

for any $\mathbf{b} \in B(A,M)^{<0}$. □

Hence one model for $\mathrm{RHom}_A(\mathcal{M},\mathcal{M})$ is given by the dg-category $\mathcal{E}$ whose hom-sets are

$$\mathcal{E}_{ij} := \mathrm{Hom}_A(B(A,M_i), B(A,M_j))$$

However there is a simpler candidate. Consider the dg-category $\mathcal{H}$ whose hom-sets are

$$\mathcal{H}_{ij} := \mathrm{Hom}_A(B(A,M_i), M_j)$$

with composition

$$(f \bullet g)(a_1 \otimes ... \otimes a_{s+t+1} \otimes m) := f(a_1 \otimes ... \otimes a_{s+1} \otimes g(1_A \otimes a_{s+2} \otimes ... \otimes a_{s+t+1} \otimes m))$$

for homogeneous maps f, g of degrees s and t. This composition was obtained as follows: $B(A,\mathbb{C})$ is naturally a coalgebra under the 'shuffle' coproduct, and $B(A,\mathcal{M})$ is a comodule over it. We are letting $f \bullet g = f(\mathbf{1} \otimes g)\mu$ where μ is the comodule map. Now for each i, j we have a quasi-isomorphism

$$\mu_j : B(A, M_j) \to M_j$$

which induces a quasi-isomorphism of chain complexes

$$\mu_j \circ : \mathcal{E}_{ij} \to \mathcal{H}_{ij}$$

(since the $B(A, M)^i$ are free). These do not form a map of dg-categories since they do not respect the compositions, but they do have a right-inverse which is a map of dg-categories:

Lemma 3.6. *Let*

$$\Psi : \mathcal{H}^t \to \mathcal{E}^t$$

be given by

$$\Psi(f)(a_1 \otimes ... \otimes a_{s+1} \otimes m) := a_1 \otimes ... \otimes a_{s-t+1} \otimes f(1_A \otimes a_{s-t+2} \otimes .. \otimes a_{s+1} \otimes m)$$

for $s \geq t$ *and*

$$\Psi(f)(a_1 \otimes ... \otimes a_{s+1} \otimes m) := 0$$

for $s < t$. *Then* Ψ *is a map of dg-categories such that* $(\mu\circ)\Psi = \mathbf{1}_{\mathcal{H}}$, *hence it is a quasi-isomorphism of dg-categories.*

Proof. Elementary (though tedious) from definitions. □

We use this model $\mathcal{H}$ for $\mathrm{RHom}_A(\mathcal{M}, \mathcal{M})$ to show that this dg-category controls the deformations of $\mathcal{M}$, at least (as in the single module case) up to module automorphisms of the M_i.

Proposition 3.7. $\mathcal{D}ef_{\mathcal{M}}$ *is a quotient by* $Aut_A(\bigoplus_i M_i)$ *of the deformation functor* $\mathcal{D}ef_{\mathcal{H}}$ *associated to the dg-category* $\mathcal{H}$.

Proof. Let R be an object of $\mathbf{Art}^r_{\mathbb{C}}$. Using the splitting $R = \mathbb{C}^r \oplus \mathfrak{m}$ we can write any R-deformation μ_R of M as

$$\mu_R = \bigoplus_i \mu_i + \tilde{\mu}_R$$

where $\tilde{\mu}_R : A \to \mathrm{End}_{\mathbb{C}}(\mathcal{M}) \underline{\otimes} \mathfrak{m}$. The R-points of $\mathcal{D}ef_{\mathcal{H}}$ are those (equivalence classes of) elements of $\mathcal{H}^1\underline{\otimes}\mathfrak{m}$ that obey the Maurer-Cartan equation. However,

$$\begin{aligned}
\mathcal{H}^1\underline{\otimes}\mathfrak{m} &= \bigoplus_{i,j} \mathrm{Hom}_A(A^{\otimes 2}\otimes M_i, M_j)\otimes \mathfrak{m}_{ij} \\
&= \mathrm{Hom}_{\mathbb{C}}(A, \bigoplus_{i,j} \mathrm{Hom}_{\mathbb{C}}(M_i, M_j)\otimes \mathfrak{m}_{ij}) \\
&= \mathrm{Hom}_{\mathbb{C}}(A, \mathrm{End}_{\mathbb{C}}(\mathcal{M})\underline{\otimes}\mathfrak{m})
\end{aligned}$$

and for an element $\tilde{\mu}_R \in \mathcal{H}^1\underline{\otimes}\mathfrak{m}$ the Maurer-Cartan equation is precisely the condition that $\bigoplus_i \mu_i + \tilde{\mu}_R$ is a map of algebras.

Now we compare the equivalence relations on each side. The class of $\mu_R = \bigoplus_i \mu_i + \tilde{\mu}_R$ in $\mathcal{D}ef_{\mathcal{M}}(R)$ is its orbit under conjugation by the the subgroup

$$\mathrm{Stab}(\bigoplus_i \mu_i)$$

of the group of invertible elements in $\mathrm{End}_{\mathbb{C}}(\mathcal{M})\underline{\otimes}R$. We have the obvious factorization

$$(1+\mathrm{End}_{\mathbb{C}}(\mathcal{M})\underline{\otimes}\mathfrak{m}) \;\to\; \mathrm{Stab}(\bigoplus_i \mu_i) \;\to\; \mathrm{Aut}_A(\bigoplus_i M_i)$$

The Lie algebra of $(1+\mathrm{End}_{\mathbb{C}}(\mathcal{M})$ is the commutator algebra of

$$\mathrm{End}_{\mathbb{C}}(\mathcal{M})\underline{\otimes}\mathfrak{m} = \mathcal{H}^0\underline{\otimes}\mathfrak{m}$$

The equivalence relation on $\mathcal{D}ef_{\mathcal{H}}(R)$ is given by integrating the 'infinitesimal gauge action' of $\mathcal{H}^0\underline{\otimes}\mathfrak{m}$, but this action is precisely the derivative of conjugation. Hence the orbits in $\mathcal{D}ef_{\mathcal{M}}(R)$ are the quotients of the orbits in $\mathcal{D}ef_{\mathcal{H}}(R)$ under the residual action by $\mathrm{Aut}_A(\bigoplus_i M_i)$. □

Corollary 3.8. *If the M_i are simple and distinct then $\mathcal{D}ef_{\mathcal{M}} = \mathcal{D}ef_{\mathcal{H}}$*

In light of Lemma 3.6 the homology of $\mathcal{H}$ is the category

$$\mathrm{Ext}_A(\mathcal{M}) = \bigoplus_{i,j} \mathrm{Ext}_A(M_i, M_j)$$

Using homological perturbation, we may put an A_∞-structure on this category (unique up to A_∞-isomorphism) such that it is A_∞-quasi-isomorphic to $\mathcal{H}$. We can now use this A_∞-algebra to compute the deformation functor of $\mathcal{M}$.

Definition 3.9. Let $(\mathcal{A}^\bullet, m_i)$ be an A_∞-algebra over $\mathbb{C}^r$. The *deformation functor* associated to $\mathcal{A}$ is the functor

$$\mathcal{D}ef_{\mathcal{A}} : \mathbf{Art}^r_{\mathbb{C}} \to \mathbf{Set}$$

which sends $(R, \mathfrak{m})$ to the set

$$\left\{ a \in \mathcal{A}^1 \underline{\otimes} \mathfrak{m};\ \sum_i m_i(a) = 0 \right\} / \sim$$

The equivalence relation $\sim$ is generated by the following map from $\mathcal{A}^0 \underline{\otimes} \mathfrak{m}$ to vector fields on $\mathcal{A}^1 \underline{\otimes} \mathfrak{m}$:

$$b : a \to a + \sum_{n \geq 1} (-1)^{n(n+1)/2} \frac{1}{n} \sum_{t=0}^{n-1} (-1)^t m_n(a^{\otimes t} \otimes b \otimes a^{\otimes n-t-1})$$

We have obtained this from the usual definition of the deformation functor of an L_∞-algebra (e.g. [26]) in the case that the L_∞-algebra is actually the commutator algebra of an A_∞-algebra.

Corollary 3.10. *$\mathcal{D}ef_{\mathcal{M}}$ is a quotient by $Aut_A(\bigoplus_i M_i)$ of the deformation functor associated to the A_∞-category $Ext_A(\mathcal{M})$.*

Proof. It is standard (e.g. [21]) that for A_∞-algebras over a field the usual commutative deformation functor is a homotopy invariant. The key point of the standard proof is that any A_∞-algebra is isomorphic to the direct sum of a minimal one and a linear contractible one, but in fact this holds when the base category is any semi-simple linear monoidal category [27] so it works over $\mathbb{C}^r$. The remainder of the proof consists of checking three things: that A_∞-morphisms induce natural transformations of deformation functors, that deformation functors commute with direct sums, and that the deformation functor associated to a linear contractible A_∞-algebra is trivial. These are easily checked to hold for our non-commutative deformation functors as well. □

3.3 Deforming a set of points

Let A be a $\mathbb{C}$-algebra, and let $p : A \to \mathbb{C}^r$ be a set of r $\mathbb{C}$-points of A. Each point $p_i : A \to \mathbb{C}$ gives a one-dimensional A-module which we call $\mathcal{O}_{p_i}$. We wish to relate the deformations of the points p to the deformations of the set of modules $\mathcal{M} := \{\mathcal{O}_{p_i}\}$.

In fact to achieve this in the way that we want we have to start with a little more data - we have to choose a splitting of the map p, so that A becomes an augmented $\mathbb{C}^r$-algebra. Then we can consider deformations of the map p in the category $\mathbf{Alg}^r_{\mathbb{C}}$, these are just $\mathbb{C}^r$-algebra maps

$$p_R : A \to R$$

where $R \in \mathbf{Art}^r_{\mathbb{C}}$ and p_R reduces to p modulo the augmentation ideal in R. In other words p_R must be a morphism of augmented $\mathbb{C}^r$-algebras, so the deformation functor of p (in $\mathbf{Alg}^r_{\mathbb{C}}$) is

$$\mathbf{Alg}^{\star r}_{\mathbb{C}}(A, -) : \mathbf{Art}^r_{\mathbb{C}} \to \mathbf{Set}$$

How does this relate to the deformation functor $\mathcal{D}ef_{\mathcal{M}}$ of the set of the modules?

Proposition 3.11. *For any* $(R, \mathfrak{m}) \in \mathbf{Art}^r_{\mathbb{C}}$ *there is a functorial isomorphism*

$$\mathcal{D}ef_{\mathcal{M}}(R) = \mathbf{Alg}^{\star r}_{\mathbb{C}}(A, R) \ / \ \{\textit{Inner } \mathbb{C}^r\textit{-algebra automorphisms}\}$$

Proof. Recall (Definition 3.2) that $\mathcal{D}ef_{\mathcal{M}}(R)$ is the set of $\mathbb{C}$-algebra maps

$$p_R : A \to \mathrm{End}_{\mathbb{C}}(\mathcal{M}) \underline{\otimes} R$$

that reduce to p modulo $\mathfrak{m}$, taken up to inner automorphism of $\mathbb{C}$-algebras. Now the hom-sets of $\mathrm{End}_{\mathbb{C}}(\mathcal{M})$ are all 1-dimensional, so

$$\mathrm{End}_{\mathbb{C}}(\mathcal{M}) \underline{\otimes} R = R$$

for any R (this fact is the reason that we defined our deformations in terms of the product $\underline{\otimes}$). So to prove the proposition we just need to check that $\mathbb{C}$-algebra maps $A \to R$ that preserve the augmentations are the same thing as augmented $\mathbb{C}^r$-algebra maps, once we have quotiented out by inner automorphisms on both sides. This is done in the following lemma. ☐

The difference between an inner $\mathbb{C}$-algebra automorphism of R and an inner $\mathbb{C}^r$-algebra automorphism is that the former is conjugation by an arbitrary element $r \in R$ whereas for the latter we must take r from the diagonal subalgebra

$$\bigoplus_i R_{ii} \subset R$$

in order to preserve the $\mathbb{C}^r$-algebra structure.

