Computations in Quantum K-theory

Andrei Okounkov
Lecture notes by Tony Feng

Contents

1 Hilbert schemes

1.1 Introduction

The lectures so far have been concerned with using geometry, such as the geometry of moduli spaces, to solve problems in representation theory. My lectures will be about the other direction: using representation theory to answer questions about the geometry of moduli spaces.

Specifically, the moduli spaces we'll be interested in parametrize *curves in* 3*-folds*: if X is a three-fold, then we consider $\mathcal{M} = \mathrm{Hilb}(X, \text{curves})$. Here, "curve" is broadly interpreted: we allow reducible curves, possibly with components that are non-reduced, and additional points.

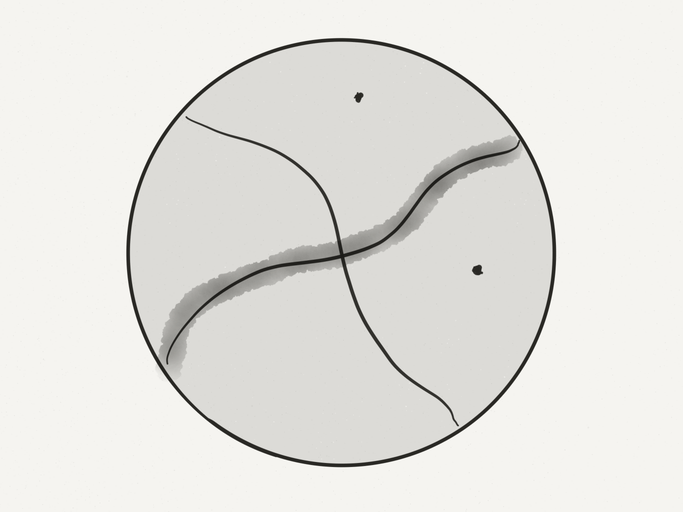

To cut down the moduli space to something reasonable, we stratify by fixed parameters describing the curve. The fixed parameters are

- some $\beta \subset H_2(X, \mathbb{Z})$, which corresponds to the linear term of the Hilbert polynomial, and
- the "constant term" $n = \chi(\mathcal{O}_C)$.

For example, in the diagram the non-reduced "doubled curve" contributes *twice* the homology class of its underlying set.

Example 1.1. The "too good to be true" example is $\mathcal{M} = \mathrm{Hilb}(\mathbb{P}^3, [L], 1)$. The only subschemes parametrized by $\mathcal{M}$ are the honest lines, since all the Euler characteristic is "spent" on the line, so that there is no room left for points. So $\mathcal{M} = \mathrm{Gr}(2, 4)$, which is a homogeneous space for GL(4). However, we'll almost always want to consider the action of the maximal torus T instead.

The moduli space $\mathcal{M} \times X$ contains a *universal subscheme* $\mathcal{Y}$. At the level of sets, every k-point in $\mathcal{M}$ corresponds to a subscheme $Y \subset X$, and the fiber of $\mathcal{Y} \to \mathcal{M}$ is precisely $Y \subset X$.

The enumerative geometry of X is equivalent to computing within the K-theory of $\mathcal{M}$.

Example 1.2. In $\mathrm{Gr}(2,4)$ we have Schubert cycles, which can be described in terms of the universal line.

- There is a divisor consisting of all lines meeting a fixed line, since to meet a line is a codimension one condition.

  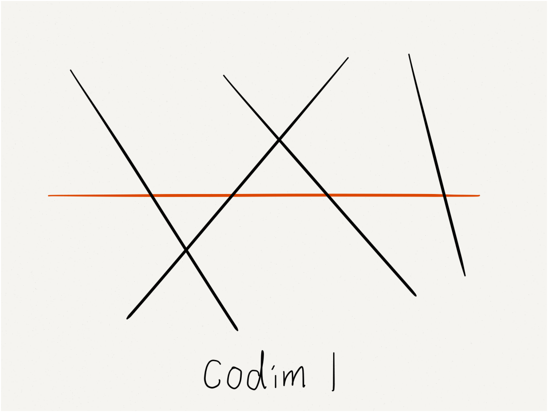

  There are two classes of codimension 2: all lines in a fixed plane, or all lines through a given point.

  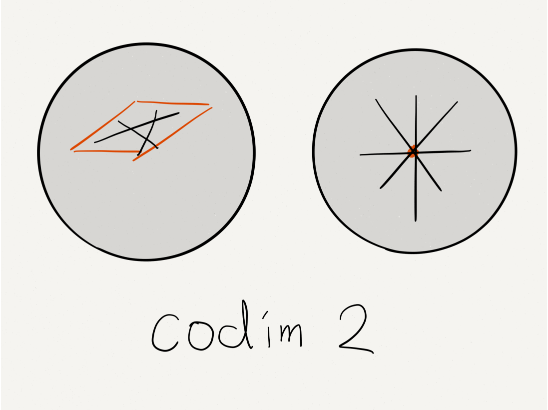

  Finally, there is a class of codimension 3, namely the pencil of lines in a given plane through a given point.

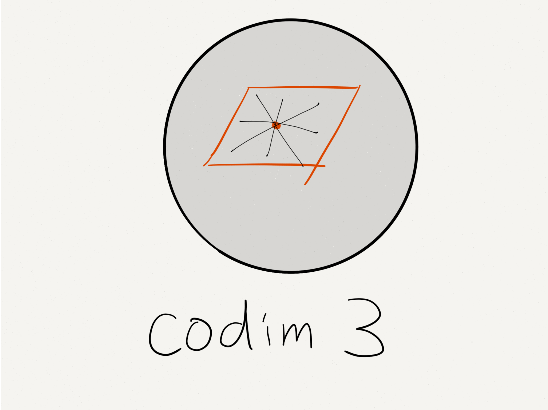

We have $\dim \mathcal{M} = 4$, which is a special case of Bezrukavnikov's formula (haha).

What is the significance of this 4? There is a dimension formula, which isn't really a dimension formula but an *expected dimension* formula:

$$\boxed{\text{expected dimension } (\mathcal{M}) = \beta \cdot c_1(X).} \tag{1}$$

♠♠♠ TONY: [Here $c_1(X) = c_1(TX)$] This follows from a standard deformation theory argument.

The special thing about $\dim x = 3$ is that this formula for the expected dimension is *independent* of the genus of the curve (i.e. the n parameter).

Example 1.3. For $\beta = 0$, we are considering $\mathrm{Hilb}(X, n \text{ points})$. The first disappointment of our formula (1) is that it says the expected dimension is 0. The actual dimension is certainly at least $3n$, because we have a subset of $\mathcal{M}$ parametrize n distinct points, each with 3 dimensions of freedom. But in fact, you can even show that $\dim \mathcal{M} \geq \text{const} \cdot n^{4/3}$. Why? The moduli space contains components which are *not* in the closure of collections of n distinct points.

Therefore, $K(\mathcal{M})$ is way too big. So we have to focus on some small part of it. Before we do this, we ask: *why* are we getting expected dimension 0?

1.2 Re-interpretation of Hilbert schemes

Let's think more generally. If $X = \operatorname{Spec} R$, then what is the "Hilbert scheme of points" of X? Well, if we have a length n subscheme then we get a surjection

$$\mathcal{O}_X \twoheadrightarrow \mathcal{O}_{\text{subscheme}}$$

such that $\dim \mathcal{O}_{\text{subscheme}} = n$. This determines the subscheme, so a point of length n is just an n-dimensional R-module *together* with a surjection from R.

Identify the n-dimensional module as $\mathbb{C}^n$. Suppose $R = \mathbb{C}\langle x_1, x_2, \ldots\rangle$/relations. To specify an n-dimensional R-module is to interpret $x_1, x_2, \ldots$ as $n \times n$ matrices, satisfying the relations. Then a surjection from R is equivalent to the data of a vector $v \in \mathbb{C}^n$ (the image of $1 \in R$) such that polynomials in the x_i applied to v generate $\mathbb{C}^n$. Finally, we have to mod out by the action of $\mathrm{GL}(n)$ on everything because the choice of basis was ambiguous.

Example 1.4. If $X = \mathbb{C}^3$, i.e. $R = \mathbb{C}[X_1, X_2, X_3]$ then we are specifying three matrices plus a cyclic vector such that $[X_i, X_j] = 0$, all modulo $\mathrm{GL}(n)$.

It is instructive to think about taking what happens if we don't impose even these commutativity relations. Note that everything makes sense even if R is *not* commutative. In particular, if you have an algebra with three generators, then it is a quotient of the *free* algebra on three generators, so we getting an embedding $\mathcal{M} \hookrightarrow \widetilde{\mathcal{M}} = \mathrm{Hilb}(\mathrm{Free}_3, n)$. The latter space parametrizes the data of 3 matrices plus a cyclic vector, modulo $\mathrm{GL}(n)$.

Exercise 1.5. Show that $\widetilde{\mathcal{M}}$ is a smooth algebraic variety of dimension $2n^2 + n$. ♠♠♠ TONY: [TODO: smoothness]

Now let's go back to thinking about the equations $[X_i, X_j] = 0$. The conditions that these are all 0 is equivalent to a *critical function equation* $\partial\varphi = 0$ where

$$\varphi = \mathrm{tr}(X_1X_2X_3 - X_1X_3X_2).$$

This is a well-defined function on $\widetilde{\mathcal{M}}$. It is easy to see that if you take the trace of this with respect to X_i, you get the commutator for the other two variables.

This situation was too good to be true in general. There was something very special about our variety. In general, what is true is that "locally you are cut out by critical points of a section of a vector bundle."

Now it is clear that since we've written the scheme in terms of variables and the critical points of a function, then the expected dimension *should* be zero, since we have at least as many equations as variables (one for each partial).

With that puzzle resolved, let's consider

$$\mathcal{O}_{\widetilde{M}} \to \mathcal{O}_{\mathcal{M}} \to 0.$$

Now, $\partial\varphi$ can be viewed as as section of $T^*\widetilde{\mathcal{M}}$. Dualizing, i.e. pairing with $\partial\varphi$, gives

$$T\widetilde{\mathcal{M}} \xrightarrow{d\varphi} \mathcal{O}_{\widetilde{M}} \to \mathcal{O}_{\mathcal{M}} \to 0.$$

This isn't necessarily exact because $\partial\varphi$ is not regular. One might imagine an extension of the complex into a resolution:

$$\ldots \wedge^2 T\widetilde{\mathcal{M}} \xrightarrow{d\varphi} T\widetilde{\mathcal{M}} \xrightarrow{d\varphi} \mathcal{O}_{\widetilde{M}} \to \mathcal{O}_{\mathcal{M}} \to 0.$$

Let's be careful, though. Being representation theorists, we would like everything here to be equivariant. However, the function φ is not invariant. Indeed, if we take $T \subset \mathrm{GL}(3)$

acting on $X = \mathbb{C}^3$, then the function transforms non-trivially under T. How precisely does it transform? The first thing in the representation theory is to clear up the difference between a representation and its dual. Here, if T has weight $\kappa = t^{(n_1,n_2,n_3)}$ acting on $\mathbb{C}^3$, then the weight of T on φ is $\kappa^{-1} = t^{(-n_1,-n_2,-n_3)}$ because functions transform *inversely* to variables.

So the correct guess for the complex (making things equivariant) is really

$$\ldots \to \wedge^2 T\widetilde{\mathcal{M}} \otimes \kappa^{-2} \xrightarrow{\kappa\otimes d\varphi} T\widetilde{\mathcal{M}} \otimes \kappa^{-1} \xrightarrow{\kappa\otimes d\varphi} \mathcal{O}_{\widetilde{M}} \to \mathcal{O}_{\mathcal{M}} \to 0. \tag{2}$$

We can interpret the cutoff

$$\ldots \to \wedge^2 T\widetilde{\mathcal{M}} \otimes \kappa^{-2} \xrightarrow{\kappa\otimes d\varphi} T\widetilde{\mathcal{M}} \otimes \kappa^{-1} \xrightarrow{\kappa\otimes d\varphi} \mathcal{O}_{\widetilde{M}} \to 0$$

as a sheaf of differential graded algebras on $\mathcal{M}$. By the Leibniz rule, the 0th cohomology, which is $\mathcal{O}_M$, acts on every cohomology group. Thus, we can view each cohomology group as a sheaf on Spec R.

Now comes the key definition. The idea is that if you want a good object to work with, then you should take the *whole complex*.

Definition 1.6. We define $\mathcal{O}^{\mathrm{vir}}_{\mathcal{M}}$ to be the complex $\bigwedge^\bullet(T\widetilde{M} \otimes \kappa^{-1})$, where (by definition) for any vector bundle V,

$$\wedge^\bullet V = \sum_i (-1)^i \wedge^i V.$$

Remark 1.7. This object is really only an approximation to the *true* object, which we will consider later.

The question that we want to address is:

Question. How many points are there in $\mathbb{C}^3$?

What does this question even mean? Our interpretation is that it is the same as understanding the function

$$\sum_n z^n \chi_{\mathrm{GL}(3)}(\mathcal{M}_n, \mathcal{O}^{\mathrm{vir}})$$

where we mean the GL(3)-equivariant Euler characteristic.

So how can we study this? The topic of my lectures is computation. Here, the computational tool is localization.

1.3 Localization formula

Localization is the subject of the first problem session. What is it about? Suppose we have a structure M, say a smooth algebraic variety, with an action of some torus T. (In our setting, T is the maximal torus of G.) Assume that T has a finite set of fixed points $\{x_i\}$. (This will be satisfied in our situation.)

Suppose that we want to compute $\chi(\mathcal{O}_M)$. This means that we want to compute the trace of how T acts on functions on M.

Imagine that your group were finite. Then the the action on functions is a "combinatorial representation," with the matrices of the action being essentially permutation matrices. Then the only things contributing to the trace are the "fixed points."

It is pretty clear here too that only the fixed points has any chance of computing the trace of the T-action on functions. So we should have something like

$$\chi(\mathcal{O}_M) = \sum_{x_i} \text{functions on formal neighborhood of } x_i$$

Now, what are functions on formal neighborhood of x_i? The "functions" are $\operatorname{Sym}^\bullet T_x^* M$. As a T-module, we get a weight decomposition

$$T_x M = \bigoplus_{w \text{ weight } \neq 1} w.$$

(Sometimes what we denote by w is written as t^w. This is just a question of notation.) If this is the space, then of course $T_x^* M = \bigoplus_w w^{-1}$, so

$$\operatorname{Sym}^\bullet T_x^* M = \prod_w \frac{1}{1 - w^{-1}}. \tag{3}$$

The most basic form of the localization formula is then

$$\chi(\mathcal{O}_M) = \sum_{x_i} \operatorname{Sym}^\bullet T_x^* M = \sum_{x_i} \prod_w \frac{1}{1 - w^{-1}}. \tag{4}$$

Example 1.8. Let's apply this to compute the Euler characteristic of $\mathbb{P}^1$. If the tangent space at 0 has weight t, then the tangent space at ∞ is scaled by t^{-1} (since the coordinate is the inverse of that at 0).