Lemma 3.12. *Let (B, p) be an augmented $\mathbb{C}^r$-algebra in $\mathbf{Alg}_{\mathbb{C}}^{\star r}$ and let $(R, \mathfrak{m}) \in \mathbf{Art}_{\mathbb{C}}^r$. Then maps of $\mathbb{C}$-algebras $B \to R$ preserving the augmentations, taken up to inner automorphism of $\mathbb{C}$-algebras, biject with maps of augmented $\mathbb{C}^r$-algebras $B \to R$ taken up to inner automorphism of $\mathbb{C}^r$-algebras.*

Proof. Obviously $\mathbb{C}^r$-algebra maps form a subset of $\mathbb{C}$-algebra maps, and if two $\mathbb{C}^r$-algebra maps are conjugate as $\mathbb{C}^r$-algebra maps then they are also conjugate as $\mathbb{C}$-algebra maps. Hence it is suffient to prove that if $f : B \to R$ is a $\mathbb{C}$-algebra map preserving the augmentations then it is conjugate to a map of $\mathbb{C}^r$-algebras. Let 1_i denote the ith direct summand of the identity (i.e. the identity arrow at the ith object) in either B or R. Since f preserves the augmentations we must have

$$f(1_i) = 1_i + m_i$$

for some elements $m_i \in \mathfrak{m}$, and since f is an algebra map we must have $\{1_i + m_i\}$ a complete orthogonal set of idempotents. If we conjugate by the element

$$\sum_i 1_i + m_i 1_i = 1 + \sum_i m_i 1_i \in R$$

(which is invertible since $\mathfrak{m}$ is nilpotent) then we get a map that sends $1_i \in B$ to $1_i \in R$ for all i and hence is a map of $\mathbb{C}^r$-algebras. □

Now we consider the equivalent description (Corollary 3.10) of our deformation functor $\mathcal{D}ef_{\mathcal{M}}$ as being the deformation functor $\mathcal{D}ef_{Ext(\mathcal{M})}$ associated to the A_∞-category $\mathrm{Ext}_A(\mathcal{M})$.

Firstly recall the construction (discussed in the Introduction) of an algebra from the degree one and two parts of an A_∞ algebra. Let the higher products on $\mathrm{Ext}_A(\mathcal{M})$ be denoted m_i. The degrees of the m_i dictate that for all i

$$m_i : \mathrm{Ext}_A^1(\mathcal{M})^{\otimes i} \to \mathrm{Ext}_A^2(\mathcal{M})$$

The direct sum of all these maps is the homotopy Maurer-Cartan function

$$HMC = \bigoplus_{i>0} m_i : T(\mathrm{Ext}_A^1(\mathcal{M})) \to \mathrm{Ext}_A^2(\mathcal{M})$$

The m_i and HMC are all maps of $\mathbb{C}^r$-bimodules. Now given a $\mathbb{C}^r$-bimodule V, its $\mathbb{C}$-linear dual $V^\vee = \mathrm{Hom}_{\mathbb{C}}(V, \mathbb{C})$ is also a $\mathbb{C}^r$-bimodule. Dualizing the map

HMC in this way we get a map

$$HMC^\vee : \mathrm{Ext}^2_A(\mathcal{M})^\vee \to \hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee)$$

Quotienting by the two-sided ideal generated by the image gives us a $\mathbb{C}^r$-algebra

$$\frac{\hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee)}{(\mathrm{Ext}^2_A(\mathcal{M})^\vee)}$$

This is naturally augmented because there is no m_0 term in HMC.

We claim that this algebra is (nearly) the ring of functions on the formal deformation space associated to $\mathrm{Ext}_A(\mathcal{M})$. Our use of the word 'space' here is a little shaky, since this is a non-commutative $\mathbb{C}^r$-algebra and we are not proposing to define Spec of it! Nevertheless the statement makes rigourous sense if we interpret it at the level of deformation functors.

At this point we must insert an extra assumption.

Proposition 3.13. *Assume that $Ext^1_A(\mathcal{O}_{p_i}, \mathcal{O}_{p_j})$ is finite-dimensional for each i, j. Then for any $(R, \mathfrak{m}) \in \mathbf{Art}^r_{\mathbb{C}}$ there is a functorial isomorphism*

$$\mathcal{D}ef_{Ext(\mathcal{M})}(R) = \mathbf{Alg}^{\star r}_{\mathbb{C}}\left(\frac{\hat{T}(Ext^1_A(\mathcal{M})^\vee)}{(Ext^2_A(\mathcal{M})^\vee)}, R\right) \;/\; \{\textit{Inner } \mathbb{C}\textit{-algebra automorphisms in } R\}$$

Proof. Ignoring the gauge group for a minute, we have that $\mathcal{D}ef_{Ext(\mathcal{M})}(R)$ is the zero locus of the homotopy Maurer-Cartan function in

$$\begin{aligned} \mathrm{Ext}^1(\mathcal{M})\underline{\otimes}\mathfrak{m} &= \bigoplus_{i,j} \mathrm{Hom}_{\mathbb{C}}(\mathrm{Ext}^1_A(\mathcal{O}_{p_i}, \mathcal{O}_{p_j})^\vee, \mathfrak{m}_{ij}) && (3.1) \\ &= \mathbf{Alg}^{\star r}_{\mathbb{C}}(\hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee), R) && (3.2) \end{aligned}$$

Here we have used our finite-dimensionality assumption on the Ext^1s. A map in $\mathbf{Alg}^{\star r}_{\mathbb{C}}(\hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee), R)$ is a zero of HMC precisely when it induces a map

$$\frac{\hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee)}{(\mathrm{Ext}^2_A(\mathcal{M})^\vee)} \to R$$

on the quotient algebra.

The gauge group is the exponential of

$$\mathrm{Ext}^0_A(\mathcal{M})\underline{\otimes}\mathfrak{m} = \bigoplus_i \mathfrak{m}_{ii}$$

which is $1+\bigoplus_i \mathfrak{m}_{ii} \subset R$, and it acts by conjugacy on $\mathbf{Alg}^{\star r}_{\mathbb{C}}(\hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee), R)$. Now an arbitrary inner $\mathbb{C}$-algebra automorphism of R is a conjugation by an element in

$$(\mathbb{C}^*)^r + \bigoplus_i \mathfrak{m}_{ii}$$

but $(\mathbb{C}^*)^r$ is in the centre of R so the orbits under the gauge group are precisely inner-$\mathbb{C}$-algebra-automorphism classes.

Functoriality follows from the functoriality of equations (3.1) and (3.2). □

We have nearly proved our main theorem. If we forget about inner automorphisms, we have shown than $\mathcal{D}ef_{\mathcal{M}}$ is the same as $\mathbf{Alg}^{\star r}_{\mathbb{C}}(A, -)$. This functor is *pro-representable*, i.e. it is represented by the completion $\hat{A}_p$ of A at the augmentation ideal (the kernel of p), which can be thought of as a directed system of objects in $\mathbf{Art}^r_{\mathbb{C}}$. On the other hand we know $\mathcal{D}ef_{\mathcal{M}} = \mathcal{D}ef_{\mathrm{Ext}(\mathcal{M})}$, which we have just shown to be pro-represented by

$$\frac{\hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee)}{(\mathrm{Ext}^2_A(\mathcal{M})^\vee)}$$

A pro-representing object is unique, so these two formal $\mathbb{C}^r$-algbras must be the same.

Now we just have to check that this argument still holds when we remember about the inner automorphisms, but this is just a matter of carefully checking the standard proof of uniqueness of a pro-representing object.

Theorem 3.14. *Let*

$$\hat{E} = \frac{\hat{T}(Ext^1_A(\mathcal{M})^\vee)}{(Ext^2_A(\mathcal{M})^\vee)}$$

then $\hat{E}$ is isomorphic as an augmented $\mathbb{C}^r$-algebra to the completion $\hat{A}_p$ of A at the kernel of p.

Proof. We have shown (Corollary 3.10, Proposition 3.11 and Proposition 3.13) that there is a natural isomorphism Ψ between

$$\frac{\mathbf{Alg}^{\star r}_{\mathbb{C}}(A, -)}{\text{Inner Aut.}} = \frac{\mathbf{Alg}^{\star r}_{\mathbb{C}}(\hat{A}_p, -)}{\text{Inner Aut.}} : \mathbf{Art}^r_{\mathbb{C}} \to \mathbf{Set}$$

and

$$\frac{\mathbf{Alg}^{\star r}_{\mathbb{C}}(\hat{E}, -)}{\text{Inner Aut.}} : \mathbf{Art}^r_{\mathbb{C}} \to \mathbf{Set}$$

Let $I \subset A$ be the kernel of p. Let $\pi_i : \hat{A}_p \to A/I^i$ be the maps in the limiting cone on the diagram

$$... \twoheadrightarrow A/I^3 \twoheadrightarrow A/I^2 \twoheadrightarrow A/I \tag{3.3}$$

Applying Ψ to the cone $\{\pi_i\}$ (and picking representatives of the resulting conjugacy classes) gives a set of maps $\Psi\pi_i : \hat{E} \to A/I^i$ forming a cone that commutes up to conjugacy. In fact since every map in (3.3) is a surjection we may inductively pick representatives such that $\Psi\pi_i$ forms a genuinely commuting cone. This cone then factors through some map $f : \hat{E} \to \hat{A}_p$, and by naturality $\Psi = \circ f$. Similarly, since $\hat{E}$ is also a limit of such a diagram, there is a map $g : \hat{A}_p \to \hat{E}$ such that $\Psi^{-1} = \circ g$.

The composition $\circ fg$ is the identity transformation, so applying it to the cone $\{\pi_i\}$ we see that for each i there is an inner automorphism α_i of A/I^i such that $\alpha_i \pi_i fg = \pi_i$. We know π_i is a surjection, so $\pi_i fg$ must also be a surjection, so by a quick diagram chase the maps $\{\alpha_i\}$ commute with the maps in (3.3) and thus lift to an automorphism $\tilde{\alpha}$ of $\hat{A}_p$. Then $\pi_i \tilde{\alpha} fg = \pi_i$ for all i and hence $\tilde{\alpha} fg = \mathbf{1}_{\hat{A}_p}$. Similarly there is an automorphism $\tilde{\epsilon}$ of $\hat{E}$ such that $\tilde{\epsilon} gf = \mathbf{1}_{\hat{E}}$ so f and g must be isomorphisms. □

It has been suggested to the author by Lieven Le Bruyn (and independently by Tom Bridgeland) that this result should generalise to the case that $\mathcal{M}$ is a set of simple (not just 1-dimensional) modules, if we weaken isomorphism to Morita equivalence.

3.4 The graded case

Let $A = A_\bullet$ be an $\mathbb{N}$-graded $\mathbb{C}^r$-algebra with $A_0 = \mathbb{C}^r$. The positively-graded part $A_{>0}$ of A gives an augmentation, so we may consider A to be an object of $\mathbf{Alg}_{\mathbb{C}}^{\star r}$. We also consider A_0 as a set of r one-dimensional A-modules.

Let $\hat{A}$ be the completion of A at $A_{>0}$. The grading on A, viewed as a $\mathbb{C}^*$ action, induces a $\mathbb{C}^*$ action on $\hat{A}$, and we can recover $A \subset \hat{A}$ as the direct sum of the eigenspaces of this action. Geometrically we may think of A_0 as a repulsive fixed point (a source) for a $\mathbb{C}^*$ action on A, so if we take an infinitesimal neighbourhood of the fixed point we can flow it outwards until we see the whole of A.

There is an induced grading on $B(A, A_0)$ which we shall call a *lower grading* to distinguish it from usual dg structure (so $B(A, A_0)$ is now bi-graded) and the differential obviously preserves this lower grading. Dualizing we get the dg-category $\mathcal{H} = \mathrm{Hom}_A(B(A, A_0), A_0)$ (see Section 3.2). Now

$$B(A, A_0)^{-i} = A^{\otimes i+1} \otimes A_0$$

so

$$\mathcal{H}^i = \mathrm{Hom}_{\mathbb{C}}(A^{\otimes i}, \mathbb{C})$$

so the grading on A induces a splitting of $\mathcal{H}$ as a direct product (not a direct sum) of lower graded pieces. The multiplication is degree zero with respect to the lower grading.

Lemma 3.15. *The category $Ext_A(A_0)$ has an induced lower grading, and there is a choice of A_∞-structure on it such that all the products, and the quasi-isomorphism $Ext_A(A_0) \to \mathcal{H}$, preserve the lower grading.*

Proof. $\mathrm{Ext}_A(A_0)$ aquires a lower grading since it is the homology of $\mathcal{H}$ and the differential on $\mathcal{H}$ has lower degree zero. Now we just apply the explicit form of the homological perturbation algorithm (see e.g. [29], [26]), noting that since the differential and multiplication on $\mathcal{H}$ have lower degree zero everything in the algorithm can be chosen to respect the lower grading. □

The following theorem was proven in [28] for the case $A_0 = \mathbb{C}$. There they assume that A is degree-wise finite-dimensional, whereas we assume that $\mathrm{Ext}^1_A(A_0)$ is degree-wise finite-dimensional. It is a consequence of the theorem that the two assumptions are equivalent.