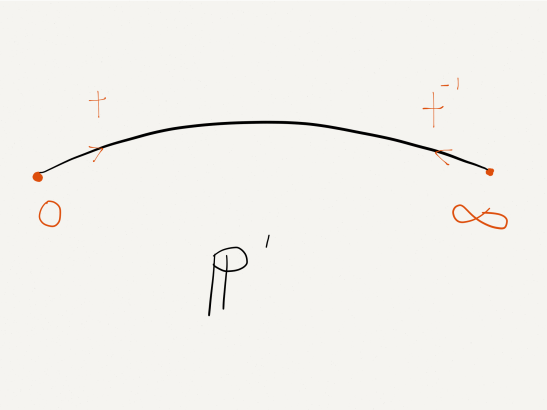

Then the weight on the *cotangent* spaces are inverted, so by (4) we have

$$\chi(\mathcal{O}_{\mathbb{P}^1}) = \frac{1}{1 - t^{-1}} + \frac{1}{1 - t} = 1.$$

Generalization to vector bundles. We can now replace the structure sheaf by something that looks at least locally like a structure sheaf, namely a vector bundle. Then what is the analogue of the localization formula? Locally the vector bundle looks trivial, so we should still have a contribution like the one from before, but there will also be an additional piece.

$$\chi(\mathcal{M}, V) = \sum_{x_i} (?) \operatorname{Sym}^\bullet T_x^* M$$

What should this additional twist be? Well, think about what's going on locally near x.

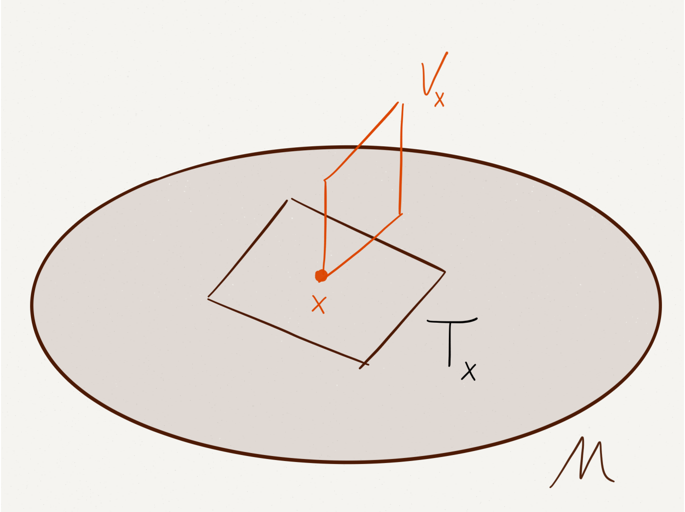

Locally we have the T action on V_x, so it's clear that we should throw in V_x:

$$\chi(\mathcal{M}, V) = \sum_{x_i} V_x \cdot o \operatorname{Sym}^\bullet T_x^* M$$

Exercise 1.9. Use this to compute $\chi(O(n))$. *Warning:* $\operatorname{Aut}(\mathbb{P}^1)$ doesn't act on $O(1)$. To get an action, you should take the maximal torus

$$\left\{ \begin{pmatrix} t_1 & \\ & t_2 \end{pmatrix} \right\} \subset \mathrm{GL}(2)$$

Then the weights are

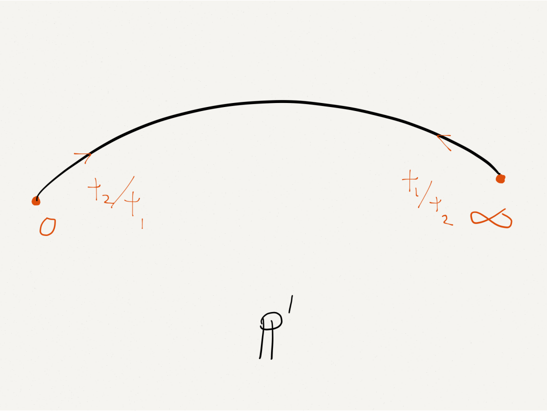

Localization formula. The restriction V_x of V to x is just the pullback in K-theory. So the localization formula may be written as

$$\boxed{\chi(M,\mathcal{F}) = \sum_{x_i} i_{x_i}^*\mathcal{F} \otimes \mathrm{Sym}^\bullet\, T_x^* M.}$$

2 The Hilbert scheme of points in $\mathbb{C}^3$

2.1 First attempt

Let's go back to the embedding $\mathcal{M} \hookrightarrow \widetilde{\mathcal{M}}$. Inside $\mathcal{M}$ we have an embedding of the fixed points $\mathcal{M}^T$, and an obvious diagram

$$\begin{tikzcd}
\mathcal{M} \arrow[r, hook] & \widetilde{\mathcal{M}} \\
\mathcal{M}^T \arrow[u, hook] \arrow[r, hook] & \widetilde{\mathcal{M}}^T \arrow[u, hook]
\end{tikzcd}$$

Now, even if $\mathcal{M}$ is not smooth, we can use localization on the ambient smooth space. We just have to figure out what $\mathcal{M}^T \hookrightarrow \widetilde{\mathcal{M}}^T$ is.

The space $\mathcal{M}_n$ parametrizes subscheme of $\mathbb{C}^3$ of length n, and you can think of this as the same as describing the ideals of $\mathbb{C}[x_1, x_2, x_3]$ of codimension n. The action of T is by scaling the variables. How could an ideal I be fixed by scaling the variables? I claim that this is only possible if I is a monomial ideal, i.e. generated by monomials. This is because the x^d are the eigenfunctions for T, with distinct weights, e.g. $x_1^3 x_2$ has weight $t_1^{-3} t_2^{-1}$. Any monomial is uniquely determined by the weights, and an invariant ideal must be generated by eigenfunctions.

In two dimensions, any such ideal can be specified by choosing some monomials from the table

$$\begin{array}{ccccc}
1 & x_1 & x_1^2 & x_1^3 & \cdots \\
x_2 & x_1x_2 & x_1^2x_2 & x_1^3x_2 & \cdots \\
x_2^2 & x_1x_2^2 & x_1^2x_2^2 & x_1^3x_2^2 & \cdots \\
x_2^3 & x_1x_2^3 & x_1^2x_2^3 & x_1^3x_2^3 & \cdots \\
\vdots & \vdots & \vdots & \vdots & \ddots
\end{array}$$

The condition of being an ideal says that if you choose a monomial to be in I, then you must also have everything below and to the right of it.

So in two dimensions, the ideals are in bijection with two-dimensional partitions; in three dimensions the ideals are in bijection with three-dimensional partitions, etc.

Now, let's go back to the complex $(\bigwedge^\bullet (T\widetilde{\mathcal{M}} \otimes \kappa^{-1}), d\varphi \otimes \kappa)$. We can take this complex and restrict it to $\widetilde{\mathcal{M}}^T$, and we claim that its *restriction* to $\widetilde{\mathcal{M}}^T$ is precisely a resolution of the structure sheaf of $\mathcal{M}^T$.