Theorem 3.16. *Choose the A_∞-structure on $Ext_A(A_0)$ to be lower graded as in the previous lemma, and assume $Ext^1_A(A_0)$ is finite dimensional in each lower degree. Then*

$$A = \frac{T(Ext^1_A(A_0)^{\vee g})}{(Ext^2_A(A_0)^{\vee g})}$$

as graded $\mathbb{C}^r$-algebras, where $\vee g$ denotes the lower-graded $\mathbb{C}$-linear dual.

Proof. First we note that the lower-grading ensures that

$$(HMC)^\vee : \mathrm{Ext}^2_A(A_0)^{\vee g} \to T(\mathrm{Ext}^1_A(A_0)^{\vee g})$$

so our statement makes sense. Without the grading it is possible that the image of $HMC^\vee$ only lies in the completed tensor algebra, as in Section 3.3.

Now we would like to use Theorem 3.14, but in Section 3.3 we required that Ext^1 have finite dimension, whereas now we are asking only for lower-degree-wise finite-dimensionality. Examining the proofs however we see that as long as we read $\vee g$ instead of $\vee$ then line (3.1) still holds and hence Theorem 3.14 still holds. Thus if we complete both sides at their positively graded parts then we have an isomorphism. This isomorphism respects the lower grading however, since the isomorphism

$$\mathcal{D}ef_{\mathrm{Ext}(A_0)} \cong \mathcal{D}ef_{\mathcal{H}}$$

respects the lower grading by Lemma 3.15 and the isomorphism

$$\mathcal{D}ef_{\mathcal{H}} \cong \mathcal{D}ef_{A_0}$$

(Lemma 3.7) clearly respects the lower grading. Hence we can identify the original algebras on both sides since they are the direct sums of the graded pieces of their completions. □

4 Superpotential algebras

As discussed in the introduction, it is well known in the physics literature that the algebras arising from quiver gauge theories on Calabi-Yau three-folds can be described by a 'superpotential'. In this section (which probably has some overlap with [14]) we show why this is a consequence of applying our results on deformation theory to the special case of Calabi-Yau 3-folds.

Let X be any complex manifold, and let E be a vector bundle on X with holomorphic structure given by $\bar{\partial}$. Then all other holomorphic structures on E are given by adding to $\bar{\partial}$ an element

$$a \in \mathrm{End}(E) \otimes \mathcal{A}_X^{0,1}$$

satisfying the Maurer-Cartan equation

$$\bar{\partial}(a) + a \wedge a = 0$$

and two such a give isomorphic holomorphic structures if they differ by a gauge transformation. This is dga deformation theory again (see Section 3) - the deformations of $(E, \bar{\partial})$ are governed by the dga

$$\mathrm{RHom}(E, E) \simeq \mathrm{End}(E, E) \otimes \mathcal{A}_X^{0,\bullet}$$

However, rather remarkably in this case the dga gives us the whole moduli space, not just a formal neighbourhood.

When X is a Calabi-Yau three-fold the Maurer-Cartan equation can be written as the derivative of a (locally-defined) function: the *Chern-Simons* function

$$CS(a) := \int_X \mathrm{Tr}(\frac{1}{2}a \wedge \bar{\partial}(a) + \frac{1}{3}a \wedge a \wedge a) \wedge \omega_{vol}$$

where ω_{vol} is a choice of holomorphic volume form on X. Thus heuristically the moduli space is the critical locus of this function. This means that we expect it to be zero-dimensional, and that the number of points in it is the Euler characteristic of the ambient space. Of course this is only heuristic, since the ambient space is the quotient of an infinite-dimensional vector space by an infinite-dimensional group. What one can do however is to construct the moduli space using algebraic geometry and then use the technology of *symmetric obstruction theories*, this leads to the definition of Donaldson-Thomas invariants [32].

If we only care about formal deformations then we have a finite-dimensional version of the above. As in Section 3.1, we can replace the dga $\mathrm{End}(E,E) \otimes \mathcal{A}_X^{0,\bullet}$ by its homology equipped with an A_∞-structure, and the Maurer-Cartan equation by the Homotopy Maurer-Cartan equation. This is still (formally) the critical locus of a function - we just have to add all the higher products into the Chern-Simons function.

Suppose now we have a 3-dimensional Calabi-Yau algebra A instead of a space. The same argument applies, so if M is an A-module then a formal neighbourhood of the moduli space of M is the critical locus of a function. In particular, if M is the module corresponding to a point of A then a formal neighbourhood of M in A is the critical locus of a function. This function is called a *superpotential* for A (at that point). We give a formal definition of superpotentials in Section 4.1, and then make this argument rigorous in Section 4.2.

In fact it is easy to construct global moduli spaces of modules over A, and the construction is entirely finite-dimensional. Furthermore if we have a (polynomial) superpotential for A, then every moduli space of A-modules is globally the critical locus of a function induced by the superpotential. This means that (rather trivially!) the moduli spaces carry symmetric obstruction theories and hence we can define invariants analogous to Donaldson-Thomas invariants. We do this in Section 4.3.

In the examples discussed in the Introduction we have both a space X and an algebra A, and they are derived equivalent, so there is some relationship between moduli spaces of sheaves on X and moduli spaces of A-modules. The physical picture (as we discussed in Section 2.2) is that moving from X to A means that we are moving in the stringy Kähler moduli space, which mathematically probably means some space of stability conditions on the triangulated category $D^b(X)$. This space (or at least a conjectural version of it) has been constructed by Bridgeland [9]. Part of this picture is obvious: passing from $D^b(X)$ to the equivalent $D^b(A)$ is just a change of T-structure, which is part of a Bridgeland stability condition. The remaining data, called the central charge, should roughly correspond on the algebra side to putting a GIT stability condition on the moduli space of A-modules. It should be possible to build a function over the whole of the space of stability conditions such that you can Taylor expand it at the point corresponding to X or the point corresponding to A and get the generating function of the Donaldson-Thomas invariants of X or the Donaldson-Thomas-type invariants of A respectively. This has been carried out in one example by Szendröi [31], and much interesting work is being done (e.g. [17], [12]), and much remains to be done, to properly understand this picture.

4.1 Definition of superpotentials

Let V be a $\mathbb{C}^r$-bimodule. A *superpotential* is simply a sum of cycles in the path algebra of V, i.e. an element

$$W \in \bigoplus_i (TV)_{ii}$$

If the sum is infinite we have a *formal superpotential*, i.e. an element of $\bigoplus_i (\hat{T}V)_{ii}$.

Roughly, we are interested in the (non-commutative) affine scheme described by the critical locus of W, but we have to take a little care with our definition of partial derivative. We choose to make our partial derivative include a sum over cyclic permutations, the alternative (as happened in the introduction) is to insist that W be cyclicly symmetric.

Definition 4.1. Let $x \in V^\vee$. For any t we define the *cyclic partial derivative* in the direction of x as the map

$$\partial_x^\circ : V^{\otimes t} \to V^{\otimes t-1}$$

$$\partial_x^\circ(v_1 \otimes ... \otimes v_t) = \sum_{s=1}^{t} x(v_s) v_{s+1} \otimes ... \otimes v_t \otimes v_1 \otimes ... \otimes v_{s-1}$$

Taking direct sums/products we get maps

$$\partial_x^\circ : TV \to TV$$

and

$$\partial_x^\circ : \hat{T}V \to \hat{T}V$$

Definition 4.2. The algebra *generated by a superpotential* W is $TV/(R)$ where R is the subspace

$$R = \partial^\circ W := \{\partial_x^\circ W \mid x \in V^\vee\} \subset TV$$

If W is a formal superpotential then it generates the algebra $\hat{T}V/(R)$ with the same definition of R.

4.2 3-dimensional Calabi-Yau algebras

As in Section 3.3, we pick an augmented $\mathbb{C}^r$-algebra $(A, p) \in \mathbf{Alg}_{\mathbb{C}}^{\star r}$ and assume that the resulting set of r one-dimensional A-modules $\mathcal{M} = \{\mathcal{O}_{p_i}\}$ has $\mathrm{Ext}^1_A(\mathcal{O}_{p_i}, \mathcal{O}_{p_j})$ finite dimensional $\forall i, j$.

Recall that a graded $\mathbb{C}$-linear category is *Calabi-Yau of dimension d* if it carries a trace map

$$\mathrm{Tr}_M : \mathrm{Hom}(M, M) \to \mathbb{C}$$

of degree $-d$ for each object M, such that the associated pairing

$$\langle\ \rangle_{M,N} : \mathrm{Hom}(M, N) \otimes \mathrm{Hom}(N, M) \to \mathbb{C}$$

is symmetric and non-degenerate. The canonical example is the derived category of a compact smooth Calabi-Yau variety with the Serre duality pairing.

When we say that a derived category $D^b(\mathcal{A})$ of some Abelian category $\mathcal{A}$ is Calabi-Yau we make the further assumption that for any object $M \in \mathcal{A}$ the natural A_∞ structure on $\mathrm{Ext}(M, M)$ may be chosen to be cyclic with respect to the Calabi-Yau pairing. Note that if $D^b(\mathcal{A})$ is the derived category of a smooth compact Calabi-Yau variety then we can satisfy the cyclicity condition: we compute RHom using the Dolbeault resolution, note that this is a cyclic dga, then observe that cyclicity is preserved under the homological perturbation algorithm. For our non-compact examples in Section 5 we will get cyclicity by another construction.

Theorem 4.3. *Suppose that $\langle \mathcal{M} \rangle \subset D^b(A)$ is Calabi-Yau of dimension 3. Then the completion $\hat{A}_p$ of A at the kernel of p is given by a formal superpotential.*

Proof. We learnt this construction from [26]. Consider the formal superpotential

$$W \in \hat{T}(\mathrm{Ext}^1_A(\mathcal{M})^\vee)$$

defined by

$$W(a_1 \otimes ... \otimes a_t) := \frac{1}{t} \langle\, a_1 \mid m_{t-1}(a_2 \otimes ... \otimes a_t) \,\rangle \tag{4.1}$$

where $\{m_i\}$ are the A_∞ products on $\mathrm{Ext}_A(\mathcal{M})$ and $\langle \cdot | \cdot \rangle$ denotes the Calabi-Yau pairing on $\langle \mathcal{M} \rangle$. Then for $x \in \mathrm{Ext}^1_A(\mathcal{M})$ we have

$$\begin{aligned} \partial^\circ_x W(a_1 \otimes ... \otimes a_t) &= \sum_{s=1}^{t+1} W(a_s \otimes ... \otimes a_t \otimes x \otimes a_1 \otimes ... \otimes a_{s-1}) \\ &= \langle\, x \mid m_t(a_1 \otimes ... \otimes a_t) \,\rangle \end{aligned}$$

by cyclicity of m_t. However, if we recall the definition of the HMC function

$$HMC = \bigoplus_{i>0} m_i : T(\mathrm{Ext}^1_A(\mathcal{M})) \to \mathrm{Ext}^2_A(\mathcal{M})$$

we see that we have shown

$$\partial^\circ_x W = HMC^\vee(x)$$

under the identification $\mathrm{Ext}^2_A(\mathcal{M})^\vee = \mathrm{Ext}^1_A(\mathcal{M})$ given by $\langle \cdot | \cdot \rangle$. Hence by Theorem 3.14 the algebra generated by W is $\hat{A}_p$. □

Combining this construction with what we know about graded algebras (Theorem 3.16) we recover a result of Bocklandt.

Theorem 4.4. *[7] Let A be a graded degree-wise finite-dimensional algebra with A_0 semi-simple and with the subcategory of $D^b(A)$ generated by the summands of A_0 Calabi-Yau of dimension 3. Then A is given by a superpotential.*

Note that we do not require that the whole of $D^b(A)$ be Calabi-Yau. Indeed this will not be the case for the algebras coming from non-compact Calabi-Yau varieties that we study in Section 5 below.

Let A be a $\mathbb{C}^r$-algebra with generators V and relations $R \subset TV$. The moduli space of A-modules and the role of GIT stability conditions for it was explained in [20], the following is a summary.

The *dimension vector* of an A-module M is just the vector

$$\mathbf{d} = (d_1, ..., d_r)$$

where d_i is the dimension of the ith summand M_i of M. To give an A-module of dimension $\mathbf{d}$ we have to give a representation of A on the $\mathbb{C}^r$-module

$$\mathbb{C}^{\mathbf{d}} = \bigoplus_i \mathbb{C}^{d_i}$$

To start with forget about the relations, and consider the set of representations of TV on $\mathbb{C}^{\mathbf{d}}$. These form the $\mathbb{C}^r$-bimodule

$$V^\vee \underline{\otimes} \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$$

since an element of this space is precisely a linear map

$$V_{ij} \to \mathrm{Hom}_{\mathbb{C}}(\mathbb{C}^{d_i}, \mathbb{C}^{d_j})$$

for all i, j. Now a relation $r \in R$ is an element of TV, and so it induces, using the composition in $\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$, a polynomial map

$$r : V^\vee \underline{\otimes} \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}}) \to \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$$

The zero locus of r is just those representations that obey the relation r. An A-module structure on $\mathbb{C}^{\mathbf{d}}$ is a representation that obeys all the relations, so the scheme of A-module structures on $\mathbb{C}^{\mathbf{d}}$ is the common zero locus Z of all such $r \in R$.

Two such modules are isomorphic if they differ by a change of basis in $\mathbb{C}^d$, which is an element of the group

$$GL(\mathbf{d}) := \prod_i GL(d_i, \mathbb{C})$$

Hence the moduli stack of A-modules of dimension $\mathbf{d}$ is

$$\mathcal{M}_{A,\mathbf{d}} = [\, Z \,/\, GL(\mathbf{d}) \,]$$

If we pick a character ξ of $GL(\mathbf{d})$ then we can instead take the GIT quotient

$$\mathcal{M}^{\xi}_{A,\mathbf{d}} = Z \,//_{\xi}\, GL(\mathbf{d})$$

A character of $GL(\mathbf{d})$ is necessarily of the form

$$\xi(g) = \prod_{i=1}^{r} \det(g_i)^{\theta_i}$$

for some r-tuple of integers (θ_i). Hence given a character ξ we can define a function

$$\Theta_\xi : K_0(A\text{−}\mathbf{mod}) \to \mathbb{Z}$$

by sending a module of dimension vector $\mathbf{d}$ to the integer

$$\sum_i \theta_i d_i$$

This function is indeed well defined on K_0 because each d_i is additive over short exact sequences.

Definition 4.5. [20] Let $\Theta : K_0(A\text{−}\mathbf{mod}) \to \mathbb{R}$ be an additive function. A module M is Θ-semistable (resp. Θ-stable) if $\Theta(M) = 0$ and every proper submodule $N \subset M$ has $\Theta(N) \geq 0$ (resp. > 0).

Two Θ-semistable modules are called *S-equivalent* if they have the same composition factors in the abelian category of Θ-semistable modules.

Theorem 4.6. *[20] $\mathcal{M}^{\xi}_{A,\mathbf{d}}$ is a coarse moduli space for Θ_ξ-semistable A-modules up to S-equivalence.*

We say $\mathbf{d}$ is *indivisible* if it is not a multiple of another integral vector.

Theorem 4.7. *[20] If $\mathbf{d}$ is indivisible and there are no strictly Θ_ξ-semistable modules of dimension $\mathbf{d}$ then $\mathcal{M}^{\xi}_{A,\mathbf{d}}$ is a fine moduli space for Θ_ξ-stable A-modules.*

Now we consider this construction when A is a superpotential algebra given by a superpotential W. We will show that in this case $\mathcal{M}^{\xi}_{A,\mathbf{d}}$ carries a natural symmetric obstruction theory. For background on symmetric obstruction theories see [4], [3].

Recall that W is required to be a sum of cycles in TV. Using the composition in $\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$ it induces a map

$$W : V^{\vee}\underline{\otimes}\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}}) \to \bigoplus_i \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{d_i})$$

Now we can take traces at each of the r vertices and sum them, getting a scalar polynomial function

$$\tilde{W} : V^{\vee}\underline{\otimes}\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}}) \to \mathbb{C}$$

$$\tilde{W} = \mathrm{Tr}(W)$$

Proposition 4.8. *The (scheme-theoretic) critical locus of $\tilde{W}$ is precisely the (scheme-theoretic) zero locus Z of the relations.*

This is just the statement that the partial derivatives of $\tilde{W}$ are the polynomials on $V^{\vee}\underline{\otimes}\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$ induced by the relations.

Proof. Pick the standard basis of $\mathbb{C}^{\mathbf{d}}$ so that elements of $\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$ are matrices. The heart of the proof is just the fact that if we take two independent matrices M and N and then partially differentiate the function

$$\mathrm{Tr}(MN)$$

holding N fixed, we get the matrix N^T. More generally if $\{M_1, ..., M_l\}$ are independent matrices and we partially differentiate the function

$$\mathrm{Tr}(M_{i_1}...M_{i_t}) \tag{4.2}$$

by varying M_j we get the transpose of the matrix

$$\sum_{i_s=j} M_{i_{s+1}}...M_{i_t}M_{i_1}...M_{i_{s-1}} \tag{4.3}$$

Now pick a basis $\{e_1, .., e_l\}$ of V and a let the dual basis of $V^{\vee}$ be $\{\epsilon_1, ..., \epsilon_l\}$. Then an element $M \in V^{\vee}\underline{\otimes}\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$ is given by a set of matrices $\{M_1, ..., M_l\}$, and evaluating the function $\tilde{W}$ at M gives a linear combination of terms of the form of (4.2). Thus taking partial derivatives of $\tilde{W}$ in all of the M_j directions gives a function

$$V^{\vee}\underline{\otimes}\mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}}) \to \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$$

which is the transpose of the corresponding sum of terms of the form (4.3). If we recall our Definition 4.1 of the cyclic partial derivative we see that this function is the transpose of

$$\partial^{\circ}_{\epsilon_j} W : V^{\vee} \underline{\otimes} \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}}) \to \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$$

Since the set R of relations is spanned by $\partial^{\circ}_{\epsilon_j} W$ the proposition is proved. □

Corollary 4.9. *If there are no strictly ξ-semistable points in $V^{\vee} \underline{\otimes} End_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$ then $\mathcal{M}^{\xi}_{A,\mathbf{d}}$ carries a symmetric obstruction theory.*

Proof. Since there are no strictly semistables the ambient space

$$\mathcal{A}^{\xi}_{\mathbf{d}} := V^{\vee} \underline{\otimes} \mathrm{End}_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}}) /\!/_{\xi}\, GL(\mathbf{d})$$

is smooth. By invariance, $\tilde{W}$ descends to a function on $\mathcal{A}^{\xi}_{\mathbf{d}}$ whose critical locus is $\mathcal{M}^{\xi}_{A,\mathbf{d}}$. The Hessian of this function gives a symmetric obstruction theory. □

Associated to this obstruction theory is a virtual fundamental class $[\mathcal{M}^{\xi}_{A,\mathbf{d}}]^{vir}$.

Definition 4.10. Let $\mathbf{d}$ be indivisible, and assume (i) there are no strictly Θ_{ξ}-semistable modules of dimension $\mathbf{d}$, and (ii) $\mathcal{M}^{\xi}_{A,\mathbf{d}}$ is compact. Then we define the *Donaldson-Thomas-type invariant*

$$\tilde{N}_{A,\mathbf{d},\Theta_{\xi}} = \int [\mathcal{M}^{\xi}_{A,\mathbf{d}}]^{vir}$$

In light of Theorem 4.7 this really is an invariant of the pair (A, Θ_{ξ}), it does not depend on our presentation of A or even on our choice of r idempotents (in fact it should probably be thought of as an invariant of $(A-\mathbf{mod}, \Theta_{\xi})$ but we do not know how to make this precise). Furthermore, by the usual obstruction theory arguments it is invariant under deformations of W that leave $\mathcal{M}^{\xi}_{A,\mathbf{d}}$ compact.

It might appear that $\tilde{N}_{A,\mathbf{d},\Theta_{\xi}}$ depends on the obstruction theory. However, as Behrend has shown in [3], the virtual count under a *symmetric* obstruction theory is in fact an intrinsic invariant equal to a weighted Euler characteristic

$$\chi(X, \nu_X) = \sum_{n \in \mathbb{Z}} n\, \chi(\{\nu_X = n\})$$

where ν is a constructible function defined by Behrend that exists on any DM stack and measures the singularity of the space. Using this we can drop the compactness assumption and define

$$\tilde{N}_{A,\mathbf{d},\Theta_\xi} = \chi(\mathcal{M}^\xi_{A,\mathbf{d}}, \nu)$$

though if $\mathcal{M}^\xi_{A,\mathbf{d}}$ is not compact then we should not expect this to be deformation invariant.

We hope to pursue these invariants further in future work. For the moment however we content ourselves with the following observation.

Lemma 4.11. *Assume there are no strictly ξ-semistable points in $V^\vee \underline{\otimes} End_{\mathbb{C}}(\mathbb{C}^{\mathbf{d}})$. Suppose that $\mathcal{M}^\xi_{A,\mathbf{d}}$ is confined (scheme-theoretically) to a compact submanifold*

$$\mathcal{N} \subset \mathcal{A}^\xi_{\mathbf{d}}$$

Then $\mathcal{M}^\xi_{A,\mathbf{d}}$ is smooth.

Proof. $\tilde{W}$ is a holomorphic function on $\mathcal{A}^\xi_{\mathbf{d}}$ so it is constant along $\mathcal{N}$, so $d\tilde{W}$ restricted to $\mathcal{N}$ is a section of the conormal bundle $N^\vee_{\mathcal{N}}$. The zero locus of this section is $\mathcal{M}^\xi_{A,\mathbf{d}}$. At any point of $\mathcal{M}^\xi_{A,\mathbf{d}}$ the symmetric obstruction theory (i.e. the Hessian of $\tilde{W}$) is an exact sequence

$$0 \to T\mathcal{M}^\xi_{A,\mathbf{d}} \to T\mathcal{N} \oplus N_{\mathcal{N}} \stackrel{Dd\tilde{W}}{\longrightarrow} T^\vee\mathcal{N} \oplus N^\vee_{\mathcal{N}} \to Ob_{\mathcal{M}^\xi_{A,\mathbf{d}}} \to 0$$

By assumption $T\mathcal{M}^\xi_{A,\mathbf{d}} \subset T\mathcal{N}$, so $N_{\mathcal{N}} \stackrel{Dd\tilde{W}}{\longrightarrow} T^\vee\mathcal{N}$ is an injection, and dually $T\mathcal{N} \stackrel{Dd\tilde{W}}{\longrightarrow} N^\vee_{\mathcal{N}}$ must be a surjection. Hence $d\tilde{W}|_{\mathcal{N}}$ is a transverse section of the conormal bundle, and $\mathcal{M}^\xi_{A,\mathbf{d}}$ is smooth. □

5 The derived categories of some local Calabi-Yaus

5.1 Ext algebras on local Calabi-Yaus

For any smooth scheme Z, there is a 'formal' way to extend Z to a Calabi-Yau, namely we use the embedding

$$\iota : Z \hookrightarrow \omega_Z$$

of Z as the zero section in the total space of its canonical bundle. In this section we prove that this procedure is reflected at the level of (A_∞-enriched) derived categories, i.e. that

$$\iota_* D^b(Z) \subset D^b(\omega_Z)$$

is a formal Calabi-Yau enlargement of $D^b(Z)$. There is a slight subtlety here: $D^b(\omega_Z)$ is not actually Calabi-Yau since ω_Z is non-compact, however if Z is compact then objects in $\iota_* D^b(Z)$ are compactly supported, so this subcategory *is* Calabi-Yau. In any case, let us first explain what we mean by 'formal Calabi-Yau enlargement'.

Let $m : V \otimes V \to V$ be any bilinear map on a vector space. We claim that m naturally extends to a bilinear map

$$m^c : (V \oplus V^\vee)^{\otimes 2} \to (V \oplus V^\vee)$$

such that the associated trilinear dual

$$\tilde{m}^c : (V \oplus V^\vee)^{\otimes 3} \to \mathbb{C}$$

is cyclically symmetric. This is straightforward: we let m^c be the direct sum of m with the two maps

$$m_1 : V \otimes V^\vee \to V^\vee$$

and

$$m_2 : V^\vee \otimes V \to V^\vee$$

obtained by dualising and cyclically permuting m (on the fourth direct summand we declare m^c to be zero). We shall call m^c the *cyclic completion* of m. Of course this construction is hardly profound, so no doubt it has been studied and named already, but unfortunately we do not have a reference for it.

In the same way we may cyclicly complete any collection of n-linear maps. Furthermore we claim that this process is sufficently natural that any algebraic structure present in the set of maps will be preserved. For example, if

$$m : V \otimes V \to V$$

is an associative unital product, then one easily checks that m^c is associative and inherits the unit of m. In fact in this case our construction is nothing more than the extension algebra associated to the (V, m)-bimodule $V^\vee$, but the point is that for this particular bimodule the extension algebra is a Frobenius algebra.