In general, if you have a virtual class and a fixed locus, you can decompose the virtual class into a "fixed" and "moving" part. The fixed part will give you a virtual class in the fixed locus, which is what we've gotten here. Then you can use localization formula on the ambient space to compute

$$\sum_n z^n \chi(\mathcal{M}_n, \mathcal{O}^{\mathrm{vir}}) = \sum_n z^n \chi(\widetilde{\mathcal{M}}_n, i_* \mathcal{O}^{\mathrm{vir}}).$$

Now, by the claim $i_* \mathcal{O}^{\mathrm{vir}} = \bigwedge^\bullet (T\widetilde{\mathcal{M}} \otimes \kappa^{-1})$, so we can write a localization formula:

$$\chi(\mathcal{M}_n, \mathcal{O}^{\mathrm{vir}}) = \sum_{\text{3D partitions } \pi} z^{|\pi|} \frac{\text{contribution from } \bigwedge^\bullet T\widetilde{\mathcal{M}}}{\text{contribution from } \mathrm{Sym}^\bullet T^* \widetilde{\mathcal{M}}}$$

The convenient way to write this contribution is in terms of the *virtual tangent space*. Define it at the partition $\pi \in \mathcal{M}$ to be

$$T_\pi^{\mathrm{vir}} = \text{Deformations} - \text{Obstructions}.$$

The obstruction class is $d\varphi \otimes \kappa \in \mathrm{Obs} = T^* \widetilde{\mathcal{M}} \otimes \kappa$ (think of this as the equations imposed on deformations by φ). So that gives

$$\begin{aligned} T_\pi^{\mathrm{vir}} &= T_\pi \widetilde{\mathcal{M}} - \kappa \otimes T_\pi^* \widetilde{\mathcal{M}} \\ &= \bigoplus \pm w_i \end{aligned}$$

where the torus weights w_i can be computed from knowledge of the action T_π and T_π^*. For every such w_i, you stick a factor into the contribution to the localization formula. What factor? It should be

$$\sum \pm w_i := \bigoplus \pm w_i \mapsto \prod (1 - w_i^{-1})^{\mp} 1.$$

Let's think about this. If you have a tangent direction, corresponding to a positive w_i, then you get something in the denominator, which is as expected. If you have an obstruction, then you get something in the numerator.

Now we need to compute $T_\pi \widetilde{\mathcal{M}} - \kappa \otimes T_\pi^* \widetilde{\mathcal{M}}$. Since $\widetilde{\mathcal{M}}$, there is nothing stopping you from computing it. You can then put everything into the generating function,and what do you get? *Garbage* - there doesn't seem to be anything nice about the function (at least that I can tell). That seems to indicate that we haven't really been working with the "right" objects here. Next time we'll refine our analysis.

2.2 The refined virtual class

Let's review our setup. We had realized $\mathcal{M} = \mathcal{M}_n = \mathrm{Hilb}(\mathbb{C}^3, n) \hookrightarrow \widetilde{\mathcal{M}}$ as the zero section of $s = d\varphi \otimes \kappa$, and $\mathrm{Obs} = T^*\widetilde{\mathcal{M}} \otimes \kappa$.

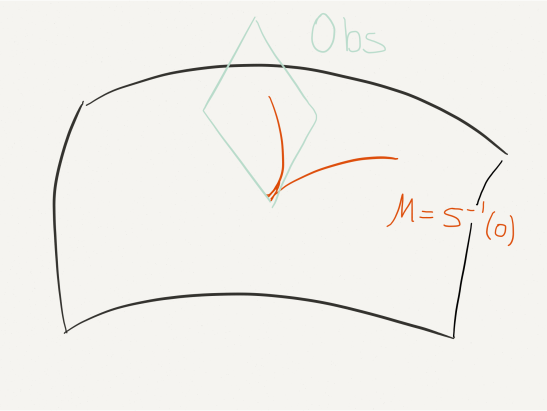

We defined a fundamental class

$$\mathcal{O}^{\mathrm{vir}} = [\ldots \xrightarrow{s} \wedge^2 \mathrm{Obs}^\vee \xrightarrow{s} \mathrm{Obs}^\vee \xrightarrow{s} \mathcal{O}_{\widetilde{\mathcal{M}}}]$$

and then considered the generating function for its Euler characteristic. We computed this by pushing $\mathcal{O}^{\mathrm{vir}}$ forward to $\widetilde{\mathcal{M}}$, and describing the result in terms of the "virtual tangent space"

$$\begin{aligned} T_x^{\mathrm{vir}} &= \underbrace{T_x\widetilde{\mathcal{M}}}_{\mathrm{Def}} - \mathrm{Obs}_x \\ &= \sum a_i - \sum b_i \end{aligned}$$

where the a_i are the weights of the deformations not counted by obstructions, and the b_i are the opposite.

In localized K-theory at a point $x \in \mathcal{M}$, we obtain

$$\mathcal{O}_x^{\mathrm{vir}} = \mathcal{O}_x \prod \frac{1 - b_i^{-1}}{1 - a_i^{-1}}. \tag{5}$$

In the exercises, you are asked to compute for $\mathrm{Hilb}(\mathbb{C}^3, n)$ and x a 3D partition. You'll find that

$$\sum_n z^n \chi(\mathcal{M}, \mathcal{O}^{\mathrm{vir}}) = \sum_{\text{3D partitions } \pi} z^{|\pi|} \prod \frac{\cdots}{\cdots}.$$

This function doesn't seem to have any nice properties, so we seek to refine our approach.

Motivation. Let M be a complex Kähler manifold. Then we have a Dolbeault resolution

$$0 \to \mathbb{C} \to \mathcal{O}_M \xrightarrow{\bar\partial} \Omega^{0,1} \xrightarrow{\bar\partial} \Omega^{0,2} \xrightarrow{\bar\partial} \ldots \to \Omega^{0,\dim M} \to 0.$$

When you look in *physics* textbooks, you never see the Dolbeault resolution. Instead, you see Dirac's resolution instead. That is the resolution obtained by tensoring the Dolbeault resolution with a *square root* of the canonical bundle (assuming a spin structure), and using the symmetrized operator $\overline{\partial} + \overline{\partial}^*$:

$$0 \to \mathbb{C} \to O_M \otimes K_M^{1/2} \xrightarrow{\overline{\partial}+\overline{\partial}^*} \Omega^{0,1} \otimes K_M^{1/2} \xrightarrow{\overline{\partial}+\overline{\partial}^*} \Omega^{0,2} \otimes K_M^{1/2} \xrightarrow{\overline{\partial}+\overline{\partial}^*} \ldots \to \Omega^{0,\dim M} \otimes K_M^{1/2} \to 0.$$

This complex is better because it is more symmetric.

So this suggests that we should replace O^{vir} with $\widehat{O}^{\mathrm{vir}} \otimes (K^{\mathrm{vir}})^{1/2}$, where K^{vir} is some appropriate "virtual canonical bundle." What should it be, precisely? Well, we have a virtual tangent bundle, so K^{vir} should be the top exterior power of the dual of the virtual tangent bundle, so

$$K^{\mathrm{vir}} = \frac{\det \mathrm{Obs}}{\det \mathrm{Def}} = \frac{\det(T^*\widetilde{M} \otimes \kappa)}{\det(T\widetilde{M})}.$$

But does a "square root" of this always exist? We see that $K^{\mathrm{vir}} \cong \kappa \cdot (\det T^*\widetilde{\mathcal{M}})^2$, where the second factor is manifestly a square and κ is a *character*, so it always has a square too.