The general statement is the following:

Proposition 5.1. *Let $\mathcal{P}$ be a cyclic operad in* **dgVect** *or* **dg**-$\mathbb{C}^r$-**bimod***, and let $U^*\mathcal{P}$ be the underlying classical operad. Let*

$$\Psi : U^*\mathcal{P} \to \mathcal{E}nd_V$$

be an algebra over $U^*\mathcal{P}$ whose underlying dg vector-space is V. Then $V \oplus V^\vee$ is naturally an algebra over $\mathcal{P}$.

For the definitions of cyclic and classical operads see [11].

Proof. We construct a natural transformation of classical operads

$$\Phi : U^*\mathcal{P} \to \mathcal{E}nd_{V\oplus V^\vee}$$

by intertwining the procedure above with the action of the cyclic groups on $\mathcal{P}$ (for the operad of Frobenus algebras this action is trivial). Let the components of Ψ be

$$\Psi_n : U^*\mathcal{P}(n) \to \mathrm{Hom}(V^{\otimes n}, V)$$

and let γ_n be a generator of the cyclic group C_{n+1}. We define

$$\Phi_n : U^*\mathcal{P}(n) \to \mathrm{Hom}((V \oplus V^\vee)^{\otimes n}, V \oplus V^\vee)$$

to be the direct sum of Ψ_n and all the compositions

$$U^*\mathcal{P}(n) \xrightarrow{\Psi_n \gamma_n^{-k}} \mathrm{Hom}(V^{\otimes n}, V) \hookrightarrow \mathrm{Hom}(V^{\otimes k-1} \otimes V^\vee \otimes V^{\otimes n-k}, V^\vee)$$

This procedure commutes with the process of gluing maps together along graphs, so it does define a natural transformation of operads.

Pick a pairing ρ on $(V \oplus V^\vee)^\vee$ compatible with the natural pairing on $V \oplus V^\vee$ (if V has finite-dimensional homology then ρ is unique up to homotopy, so this choice should not worry us). Then $\mathcal{E}nd^\rho_{V\oplus V^\vee}$ is a cyclic operad and there is a natural transformation

$$\mathcal{E}nd_{V\oplus V^\vee} \to U^*\mathcal{E}nd^\rho_{V\oplus V^\vee}$$

The composition of this with Φ is equivariant with respect to the actions of the cyclic groups by construction, so it lifts to a map

$$\mathcal{P} \to \mathcal{E}nd^\rho_{V\oplus V^\vee}$$

□

The lemma remains true if we fix the dimension of the pairing by forming the *n-dimensional cyclic completion*

$$V \oplus V^{\vee}[-n]$$

In particular we may form the n-dimensional cyclic completion of an A_∞-algebra.

Theorem 5.2. *Let Z be a smooth proper scheme of dimension $n-1$, and let*

$$\iota : Z \to \omega$$

be the embedding of Z into its canonical bundle. Then for any $S \in D^b(Z)$, the A_∞-algebra

$$\mathit{Ext}_\omega(\iota_* S, \iota_* S)$$

is the n-dimensional cyclic completion of $\mathit{Ext}_Z(S,S)$.

Proof. Let $\pi : \omega \to Z$ be the projection. We have a tautological exact sequence (which we draw right to left to ensure a happy typographical coincidence later on):

$$0 \leftarrow \iota_* \mathcal{O}_Z \leftarrow \mathcal{O}_\omega \xleftarrow{\tau} \pi^* \omega^\vee \leftarrow 0 \tag{5.1}$$

so there are quasi-isomorphisms of chain complexes

$$\begin{aligned} \mathrm{RHom}_\omega(\iota_* S, \iota_* S) &= \mathrm{RHom}_\omega(\pi^* S \xleftarrow{\tau} \pi^*(S \otimes \omega^\vee), \ \ \iota_* S) & (5.2) \\ &= \mathrm{RHom}_Z(S,S) \oplus \mathrm{RHom}_Z(S \otimes \omega^\vee, S)[-1] & (5.3) \end{aligned}$$

since τ vanishes along the zero section. This latter admits a dga structure if we identify $\mathrm{RHom}_Z(S,S)$ with $\mathrm{RHom}_Z(S \otimes \omega^\vee, S \otimes \omega^\vee)$ and declare the product of two elements in $\mathrm{RHom}(S \otimes \omega^\vee, S)$ to be zero. This dga structure is cyclic with respect to the Serre duality pairing, so it is in fact the n-dimensional cyclic completion of $\mathrm{RHom}_Z(S,S)$.

We claim that under this dga structure the equation (5.3) is actually a dga quasi-isomorphism. To see this take a dga model of the LHS of the form

$$\mathrm{RHom}_\omega(\pi^* S \xleftarrow{\tau} \pi^*(S \otimes \omega^\vee), \ \ \pi^* S \xleftarrow{\tau} \pi^*(S \otimes \omega^\vee))$$

What we mean here is that we should apply RHom termwise. This gives us a two-by-two term complex that looks like the square on the RHS of the following diagram (we suppress the π^*'s to keep the width manageable):

$$\begin{array}{ccccc}
\mathrm{RHom}_Z(S,S) & \dashleftarrow & \mathrm{RHom}_\omega(S,S) & \xleftarrow{\ \tau\ } & \mathrm{RHom}_\omega(S, S\otimes\omega^\vee)[1] \\
\Big\downarrow{\scriptstyle 0} & & \Big\downarrow{\scriptstyle \tau} & & \Big\downarrow{\scriptstyle \tau} \\
\mathrm{RHom}_Z(S\otimes\omega^\vee, S)[-1] & \dashleftarrow & \mathrm{RHom}_\omega(S\otimes\omega^\vee, S)[-1] & \xleftarrow{\ -\tau\ } & \mathrm{RHom}_\omega(S\otimes\omega^\vee, S\otimes\omega^\vee)
\end{array}$$

The dga structure on this two-by-two term complex is easy to describe (by a happy typographical coincidence...): treat an element of the RHS square in the above diagram as a two-by-two matrix, then the product is precisely matrix multiplication. The dashed arrows are the cokernels of each row, we know what these are from the tautological exact sequence (5.1). The two dashed arrows together form the quasi-isomorphism of (5.3).

The dashed arrows do not form a map of dgas. However, consider the map going in the opposite direction, from the LHS two-term complex to the RHS two-by-two term complex, obtained by summing the maps

$$\pi^* : \mathrm{RHom}_Z(S,S) \to \mathrm{RHom}_\omega(S,S)$$

$$\pi^* : \mathrm{RHom}_Z(S,S) \to \mathrm{RHom}_\omega(S\otimes\omega^\vee, S\otimes\omega^\vee)$$

and

$$\pi^* : \mathrm{RHom}_Z(S\otimes\omega^\vee, S)[-1] \to \mathrm{RHom}_\omega(S\otimes\omega^\vee, S)[-1]$$

It is easy to check that this map respects the product structures on each side, so it is a map of dgas. It is also a right inverse to the quasi-isomorphism given by the dashed arrows, hence it is a dga quasi-isomorphism. We conclude that $\mathrm{RHom}_\omega(\iota_*S, \iota_*S)$ is the n-dimensional cyclic completion of $\mathrm{RHom}_Z(S,S)$.

Now we apply homological perturbation, but we may pick our propagator for $\mathrm{RHom}_\omega(\iota_*S, \iota_*S)$ to be the cyclic completion of the propagator we pick for $\mathrm{RHom}_Z(S,S)$, and we have already established that taking cyclic completions is compatible with summing over graphs, so it is compatible with homological perturbation. □

5.2 Completing the algebra of an exceptional collection

Now we specialize to the case discussed in the introduction, where Z is a surface and ω is a local CY three-fold. Furthermore we assume that we have been given a finite full strong exceptional collection of objects $\{T_i\} \subset D^b(Z)$. Known examples include the del Pezzo [23] and ruled surfaces [24].

This leads to the following description of the derived category: if we denote the direct sum of the collection by $T = \oplus_{i=1}^r T_i$ then T is a tilting object, i.e.

$$\mathrm{RHom}(T, -) : D^b(Z) \xrightarrow{\sim} D^b(\mathrm{End}_Z(T)) \tag{5.4}$$

is a triangulated equivalence. The fact that we use only $\mathrm{End}_Z(T)$ instead of $\mathrm{RHom}_Z(T,T)$ is because there are no higher Ext's (the collection is strong), and the fact that this is an equivalence is because the collection generates the whole derived category (it is full). The astute reader will have noted that this functor produces right modules not left modules, so to be consistent with the rest of this paper we should (but won't) replace $\mathrm{End}_Z(T)$ with its opposite algebra.

We make one further assumption on $\{T_i\}$, that for any i, j and $p > 0$ we have

$$\mathrm{Ext}^k_Z(T_i, T_j \otimes \omega_Z^{-p}) = 0$$

for all $k > 0$. The case $p = 0$ is just the meaning of the word 'strong'. We shall call this a *simple* collection (after [8]), it is a generalization of what Bondal and Polishchuk call 'geometric' [6], though their definition is in terms of mutations of the collection. However they show that a geometric collection can only exist on a variety where

$$\mathrm{rk}\ K(Z) = \dim(Z) + 1$$

in which case (as they also show) their definition is equivalent to ours. Simple collections exist on all the del Pezzos and on the ruled surfaces at least up to $\mathbb{F}_2$.

The algebra $A := \mathrm{End}_Z(T)$ has a very simple form. Since T is the direct sum of r objects it is clear that A may be described as the path algebra of a quiver with r nodes plus some relations. Also, by the axioms for a full strong exceptional collection, we may order the T_i so that the Hom's only go in one direction (say increasing i) and so the quiver is directed. Thus we may give A an $\mathbb{N}$-grading by declaring $\mathrm{Hom}_Z(T_i, T_j)$ to be of degree $j - i$. The degree zero piece is just $\oplus_i \mathrm{End}_Z(T_i) = \mathbb{C}^r$. Hence A is of the correct form for Theorem 3.16 to apply, so it is given by a presentation

$$m^\vee : \mathrm{Ext}^2_A(A_0, A_0)^\vee \to T\mathrm{Ext}^1_A(A_0, A_0)^\vee \tag{5.5}$$

for m a (graded) A_∞-structure on $\mathrm{Ext}_A(A_0, A_0)$.

Now form the local Calabi-Yau ω_Z. It is straight-forward [8] to show that the pull-ups of the T_i to ω_Z generate $D^b(\omega_Z)$. The projection formula gives

$$\mathrm{Ext}_\omega(\pi^*T, \pi^*T) = \bigoplus_{p\geq 0} \mathrm{Ext}_Z(T, T \otimes \omega_Z^{-p}) \tag{5.6}$$

so by our condition all the higher self-Ext's of π^*T vanish. Thus π_*T is also a tilting object and

$$D^b(\omega) \cong D^b(\mathrm{End}_\omega(\pi^*T))$$

The question we asked in the introduction was: given a presentation of the form (5.5) for A, can we give a presentation of

$$\tilde{A} := \mathrm{End}_\omega(\pi_*T)?$$

The answer, we claimed, is that $\tilde{A}$ is given by the cyclic completion of the quiver corresponding to the presentation (5.5).

We now fill in the final details in the justification of this answer. Using (5.6) we see we may also give $\tilde{A}$ an $\mathbb{N}$-grading, if we declare

$$\mathrm{Hom}_Z(T_i, T_j \otimes \omega_Z^{-p})$$

to have degree $j - i + rp$. The degree-zero piece is still $\mathbb{C}^r$, so $\tilde{A}$ also admits a presentation of the form

$$m^\vee : \mathrm{Ext}^2_{\tilde{A}}(\tilde{A}_0, \tilde{A}_0)^\vee \to T\mathrm{Ext}^1_{\tilde{A}}(\tilde{A}_0, \tilde{A}_0)^\vee \tag{5.7}$$

Let

$$S = \bigoplus_{i=1}^{r} S_i \in D^b(Z)$$

denote the image of $A_0 \in D^b(A)$ under the derived equivalence (5.4). The S_i form a dual exceptional collection to the T_i. One checks [8] that ι_*S is the object in $D^b(\omega_Z)$ corresponding to $\tilde{A}_0$. We claim (Lemma 5.3 below) that a derived equivalence obtained by tilting is necessarily an A_∞-equivalence, so

$$\mathrm{Ext}_A(A_0, A_0) = \mathrm{Ext}_Z(S, S)$$

and

$$\mathrm{Ext}_{\tilde{A}}(\tilde{A}_0, \tilde{A}_0) = \mathrm{Ext}_\omega(\iota_*S, \iota_*S)$$

as A_∞-algebras. Thus by Theorem 5.2 $\mathrm{Ext}_{\tilde{A}}(\tilde{A}_0, \tilde{A}_0)$ is the 3-dimensional cyclic completion of $\mathrm{Ext}_A(A_0, A_0)$. This means that the presentation (5.7) is really the map

$$\mathrm{Ext}^2_A(A_0, A_0)^\vee \oplus \mathrm{Ext}^1_A(A_0, A_0) \to T\left(\mathrm{Ext}^1_A(A_0, A_0)^\vee \oplus \mathrm{Ext}^2_A(A_0, A_0)\right)$$

given by dualising the cyclic completion of the map m in (5.5). This corresponds precisely to the process we described of cyclicly completing the quiver.

This presentation may be encoded in a superpotential using the construction from Theorem 4.3.

There is one element missing in this story - really we should give a criterion for an arbitrary presentation of A to arise from an A_∞ structure in the manner of (5.5). It seems plausible that the proof in [28] might yield such a criterion. As it is we shall just assume that any reasonable presentation does arise in this way.

Lemma 5.3. *Suppose that the derived category $D^b(\mathcal{C})$ of some abelian category $\mathcal{C}$ has a tilting object $T = \oplus T_i$, and let $A = End_Z(T)$. Then the derived equivalence*

$$\Psi = RHom_{\mathcal{C}}(T, -) : D^b(\mathcal{C}) \xrightarrow{\sim} D^b(A)$$

is in fact a dg (or A_∞) equivalence, i.e. for any $E \in \mathcal{C}$ we have

$$RHom_{\mathcal{C}}(E, E) \simeq RHom_A(\Psi E, \Psi E)$$

as dgas.

Proof. Sending E through Ψ and back produces a resolution $\mathcal{E} \simeq E$ in terms of the T_i. There are no higher Exts between the T_i, so

$$\mathrm{RHom}_{\mathcal{C}}(E, E) = \mathrm{Hom}_{\mathcal{C}}(\mathcal{E}, \mathcal{E})$$

For the same reason $\Psi E = \mathrm{Hom}_{\mathcal{C}}(T, \mathcal{E})$, which is a complex $\mathcal{F}$ of the projective modules $\{A_i := \Psi T_i\}$ that is isomorphic to $\mathcal{E}$. Thus

$$\mathrm{RHom}_A(\Psi E, \Psi E) = \mathrm{Hom}_A(\mathcal{F}, \mathcal{F}) = \mathrm{Hom}_{\mathcal{C}}(\mathcal{E}, \mathcal{E})$$

□

A Appendix - A_∞ Algebras

A.1 Topological motivation

The notion of an A_∞-algebra (or *strongly homotopy associative* algebra) was introduced by Stasheff [30] to answer the following question: what is the algebraic structure present in the loop space of a topological space?

Let X be a (pointed) topological space, and let $\mathcal{L}X$ be the space of loops (at the base-point) in X. Concatenation of loops gives a map

$$\circ : \mathcal{L}X \times \mathcal{L}X \to \mathcal{L}X$$

Everybody knows that if we take π_0 of this map then we get an associative product (in fact a group), but the map itself is not associative, it is only homotopy-associative. We can make this precise: given any three points $a, b, c \in \mathcal{L}X$ there is a standard homotopy between $(a \circ b) \circ c$ and $a \circ (b \circ c)$ given by linearly reparametrizing.

So we have a 1-chain in $\mathcal{L}X$, which we will call $m_3(a, b, c)$, whose boundary measures the failure of associativity of a, b and c. If we extend this process from points to 0-chains we get a map

$$m_3 : C_0(\mathcal{L}X)^{\otimes 3} \to C_1(\mathcal{L}X)$$

such that

$$\partial(m_3(a \otimes b \otimes c)) = (a \circ b) \circ c - a \circ (b \circ c)$$

The RHS of this is sometimes called the 'associator' of a, b, c.

There is no need to restrict ourselves to 0-chains. For consistency let us rename the map $\circ$ as m_2, it is a degree zero bilinear map

$$m_2 : C_*(\mathcal{L}X)^{\otimes 2} \to C_*(\mathcal{L}X)$$

Suppose we have some 1-simplex $\alpha : \Delta_1 \to \mathcal{L}X$, and b, c are points in $\mathcal{L}X$. At each point $s \in \Delta_1$ we have our path $m_3(\alpha_s \otimes b \otimes c)$ between $m_2(m_2(\alpha_s \otimes b) \otimes c)$ and $m_2(\alpha_s \otimes m_2(b \otimes c))$, and we can put them all together to get a 2-chain in $\mathcal{L}X$, which we call $m_3(\alpha \otimes b \otimes c)$ (see Fig. 1).

The boundary of this 2-chain has four pieces, the top and bottom are the asso-

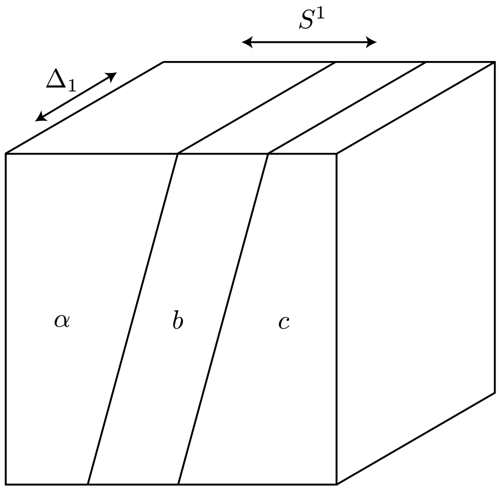

Figure 1: $m_3(\alpha \otimes b \otimes c)$

ciator

$$m_2(m_2(\alpha \otimes b) \otimes c) - m_2(\alpha \otimes m_2(b \otimes c))$$

and the front and back are

$$m_3(\alpha_1 \otimes b \otimes c) - m_3(\alpha_0 \otimes b \otimes c) = m_3(\partial\alpha \otimes b \otimes c)$$

It should now be easy to see that this works for arbitrary simplices α, β, γ in $\mathcal{L}X$, so we have a degree 1 trilinear map

$$m_3 : C_*(\mathcal{L}X)^{\otimes 3} \to C_*(\mathcal{L}X)$$

such that

$$\partial m_3(\alpha \otimes \beta \otimes \gamma) - m_3(\partial(\alpha \otimes \beta \otimes \gamma)) = m_2(m_2(\alpha \otimes \beta) \otimes \gamma) - m_2(\alpha \otimes m_2(\beta \otimes \gamma))$$

The LHS of this is by definition the boundary of the map m_3 in the chain complex

$$\mathrm{Hom}(C_*(\mathcal{L}X)^{\otimes 3}, C_*(\mathcal{L}X))$$

so m_3 is just a 1-chain in this complex linking the point $m_2(\mathbf{1} \otimes m_2)$ to the point $m_2(m_2 \otimes \mathbf{1})$. An alternative way to see this is the following: m_3 is really a 1-cell in the space of continuous maps from $\mathcal{L}X^{\times 3}$ to $\mathcal{L}X$, with end-points $m_2(m_2 \times \mathbf{1})$ and $m_2(\mathbf{1} \times m_2)$, and (at least roughly speaking)

$$C_1(Cts(\mathcal{L}X^{\times 3}, \mathcal{L}X)) = \mathrm{Hom}_1(C_*(\mathcal{L}X)^{\otimes 3}, C_*(\mathcal{L}X))$$

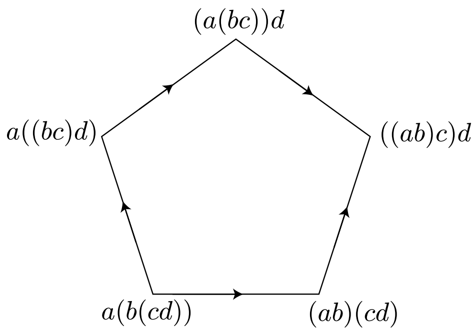

Figure 2: The two-dimensional associahedron

Now we can ask about the ‘associativity’ of m_3. If we take four loops a, b, c and d then there are five different ways to compose them, given by the vertices of the pentagram in Fig. 2.

The edges of the pentagram are the homotopies given by m_3, and every point in them corresponds to some partition of S^1 into four pieces. The space of all such partitions is just an affine subspace of $I^4 \subset \mathbb{R}^4$, so we can pick some 2-cell (independent of a, b, c, d) in this affine space to fill in the interior of the pentagon. This choice of 2-cell will define our next map, m_4. This choice is by no means unique, however our choice of m_3 wasn’t really unique either - we could have picked some other path with the same end-points as our linear one.

Our choice induces, for all a, b, c, d, a 2-cell $m_4(a \otimes b \otimes c \otimes d)$ in $\mathcal{L}X$ whose boundary is the five 1-cells provided by m_3. As before we can extend this to chains of arbitrary degree, so we have a degree 2 quadrilinear map

$$m_4 : C_*(\mathcal{L}X)^{\otimes 4} \to C_*(\mathcal{L}X)$$

Again we could view m_4 as an element in

$$C_2(Cts(\mathcal{L}X^{\times 4}, \mathcal{L}X)) = \mathrm{Hom}_2(C_*(\mathcal{L}X)^{\otimes 3}, C_*(\mathcal{L}X))$$

whose boundary is the expression in m_2 and m_3 given by the boundary of the pentagon. Thus m_3 is ‘associative’ up to the homotopy given by m_4.

What is the rule for m_4? If we take five loops then there are lots of different ways to compose them, corresponding to all the ways we can insert brackets into a product of five variables. These form the vertices of a three-dimensional polyhedron, called the third *Stasheff associahedron* (the first and second are the interval and the pentagon). The edges of the polyhedron are the 1-cells given by

m_3, and the faces are pentagons given by m_4 and rectangles given by products of m_3 1-cells. Now we proceed as before: we sit this pentagon inside the affine space of all partitions of S^1 into five pieces, and pick some 3-cell that it bounds. This induces a 3-cell in $Cts(\mathcal{L}X^{\times 5}, \mathcal{L}X)$, hence a degree 3 map

$$m_5 : C_*(\mathcal{L}X)^{\otimes 5} \to C_*(\mathcal{L}X)$$

whose boundary is the expression in m_2, m_3 and m_4 given by the boundary of the third associahedron. Continuing, we get an infinite sequence of maps

$$m_n : C_*(\mathcal{L}X)^{\otimes n} \to C_*(\mathcal{L}X)$$

of degree $n-2$ such that the boundary of m_n is some expression in the lower products corresponding to the boundary of the $(n-2)$-dimensional associahedron. This structure is the prototype for an A_∞-algebra.

A.2 Formal definition

Let V be a $\mathbb{Z}$-graded vector space. An A_∞*-structure* on V is a sequence of linear maps

$$m_n : V^{\otimes n} \to V$$

of degree $2-n$, for $n \geq 1$, such that for all $n \geq 1$

$$\sum_{n=r+s+t} (-1)^{r+st} m_{r+1+t}(\mathbf{1}^{\otimes r} \otimes m_s \otimes \mathbf{1}^{\otimes t}) = 0 \tag{A.1}$$

The first relation says that $m_1 \circ m_1 = 0$, i.e. m_1 is a differential, so V is a chain-complex (note that we have switched to cohomological grading). The second relation says

$$m_1 \circ m_2 = m_2(m_1 \otimes \mathbf{1} + \mathbf{1} \otimes m_1)$$

which says that m_1 is a derivation with respect to the product m_2. The remaining relations are the relations discussed in the previous section, for example the third one is

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

which says that m_2 is associative up to a homotopy given by m_3.

There are some important special cases. If all the m_n vanish for $n \geq 2$ then we just have a chain-complex. If they all vanish for $n \geq 3$ then we have a differential-graded-algebra (dga), which is a common object e.g. the de Rham cochains on a manifold.

There is a slick way to repackage this data. We shift the grading on V by one to get $W := V[1]$, and form the reduced tensor coalgebra

$$\bar{T}W := \sum_{i \geq 1} W^{\otimes i}$$

The coalgebra structure on this is given by the 'shuffle' coproduct, which takes an element to the sum over all ways in which we can break it in two:

$$\Delta : \bar{T}W \to \bar{T}W \bigotimes \bar{T}W$$

$$\Delta(w_1 \otimes ... \otimes w_n) = \sum_{i=1}^{n-1} (w_1 \otimes ... \otimes w_i) \bigotimes (w_{i+1} \otimes ... \otimes w_n)$$

(and $\Delta(w_1) = 0$).