Putting our new candidate virtual bundle into (5),

$$\widehat{O}_x^{\mathrm{vir}} = \underbrace{O_x \prod \frac{1-b_i}{1-a_i}}_{O_x^{\mathrm{vir}}} \otimes (K^{\mathrm{vir}})_x^{1/2}.$$

I prefer to put this in a more symmetric form. Recalling that $\prod a_i$ and $\prod b_i$ are precisely $\det \mathrm{Def}$ and $\det \mathrm{Obs}$, we can rewrite the above as

$$\boxed{\widehat{O}_x^{\mathrm{vir}} = \prod \frac{b_i^{1/2} - b_i^{-1/2}}{a_i^{1/2} - a_i^{-1/2}}.} \tag{6}$$

2.3 Nekrasov's formula

Now we study the generating function

$$\boxed{Z := \sum_n z^n \chi(\mathcal{M}, \widetilde{O}^{\mathrm{vir}}).}$$

You can interpret z^n as an element of the K-theory of some torus, and $\widehat{O}^{\mathrm{vir}}$ as a representation of the maximal torus $T \subset \mathrm{GL}(3)$, say with coordinates

$$T = \left\{ \begin{pmatrix} t_1 & & \\ & t_2 & \\ & & t_3 \end{pmatrix} \right\}.$$

Nekrasov's formula says that Z should be $\mathrm{Sym}^\bullet$ of something, but what is that something?

Define $t_4 = \frac{z}{\sqrt{t_1 t_2 t_3}}$ and $t_5 = \frac{1}{z\sqrt{\kappa}}$. Recalling that $\kappa = t_1 t_2 t_3$, we then have

$$t_1 t_2 t_3 t_4 t_5 = t_1 t_2 t_3 \frac{1}{t_1 t_2 t_3} = 1.$$

This maybe motivates us to think of the situation as arising from some *larger* (five-dimensional) torus action on some larger space, and that the answer will take the form

$$\mathrm{Sym}^\bullet \frac{?}{\prod_{i=1}^5 (t_i^{1/2} - t_i^{-1/2})}.$$

To guess what ? is, we should ask ourselves, for π a 3D partition what is T_π? Remember to think of the (virtual) tangent space as deformations minus obstrutions. Obviously it has three directions t_1, t_2, and t_3. What about the obstructions? For these we have to subtract off $t_1 t_2 + t_1 t_3 + t_2 t_3$ ♠♠♠ TONY: [I didn't really understand the motivation why]. So the numerator ? above should correspond to $t_1 t_2 - t_1 t_3 - t_2 t_3$; in symmetrized terms,

$$\mathrm{Sym}^\bullet \frac{\prod_{1<j\leq 3} (t_i t_j)^{1/2} - (t_i t_j)^{-1/2}}{\prod_{i=1}^5 (t_i^{1/2} - t_i^{-1/2})}.$$

Here is another way to think about it. Let Z be the total space of the vector bundle $\mathcal{L}_4 \oplus \mathcal{L}_5$. Since $\mathcal{L}_4 \otimes \mathcal{L}_5 = K_X$, we have $K_Z = \mathcal{O}_Z$.

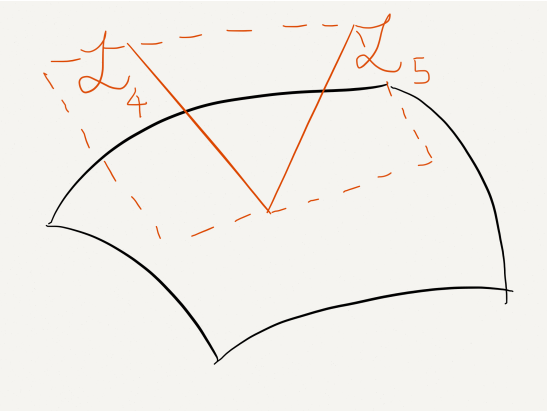

Now Z is five-dimensional, and has a five-dimensional torus action by the weights $t_1, \ldots, t_5$. It turns out that

$$\frac{(t_i t_j)^{1/2} - (t_i t_j)^{-1/2}}{\prod_{i=1}^5 (t_i^{1/2} - t_i^{-1/2})} = \chi(Z, \text{sheaf})$$

where the sheaf is a constant plus $T_Z^* - T_Z$. You can check this by localizing $T_Z^* - T_Z$ at a point: you'll see that the right hand side is the left hand side plus something on which z doesn't act.

Theorem 2.1 (Nekrasov's formula)**.** *In the preceding notation, we have*

$$Z := \sum_n z^n \chi(\mathcal{M}, \widehat{\mathcal{O}}^{\text{vir}}) = \text{Sym}^\bullet \frac{\prod_{1<j\le 3}(t_i t_j)^{1/2} - (t_i t_j)^{-1/2}}{\prod_{i=1}^5 (t_i^{1/2} - t_i^{-1/2})}.$$

Remark 2.2. $\text{Sym}^\bullet$ replaces weights with an infinite product, so the right hand side is some massive infinite product.

2.4 Proof of Nekrasov's formula

2.4.1 Step 1

We will show that

$$Z = S^\bullet \chi(\mathbb{C}^3, \mathcal{F})$$

for some sheaf

$$\mathcal{F} = z\mathcal{F}_1 + z^2\mathcal{F}_2 + \ldots \in K_{\text{equiv}}(\mathbb{C}^3)[[z]].$$

We will construct the $\mathcal{F}_i$ inductively. The first term is

$$\mathcal{F}_1 = -\widehat{\mathcal{O}}^{\text{vir}}_{\text{Hilb}_1(\mathbb{C}^3)}.$$

The key statement has to do with the "Hilbert-Chow map"

$$\text{Hilb}(\mathbb{C}^3, n) \xrightarrow{p} S^n\mathbb{C}^3.$$

Then we can push forward along p in K-theory, and the claim is that $p_*\widehat{\mathcal{O}}^{\text{vir}}$ *factors*. The upshot is that this implies

$$\chi(\mathbb{C}^3, \mathcal{F}_k) = \text{Sym}^\bullet \frac{???}{\prod_{i=1}^3 (1 - t_i^{-1})}$$

This follows from a very general argument, that also applies for example to $\mathcal{O}^{\text{vir}}$. We're going to just assume it and proceed.

2.4.2 Step 2

Now we have to do something specific to the sheaf $\widehat{\mathcal{O}}^{\text{vir}}$ that we so carefully constructed. The next step is to determine the numerator of the factorization discussed in the first step. We first claim that it has a factor:

$$\chi(\mathbb{C}^3, \mathcal{F}_k) = \frac{\square \cdot ??}{\prod_{i=1}^3 (1 - t_i^{-1})}.$$

Why? It is an elementary exercise to show if $t_1 t_2 = 1$, then $Z = 1$ identically. By symmetry, this holds for all three pairs, so the numerator should be divisible by $(t_1t_2 - 1)(t_2t_3 - 1)(t_3t_1 - 1)$. We can rewrite everything in a symmetric manner: since $\kappa = t_1t_2t_3$ we can write $t_1t_2 = \kappa/t_3$, etc. So there is a Laurent polynomial $\square$ such that

$$Z = \text{Sym}^\bullet \left(\square \cdot \prod_i \frac{(\kappa/t_i)^{1/2} - (t_i/\kappa)^{1/2}}{t_i^{1/2} - t_i^{-1/2}} \right). \tag{7}$$

2.4.3 Step 3

Next we claim that the Laurent polynomial $\square \in \mathbb{Z}[t_1^{\pm 1/2}, t_2^{\pm 1/2}, t_3^{\pm 1/2}][[z]]$ from (7) depends only on $\kappa^{1/2}$. This is called "rigidity." The idea is to send $t_i^{\pm 1} \to \infty$ while keeping κ constant.

On one hand, the localization formula (see 6) tells us that Z has the form

$$Z = \sum z^{|\pi|} \prod_{a_i \text{ weight of } \pi} \frac{(\kappa/a_i)^{1/2} - (a_i/\kappa)^{1/2}}{a_i^{1/2} - a_i^{-1/2}}$$

The inner expressions

$$\frac{(\kappa/a_i)^{1/2} - (a_i/\kappa)^{1/2}}{a_i^{1/2} - a_i^{-1/2}} \tag{8}$$

have the property that if $a_i \to 0$ or ∞ then they remain bounded. Of course, the same holds for

$$\prod_i \frac{(\kappa/t_i)^{1/2} - (t_i/\kappa)^{1/2}}{t_i^{1/2} - t_i^{-1/2}}$$

Therefore, specializing $t_i^{\pm 1} \to \infty$ so that κ remains constant, we find that both sides of (7) must remain bounded, hence $\square$ must remain bounded. But as it is a Laurent *polynomial*, this is only possible if it is independent of the all the t_i.

2.4.4 Step 4

We're now going to determine $\square \in \mathbb{Z}[t_1^{\pm 1/2}, t_2^{\pm 1/2}, t_3^{\pm 1/2}][[z]]$ by a judicious choice of specialization. The nice thing about Laurent polynomials is that they are *independent* of how we send the variables to infinity.

Consider (8) again. If $a_i \to \infty$ then (8) tends to $-\kappa^{-1/2}$, and if $a_i \to 0$ then it tends to $-\kappa^{1/2}$. So

$$\prod \frac{(\kappa/a_i)^{1/2} - (a_i/\kappa)^{1/2}}{a_i^{1/2} - a_i^{-1/2}} \to (-\kappa^{1/2})^{\text{virtual index of fixed point}}$$

where the virtual index is $\#\{a_i \to 0\} - \#\{a_i \to \infty\}$.

A nice choice of specialization is $t_1, t_3 \to \infty$ with t_3 negligible compared to t_3, and that determines t_2 by the constraints that $\kappa = t_1 t_2 t_3$ be constant. For instance, imagine

$t_1 = s^{10^{10}-1}$, $t_2 = s^{-10^{10}}$, and $t_3 = s$ as $s \to \infty$.

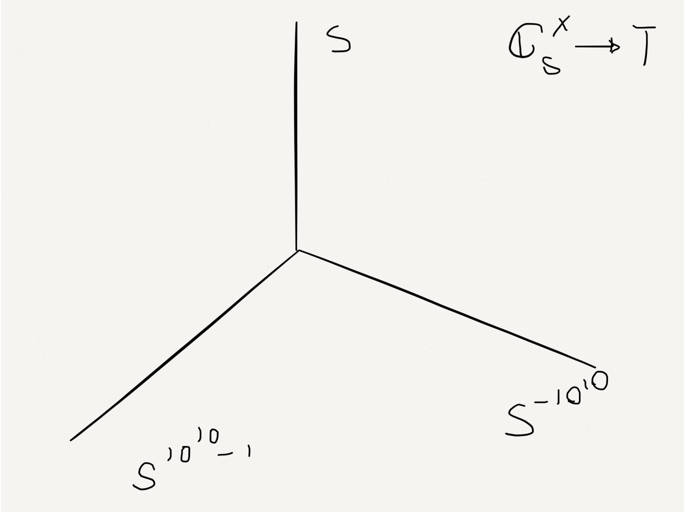

Imagine each partition π as being stacked out of cubes in the first quadrant. Since t_3 is basically negligible, the boundary between whether $a_i \to 0$ or $a_i \to \infty$ is completely determined by the relative values of i_1 and i_2. For a cube at location (i_1, i_2, i_3) the index is the number of boxes with $i_1 - i_2 \geq 0$ minus the number of boxes with $i_1 - i_2 < 0$.

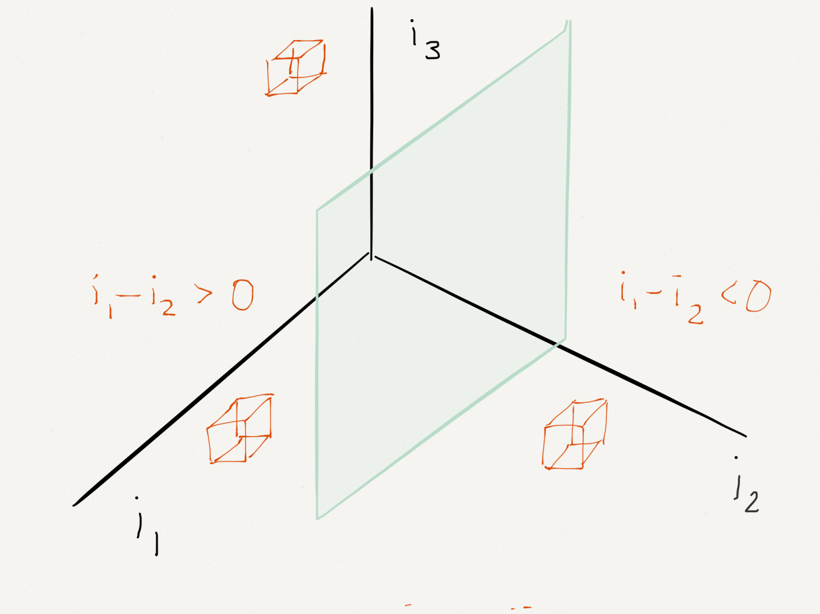

So the sum is then

$$\sum_{\pi} \prod_{\square=(i_1,i_2,i_3)} q_{i_1-i_2}$$

where $\ldots = q_{-2} = q_{-1} = \kappa^{-1/2}z$, and $q_0 = q_1 = \ldots = \kappa^{1/2}z$. It is an exercise to prove that this is equal to

$$\prod_{a\leq 0\leq b} \frac{1}{1 - q_a q_{a+1} \ldots q_b} = \mathrm{Sym}^\bullet \text{ (what we want).}$$

Now we are done.

Remark 2.3. If you take all $q_i = z$ in the formula, then you get $\prod \frac{1}{(1-z^n)^n}$.

3 Introduction to enumerative K-theory

3.1 Initial discussion

We are now ready to move on from counting *points* to counting *curves*.

To being, we take a special threefold of the form $X = C \times \mathbb{C}^2$. This doesn't have so many interesting cycles, since any connected, proper curve must be collapsed to a point via the map to $\mathbb{C}^2$ (since it's affine). However, the Hilbert scheme $\mathrm{Hilb}(X, d[C], n)$ is nontrivial, because we can things such as nonreduced curves "rotating" around the fibers, and extraneous points

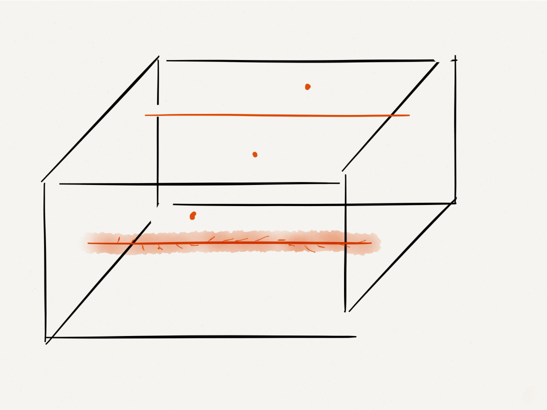

We replace $\mathrm{Hilb}(X, d[C], n)$ with a better space $PT(X, d[C], n)$ (PT stands for Pandharipande-Thomas). The Hilbert scheme parametrizes surjections

$$\mathcal{O}_X \to \mathcal{O}_C \to 0.$$

Consider for the moment the case where X is a surface, say $\mathbb{C}^2$. Then we can view $\mathcal{O}_C$ as being specified by choosing certain polynomials which are not in the ideal of C. For simplicity, imagine a monomial ideal as in the picture.