A sequence of n-linear maps m_n of degree $(2-n)$ on V is the same thing as a single linear map of degree 1

$$m : \bar{T}W \to W$$

Such maps are in bijection with coderivations on $\bar{T}W$, i.e degree 1 maps

$$\tilde{m} : \bar{T}W \to \bar{T}W$$

such that

$$\Delta\tilde{m} = (\tilde{m} \otimes \mathbf{1} + \mathbf{1} \otimes \tilde{m})\Delta$$

The bijection sends m to the coderivation whose component

$$W^{\otimes n} \to W^{\otimes u}$$

is

$$\sum_{r+1+t=u} \mathbf{1}^{\otimes r} \otimes m_{n-r-t} \otimes \mathbf{1}^t$$

Then the A_∞ relations (A.1) are simply the statement that

$$\tilde{m}^2 = 0$$

i.e. $\tilde{m}$ is a coalgebra differential.

Kontsevich and Soibelman [22] have interpreted this structure slightly more geometrically: they view W as the germ of a non-commutative graded manifold, and $\tilde{m}$ as a vector field. This is probably a good point of view to take when using A_∞-algebras to describe deformation spaces.

The easiest way to describe a morphism of A_∞-algebras is via this dual coalgebra description. Let (V, m_*) and (W, l_*) be two A_∞-algebras, and $\tilde{m}$, $\tilde{l}$ be the corresponding differentials on their reduced free tensor coalgebras. Then a *morphism* of A_∞-algebras is a morphism of differential coalgebras

$$f : (\bar{T}(V[1]), \tilde{m}) \to (\bar{T}(W[1]), \tilde{l})$$

We can unpack this: f is determined by its components

$$f_n : V[1]^{\otimes n} \to W[1]$$

so we have a sequence of n-linear maps from V to W of degree $(1-n)$, obeying some relations. The first relation says that f_1 is a chain map, i.e. it commutes with the differentials. The second says that it commutes with the bilinear products up to a homotopy given by f_3. In general f_n provides a homotopy making the lower f_i's compatible with the products in some way.

We say that f is a *quasi-isomorphism* if its first component f_1 is a quasi-isomorphism of chain-complexes, i.e. it induces an isomorphism on homology. This is the correct notion of equivalence between A_∞-algebras.

One can show that every A_∞-algebra is quasi-isomorphic to a dga. Thus working with A_∞-algebras is not essentially different from working with dgas, but it does have various advantages because the notion of a morphism of A_∞-algebras is so much more flexible. For example, every quasi-isomorphism of A_∞-algebras is actually a homotopy equivalence, i.e. it has an inverse up to homotopy. We will see another great advantage in the next section.

A.3 Homological perturbation and Chern-Simons

If we have a dga V, its homology $H(V)$ is a graded algebra. A graded algebra is just a special case of a dga, so one can ask: are V and $H(V)$ quasi-isomorphic?

Of course V and $H(V)$ have the same homology, but this is not enough - we need a morphism of dgas between V and $H(V)$ that induces an isomorphism on homology. In general this is not possible (if it is then V is called *formal*), because we have lost too much information by passing to homology. However, there is a way to put that information back in - we must make $H(V)$ into an A_∞-algebra.

Theorem A.1. *([18] etc.) Let V be a dga (or more generally an A_∞-algebra). Then there is an A_∞-structure on $H(V)$, with $m_1 = 0$ and m_2 the usual product induced from V, such that V and $H(V)$ are quasi-isomorphic. This structure is unique up to A_∞-isomorphism.*

The idea of this theorem goes as follows. Let m denote the product on V and also the induced product on $H(V)$. We can certainly write down a quasi-isomorphism of chain-complexes

$$f_1 : H(V) \to V$$

Now we can ask if it is compatible with the products, i.e. pick $a, b \in H(V)$ and consider the element

$$m(f_1 a, f_1 b) - f_1(m(a, b))$$

This won't be zero in general, but it will be exact. So now pick an element $f_2(a, b)$ such that

$$\partial f_2(a, b) = m(f_1 a, f_1 b) - f_1 m(a, b)$$

Do this for all a, b in such a way that f_2 is bilinear. This is the first relation on an A_∞-morphism. The next relation (in our case) is

$$\partial f_3 = f_2(m \otimes \mathbf{1} + \mathbf{1} \otimes m) + m(f_1 \otimes f_2 + f_2 \otimes f_1)$$

Now we cannot find such an f_3 in general, because the RHS may not be exact. But if we allow $H(V)$ to have an m_3 product, then the LHS of this relation becomes

$$\partial f_3 + f_1 m_3$$

and so, if m_3 is chosen correctly, we may find an f_3. If we continue we find we are able to find our sequence of maps f_i obeying the A_∞ relations, provided that we also introduce the correct higher products m_i for $H(V)$. This procedure is called *homological perturbation.*

Merkulov [29] wrote down an elegant algorithm for computing the higher products on $H(V)$. We start by picking a homotopy equivalence of chain-complexes between $H(V)$ and V, i.e. we pick quasi-isomorphisms of chain-complexes

$$f_1 : H(V) \to V$$

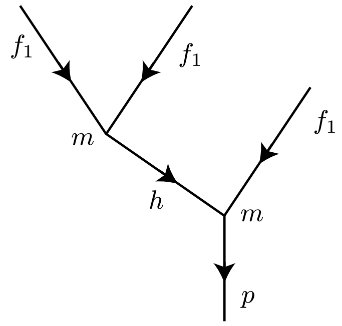

Figure 3: One of the two summands of m_3

and

$$p : V \to H(V)$$

such that $pf_1 = \mathbf{1}$, and we also pick a degree -1 map

$$h : V \to V$$

such that

$$f_1 p - \mathbf{1} = \partial h + h\partial$$

For example, suppose V is the de Rham complex on a compact Riemannian manifold. Then the Hodge decomposition is

$$V = \mathcal{H} \oplus d(H) \oplus d^\dagger(H)$$

where $\mathcal{H}$ are the harmonic forms. Then the inclusion of $\mathcal{H}$ and the projection onto $\mathcal{H}$ provide the quasi-isomorphisms f_1 and p, and the adjoint $d^\dagger$ of the differential provides the homotopy h.

A *rooted trivalent planar tree* is a simply-connected graph embedded in the plane with all vertices having valency three and one external edge chosen as the root. Given such a tree T, with $n+1$ external edges, we can use it to define an n-linear map

$$m_T : H(V) \to H(V)$$

We label all the external edges (except the root) with f_1, all the vertices with m, all the internal edges with h, and the root with p (see Fig. 3 for an example). Composing all these maps in the way indicated by the tree gives m_T. Now we let

$$m_n = \sum_T m_T$$

where the sum is over trees with $n + 1$ external edges. This gives the A_∞-structure on $H(V)$. We can generalize this to the case that V is an A_∞ algebra by allowing T to have vertices of higher valency and labelling these with the higher products of V. There is also a similar description of the components of an A_∞ quasi-isomorphism $H(V) \to V$.

There is a special case where this algorithm has a nice physical description. Suppose that V has a symmetic closed pairing $\langle \ | \ \rangle$ which is non-degenerate on homology, and that the expression $\langle a|m(b \otimes c)\rangle$ is cyclically symmetric. Suppose further that the degree of the pairing is three, i.e. $\langle a|b\rangle = 0$ unless $\deg(a)+\deg(b) = 3$. An example of this is the Poincaré pairing on the de Rham complex of a 3-dimensional compact manifold, or the complex $\mathcal{A}^{0,*}$ of Dolbeault cochains on a Calabi-Yau threefold. Then we can define the *Chern-Simons* function on V:

$$S(a) := \langle a|\partial a + \hbar a^2\rangle$$

We treat V as a space of fields for some theory, ∂ as a BRST operator, and $H(V)$ as the space of physical states. Then the n-point correlator on $H(V)$ is given by the path-integral expression

$$\langle a_1 ... a_n\rangle := \int_{a\in V} e^{S(a)} f_1(a)...f_n(a)$$

We can try to evaluate this by expanding it using Feynman graphs. Each term will have a power of $\hbar$ equal to the number of loops in the graph, so if we now set $\hbar = 0$ we are left with just the trees. This is called taking the *semi-classical limit*, and the resulting correlators are the *tree level correlators*. These Feynman trees are precisely the labelled trees that we used to produce the maps m_T, so we have that

$$\langle a_1 .... a_n\rangle_{tree} = \langle a_1|m_{n-1}(a_2 \otimes ... \otimes a_n)\rangle$$

A.4 Moduli of curves and TCFT

In this section we'll sketch another connection between A_∞-algebras and physics. All the ideas and results of this section come from Costello [11], where they are explained rigourously.

An *open string worldsheet* is a Riemann surface with boundary, with marked points on the boundary. The marked points are distinct, ordered, and partitioned into 'incoming' and 'outgoing' points. We can form a nice compact moduli space of these if we allow surfaces with nodal singularities on their boundary, and only allow stable surfaces (ones with finitely many automorphisms). Let $\mathcal{M}_{i,j}$ be the moduli space of worldsheets with i incoming and j outgoing marked points. These spaces form the morphism spaces of a topological category $\mathcal{M}$, whose objects are the natural numbers, and with the composition map

$$\mathcal{M}_{i,j} \times \mathcal{M}_{j,k} \to \mathcal{M}_{i,k}$$

given by gluing all the outgoing points of one surface to all the incoming points of the next surface. The category has a monoidal product on it given by taking disjoint union of surfaces (and hence cross-product of the moduli spaces).

An (open) *Conformal Field Theory* is supposed to be some kind of monoidal functor $\Psi : \mathcal{M} \to \mathbf{Vect}$. The image of the object $1 \in \mathcal{M}$ is a vector space V which is the space of states of the theory. Then for each worldsheet $\Sigma \in \mathcal{M}_{i,j}$ we have a linear map

$$\Psi(\Sigma) : V^{\otimes i} \to V^{\otimes j}$$

which is the correlator of the theory over that worldsheet. The fact that Ψ is a functor means that the correlators are local, i.e. they can be computed by cutting Σ into pieces. Really we should also include some set of 'branes', and label each piece of the boundary of Σ that lies between successive marked points by some choice of brane.

Unfortunately there is not yet a satisfactory definition along these lines. We can achieve a rigourous (and much simpler) object if we replace each $\mathcal{M}_{i,j}$ with its connected components, so we get the category $H_0(\mathcal{M})$ whose morphisms are topological surfaces with boundary and marked points on the boundary. A monoidal functor $H_0(\mathcal{M}) \to \mathbf{Vect}$ is an (open) 2D *Topological Field Theory*, which turns out to be the same thing as a Frobenius algebra. If we want to be a bit more sophisticated we could try and take account of all the higher homology $H_*(\mathcal{M})$ of our moduli spaces. In fact it is better to work at the the chain level, so we make the following definition:

An (open) *Topological Conformal Field Theory* is a monoidal functor

$$C_*(\mathcal{M}) \to \mathbf{dgVect}$$

($\mathbf{dgVect}$ is the category of differential-graded vector spaces, i.e. chain-complexes). Here C_* is some appropriate chain-complex valued functor that computes homology groups. This means that we have the following:

- A chain-complex $V := \Psi(1)$
- To every surface Σ we have a linear map $\Psi(\Sigma) : V \to V$
- To every bordism Γ between Σ_0 and Σ_1 we have a homotopy $\Psi(\Gamma)$ between $\Psi(\Sigma_0)$ and $\Psi(\Sigma_1)$.
- ...

Now we restrict to a much smaller subcategory of $\mathcal{M}$. Let $\mathcal{D} \subset \mathcal{M}$ be the subcategory of surfaces Σ such that each connected component of Σ is a (possibly nodal) disc with exactly 1 outgoing point. What do the moduli spaces in $\mathcal{D}$ look like? Clearly we only need to understand $\mathcal{D}_{i,1}$, as all the other components are built by taking cross-products of these.

Claim A.2. The moduli space $\mathcal{D}_{i,1}$ is the product of the $(i-2)$-dimensional associahedron with the symmetric group S_i.

Firstly note that for $i < 2$ all the surfaces are unstable so the moduli space is empty. For $i = 2$ the moduli space is any surface is a point, because the Möbius group acts 3-transitively on the disc. However there are two ways to order the incoming points, so the moduli space is S_2.

What happens for $i = 3$? We can use the automorphisms to fix the positions of the outgoing point and the first two incoming points. We can vary the position of the third incoming point, so we have a 1-dimensional moduli space. When the third point reaches either of the two marked points nearest to it, the disc becomes nodal, with two irreducible components each having two marked points. Hence the moduli space of any such surface is an interval (the 1-dimensional associahedron), with the nodal surfaces at the end points. In addition there are $\#S_3$ ways to order the incoming points around the boundary.