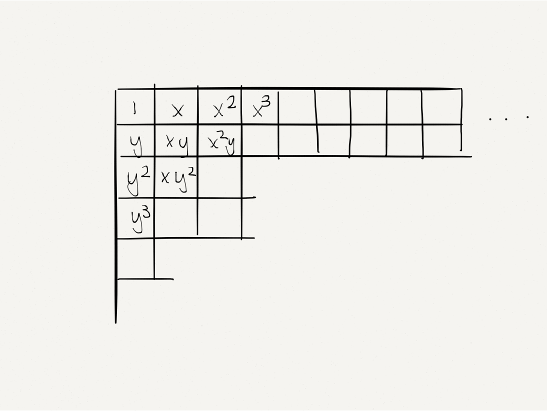

The surjection $\mathcal{O}_X \to \mathcal{O}_C$ can be viewed as "hitting" the top left box with $1 \in \mathcal{O}_X$.

What part of this picture don't we like? The hanging boxes, which form a 0-dimensional subsheaf of $\mathcal{O}_C$.

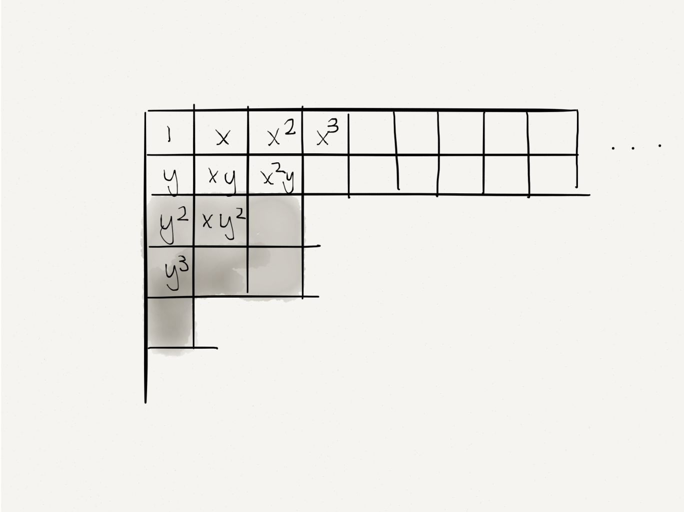

Curves with no zero-dimensional subsheaf are Cohen-Macaulay, but this is not: the hanging block corresponds to an embedded point. The idea of Pandariphande-Thomas is to prohibit these hanging blocks. But that's not a closed condition, so we have to allow something else, and that something else is for $\mathcal{O}_X \to C$ to be *non-surjective*. That is, we can cut off the hanging boxes at the price of allowing a *kernel*.

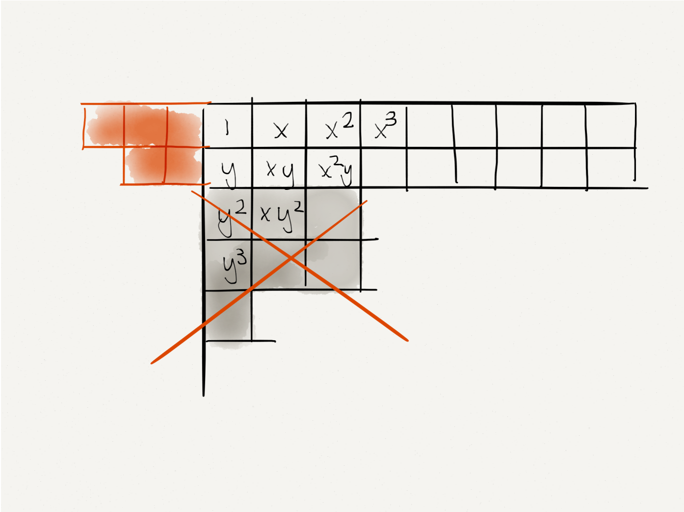

So we consider instead complexes

$$\mathcal{O}_X \to \underbrace{\mathcal{F}}_{\text{pure dim 1}} \to \underbrace{\text{coker}}_{\text{dim } 0} \to 0.$$

Here "pure dimension 1" means that $\mathcal{F}$ has no subsheaf supported on a set of dimension 0.

Now we go back to the situation $X = C \times \mathbb{C}^2$. Let $\pi\colon X \to C$ denote the projection. If $\mathcal{F}$ is a torsion-free sheaf on X, then $\pi_*\mathcal{F}$ is a torsion-free sheaf on a smooth curve, hence a vector bundle $\mathcal{V}$. Multiplication by coordinates in $\mathbb{C}^2$ give $X_1, X_2 \in \operatorname{End}(\mathcal{V})$. Given a section s of $\mathcal{V}$, we demand that

$$P(X_1, X_2)s \to \mathcal{V} \to \underbrace{\text{coker}}_{\dim 0} \to 0.$$

Now we recognize this as $\text{QMaps}(C \to \text{Hilb}(\mathbb{C}^2))$.

More generally, we can replace X by the total space of a sum of two line bundles $\mathcal{L}_1 \oplus \mathcal{L}_2$ over C.

$$\begin{tikzcd} X = \mathcal{L}_1 \oplus \mathcal{L}_2 \arrow[d, "\pi"'] \\ C \end{tikzcd}$$

What is the corresponding PT space? Now multiplication by x_i is not an endomorphism but a map $\mathcal{V} \to \mathcal{V} \otimes \mathcal{L}_i^{-1}$ (since x_i is really a "function" on $\mathcal{L}_i$, i.e. an element of $\mathcal{L}_i^*$). Everything else is the same. This leads to the notion of "twisted quasi-map to Hilb."

3.2 Nakajima varieties

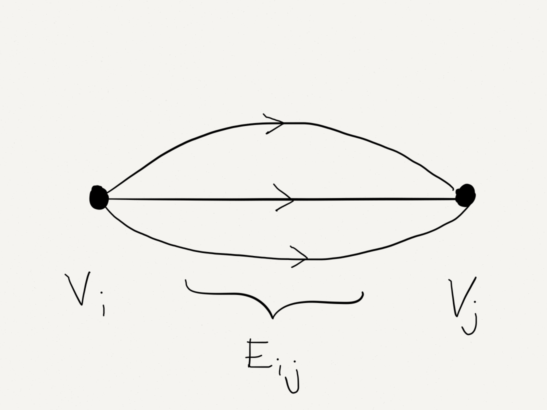

A representation of the quiver consists of

- a vector space V_i for each vertex,
- for each i, j an element of $\operatorname{Hom}(V_i, V_j \otimes E_{ij})$ where E_{ij} is a multiplicity space with dimension equal to the number of edges from i to j.

The corresponding *Nakajima variety M* is a vector bundle version. By definition, a map to M is equivalent to the data of

- vector bundles $\mathcal{V}_i$ for each vertex (varying in moduli),

- for i, j an element of $\mathrm{Hom}(\mathcal{V}_i, \mathcal{V}_j \otimes \mathcal{E}_{ij})$ where E_{ij} is as before.

Remark 3.1. Nakajima considers "double" the data, consisting also of maps in $\mathrm{Hom}(V_j, V_i \otimes E_{ji})$ where $E_{ji} = E_{ij}^\vee \otimes \hbar^{-1}$. It is better to keep this.