For $i = 4$ we can again fix three of the marked points, and vary the remaining two, so the space is 2-dimensional. It has 1-dimensional boundary strata where the surface has a single node, and zero-dimensional boundary strata where the surface has two nodes. There are five such surfaces with two nodes, corresponding to the five ways of putting brackets between the four incoming marked points. We draw the moduli space in Fig. 4, it is isomorphic to the two-dimensional associahedron of Fig. 2.

The claim should now be clear. The zero-dimensional boundary strata of $\mathcal{D}_{i,1}$ are discs with $(i-2)$ nodes, they correspond to all the ways of bracketing the

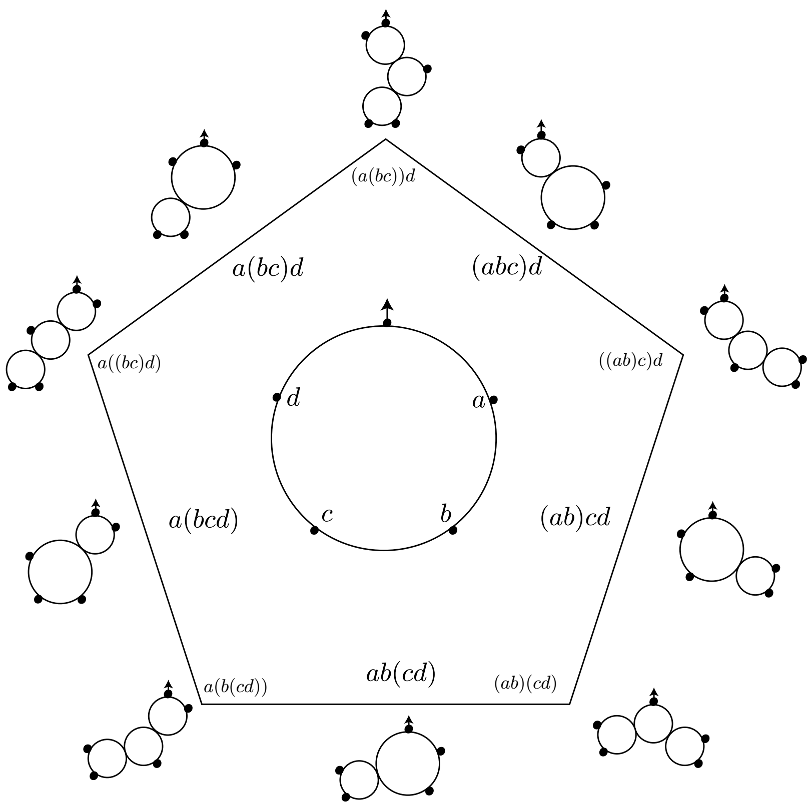

Figure 4: The moduli space $\mathcal{D}_{4,1}$

incoming points. Discs with fewer nodes live on higher-dimensional boundary strata and correspond to inserting fewer brackets, these are the faces of the associahedra. The factor of S_i corresponds to reordering the incoming points.

Now we can describe $C_*(\mathcal{D})$. For each $i \geq 2$ pick a single $(i-2)$-cell Γ_i that forms the interior of one component of $\mathcal{D}_{i,1}$. Then the set of Γ_i (or rather their orbits under S_i) generate the whole of $C_*(\mathcal{D})$ as a monoidal category, because everything in the boundary of $\mathcal{D}_{i,1}$ is produced by composing surfaces with fewer incoming points (and $\mathcal{D}_{i,j}$ is produced by taking products of components of $\mathcal{D}_{*,1}$).

Corollary A.3. *A (degree-reversing) monoidal functor*

$$\Psi : C_*(\mathcal{D}) \to \mathbf{dgVect}$$

is precisely an A_∞-algebra.

The generating $(i-2)$-cell Γ_i gets sent to a degree $(2-i)$ map

$$\Psi(\Gamma_i) : V^{\otimes i} \to V$$

This is the ith A_∞ product. The boundary of Γ_i is the expression in the lower Γ_j's given by the boundary of the $(i-2)$-dimensional associahedron, so the boundary of $\Psi(\Gamma_i)$ must the same expression in some lower products $\Psi(\Gamma_j)$'s. These are precisely the A_∞ relations (A.1).

If we'd allowed branes we'd have ended up with an A_∞-category, which is an A_∞-algebra with more object (i.e. a groupoid compared to a group, or a linear category compared to an algebra).

It turns out that we only need a little more to get a functor defined on the original category $C_*(\mathcal{M})$, ie. an open TCFT. Consider now the subcategory $\hat{\mathcal{D}} \subset \mathcal{M}$ consisting of surfaces each of whose irreducible components is a disc. Obviously $\hat{\mathcal{D}}$ is bigger than $\mathcal{D}$, but it is still just a small (i.e. high codimension) piece of the boundary of $\mathcal{M}$. However, $\mathcal{M}$ can actually be deformation retracted onto $\hat{\mathcal{D}}$. The trick is to use the canonical hyperbolic metric on each surface to flow the boundary inwards. At some point the boundary intersects itself and we get a new nodal singularity - we pull this node apart and start the flow again. Eventually all the irreducible components are discs.

This means that $C_*(\hat{\mathcal{D}})$ is quasi-isomorphic to $C_*(\mathcal{M})$, so our TCFT only needs to be defined on $C_*(\hat{\mathcal{D}})$. Now suppose we had some way of turning incoming marked points into outgoing points. Then we could get all of $\hat{\mathcal{D}}$ by gluing together surfaces in $\mathcal{D}$, and our full TCFT structure would be determined by its value on $C_*(\mathcal{D})$. This happens precisely when our A_∞-algebra V has a homologically non-degenerate pairing on it, i.e. we have a quasi-isomorphism between V and V^*, then we can freely switch between the inputs and outputs of linear maps. We also need that the pairing is cyclically symmetric with respect to the A_∞-structure, i.e. the expression

$$\langle a_0 | m_n(a_1 \otimes .... \otimes a_n) \rangle$$

is always cyclically symetric. An A_∞-algebra with such a cyclic non-degenerate pairing is called *Calabi-Yau*.

Theorem A.4. *A Calabi-Yau A_∞-algebra is precisely an open TCFT (with one brane).*

References

[1] Paul S. Aspinwall and Lukasz M. Fidkowski, *Superpotentials for quiver gauge theories*, J. High Energy Phys. (2006), no. 10, 047, 25 pp. (electronic). MR MR2266656 (2007g:81091)

[2] Paul S. Aspinwall and Sheldon Katz, *Computation of superpotentials for D-branes*, Comm. Math. Phys. **264** (2006), no. 1, 227–253. MR MR2212222 (2007b:81178)

[3] K. Behrend, *Donaldson-Thomas type invariants via microlocal geometry*, 2005, arXiv:math/0507523v2.

[4] K. Behrend and B. Fantechi, *The intrinsic normal cone*, Invent. Math. **128** (1997), no. 1, 45–88. MR MR1437495 (98e:14022)

[5] A. A. Beĭlinson, *Coherent sheaves on $\mathbf{P}^n$ and problems in linear algebra*, Funktsional. Anal. i Prilozhen. **12** (1978), no. 3, 68–69. MR MR509388 (80c:14010b)

[6] Aaron Bergman and Nicholas Proudfoot, *Moduli spaces for D-branes at the tip of a cone*, J. High Energy Phys. (2006), no. 3, 073, 9 pp. (electronic). MR MR2222961

[7] R. Bocklandt, *Graded Calabi Yau algebras of dimension 3*, 2006, arXiv:math/0603558v1.

[8] Tom Bridgeland, *T-structures on some local Calabi-Yau varieties*, J. Algebra **289** (2005), no. 2, 453–483. MR MR2142382 (2006a:14067)

[9] ______, *Spaces of stability conditions*, 2006, arXiv:math/0611510v1.

[10] Tom Bridgeland, Alastair King, and Miles Reid, *The McKay correspondence as an equivalence of derived categories*, J. Amer. Math. Soc. **14** (2001), no. 3, 535–554 (electronic). MR MR1824990 (2002f:14023)

[11] Kevin Costello, *The A_∞ operad and the moduli space of curves*, 2004, arXiv:math/0402015v2.

[12] Alistair Craw and Gregory G. Smith, *Projective toric varieties as fine moduli spaces of quiver representations*, 2006, arXiv:math/0608183v2.

[13] M.R Douglas and G. Moore, *D-branes, quivers, and ALE instantons*, 1996, arXiv:hep-th/9603167v1.

[14] Victor Ginzburg, *Calabi-Yau algebras*, 2006, arXiv:math/0612139v3.

[15] V. K. A. M. Gugenheim and J. D. Stasheff, *On perturbations and A_∞-structures*, Bull. Soc. Math. Belg. Sér. A **38** (1986), 237–246 (1987). MR MR885535 (88m:55023)

[16] Amihay Hanany and Edward Witten, *Type IIB superstrings, BPS monopoles, and three-dimensional gauge dynamics*, Nuclear Phys. B **492** (1997), no. 1-2, 152–190. MR MR1451054 (98h:81096)

[17] Dominic Joyce, *Holomorphic generating functions for invariants counting coherent sheaves on CalabiYau 3-folds*, 2006, arXiv:hep-th/0607039v1.

[18] T. V. Kadeishvili, *The algebraic structure in the homology of an $A(\infty)$-algebra*, Soobshch. Akad. Nauk Gruzin. SSR **108** (1982), no. 2, 249–252 (1983). MR MR720689 (84k:55009)

[19] Bernhard Keller, *A-infinity algebras in representation theory*, Representations of algebra. Vol. I, II, Beijing Norm. Univ. Press, Beijing, 2002, pp. 74–86. MR MR2067371 (2005b:16021)

[20] A. D. King, *Moduli of representations of finite-dimensional algebras*, Quart. J. Math. Oxford Ser. (2) **45** (1994), no. 180, 515–530. MR MR1315461 (96a:16009)

[21] Maxim Kontsevich, *Deformation quantization of Poisson manifolds*, Lett. Math. Phys. **66** (2003), no. 3, 157–216. MR MR2062626 (2005i:53122)

[22] Maxim Kontsevich and Yan Soibelman, *Notes on A-infinity algebras, A-infinity categories and non-commutative geometry. I*, 2006, arXiv:math/0606241v2.

[23] S. A. Kuleshov and D. O. Orlov, *Exceptional sheaves on Del Pezzo surfaces*, Izv. Ross. Akad. Nauk Ser. Mat. **58** (1994), no. 3, 53–87. MR MR1286839 (95g:14048)

[24] A. B. Kvichansky and D. Yu. Nogin, *Exceptional collections on ruled surfaces*, Helices and vector bundles, London Math. Soc. Lecture Note Ser., vol. 148, Cambridge Univ. Press, Cambridge, 1990, pp. 97–103. MR MR1074785 (91m:14068)

[25] O. A. Laudal, *Noncommutative deformations of modules*, Homology Homotopy Appl. **4** (2002), no. 2, part 2, 357–396 (electronic), The Roos Festschrift volume, 2. MR MR1918517 (2003e:16005)

[26] Calin I. Lazaroiu, *String field theory and brane superpotentials*, J. High Energy Phys. (2001), no. 10, Paper 18, 40. MR MR1877563 (2003h:81179)

[27] Kenji Lefèvre-Hasegawa, *Sur les A_∞-catégories*, 2003, PhD thesis.

[28] D.-M. Lu, J.H. Palmieri, Q.-S. Wu, and J.J. Zhang, *A-infinity structure on Ext-algebras*, 2006, arXiv:math/0606144v1.

[29] S. A. Merkulov, *Strong homotopy algebras of a Kähler manifold*, Internat. Math. Res. Notices (1999), no. 3, 153–164. MR MR1672242 (2000h:32026)

[30] James Dillon Stasheff, *Homotopy associativity of H-spaces. I, II*, Trans. Amer. Math. Soc. 108 (1963), 275-292; ibid. **108** (1963), 293–312. MR MR0158400 (28 #1623)

[31] B. Szendroi, *Non-commutative Donaldson-Thomas theory and the conifold*, (2007), arXiv:0705.3419v2.

[32] R. P. Thomas, *A holomorphic Casson invariant for Calabi-Yau 3-folds, and bundles on $K3$ fibrations*, J. Differential Geom. **54** (2000), no. 2, 367–438. MR MR1818182 (2002b:14049)

[33] Michel van den Bergh, *Non-commutative crepant resolutions*, The legacy of Niels Henrik Abel, Springer, Berlin, 2004, pp. 749–770. MR MR2077594 (2005e:14002)