The tangent bundle. What is the tangent bundle of M? In terms of the tautological bundles on M, labelled V_i, the class of TM in K-theory is

$$\boxed{TM = \bigoplus_{i,j} \mathrm{Hom}(V_i, V_j \otimes E_{ij}) - (1+\hbar) \bigoplus_i \mathrm{Hom}(V_i, V_i).}$$

♠♠♠ TONY: [some reference was mentioned, couldn't understand the names... Badesh, Jonatan????]

From this we see that there is a polarization

$$T = T^{1/2} + \hbar^{-1} \cdot (T^{1/2})^\vee,$$

where

$$T^{1/2} = \bigoplus_{\substack{\text{half of}\\ \text{double arrows}}} R\,\mathrm{Hom}(V_i, V_j \otimes E_{ij}) - \bigoplus_i R\,\mathrm{Hom}(V_i, V_i).$$

This is upgraded into a bundle version on the moduli space of stable maps to M:

$$\mathcal{T}^{\mathrm{vir}} = \mathcal{T}^{1/2} + \hbar^{-1} \cdot (\mathcal{T}^{1/2})^\vee.$$

where

$$\mathcal{T}^{1/2} = \bigoplus_{\substack{\text{half of}\\ \text{double arrows}}} \mathrm{Hom}(\mathcal{V}_i, \mathcal{V}_j \otimes \mathcal{E}_{ij}) - \bigoplus_i \mathrm{Hom}(\mathcal{V}_i, \mathcal{V}_i).$$

Now specialize to the case $C = \mathbb{P}^1$, with the action of $\mathbb{C}_q^\times$. Let $p_1 = 0$ and $p_2 = \infty$. Then $K_C = -[p_1] - [p_2]$. By localization,

$$T^{\mathrm{vir}} = \mathcal{T}_{p_1}^{1/2} + \mathcal{T}_{p_2}^{1/2} + \square - \hbar^{-1} \cdot \square^\vee.$$

Thus, we should set

$$\boxed{\widehat{O}^{\mathrm{vir}} = O^{\mathrm{vir}} \otimes \left(K^{\mathrm{vir}} = \frac{\det \mathcal{T}_{p_2}^{1/2}}{\det \mathcal{T}_{p_1}^{1/2}}\right)^{1/2}.}$$

We apologize for the confusing notation. The $1/2$ above the tangent spaces refer to tangent sheaf polarizations, and the outer $1/2$ is an honest square root.

Goal. Define vectors, operators, etc. in $K(M)$ by pushing forward $\widehat{O}_{\mathrm{vir}}$ under evaluation maps.

For all but finitely many points, the quasimap takes values in the stable locus. However, at unlucky points we hit the singular locus. Let's investigate by example.

3.3 I-function

We consider the open subset of $\text{QMaps}(\mathbb{P}^1 \to X)$ which are non-singular at $p_2 = \infty$.

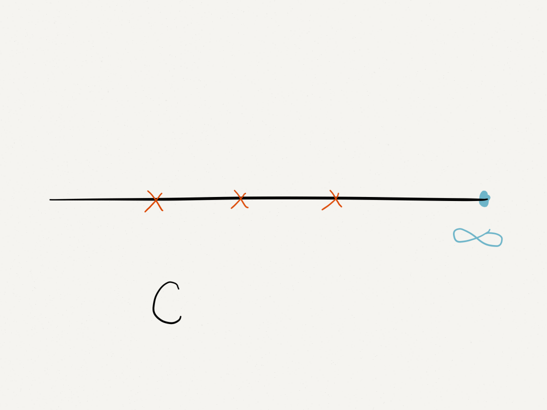

We would like to push forward

$$\text{ev}_*(\text{QMaps}(\mathbb{P}^1 \to X)_{\text{nonsing at } p_2}, \widehat{\mathcal{O}}^{\text{vir}}).$$

Now we can use equivariant localization. We can make a torus $\mathbb{C}_q^\times$ act on the domain of C. Then for any fixed point in the moduli space, the singularities have to form a finite $\mathbb{C}^\times$-invariant subset, which are then necessarily concentrated at the origin. This fixed locus *is* proper, so we can define localization.

The *degree* of a quasi-map is the vector $(\deg \mathcal{V}_i) \in \mathbb{Z}^I$. If we consider

$$\text{ev}_*(\text{QMaps}(\mathbb{P}^1 \to X)_{\text{nonsing at } p_2.}, \widehat{\mathcal{O}}^{\text{vir}} z^{\text{deg}})$$

then we get an element of $K_{\text{equiv}}(X)_{\text{loc}}[[z]]$.

There is a balance between numerators and denominators. If you don't have numerators, then your function is stupid.

The evaluation map at ∞ only makes sense on the open locus away from curves singular at x, so we only have a rational map from $\text{QMaps}(C, M)$.

$$\begin{tikzcd}
 & \text{QMaps} \arrow[dl] \arrow[dr, dotted] & \\
\text{QMaps}_{\text{nonsing at } p} \arrow[rr] & & X
\end{tikzcd}$$

However, there is *another* compactification in which it will extend to a morphism on the entire compatification:

$$\begin{tikzcd}
 & \text{QMaps}_{\text{relative to } p} \arrow[dl] \arrow[dr, dotted] & \\
\text{QMaps}_{\text{nonsing at } p} \arrow[rr] & & X
\end{tikzcd}$$

Let's explore how this alternate compactification looks in a special case. Let $X = T^*\mathbb{P}^0 = \mathbb{C}/\mathbb{C}^\times$. This is the quiver variety corresponding to the pictured quiver:

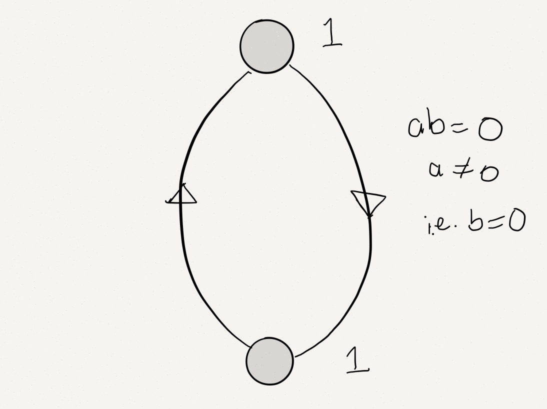

Now what's a quasi-map to X? For every vertex you get a bundle of that rank, so we get a line bundle $\mathcal{L}$ on C. Then, we get a non-zero arrow (corresponding to a) from the trivial bundle to $\mathcal{L}$, which is the same as a non-zero section

$$\mathcal{O}_C \xrightarrow{s} \mathcal{L}.$$

Since $\mathcal{L} \cong \mathcal{O}(D)$, we find that

$$\mathrm{QMaps}(X) = \{\mathcal{O}_C \to \mathcal{O}_C(D)\} = \bigsqcup_n \mathrm{Sym}^n\, C.$$

Going back to the picture

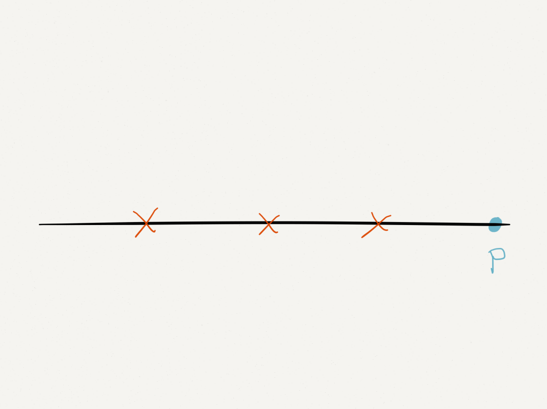

the singularities are the zeros of the section s, and the open subset $\mathrm{QMaps}_{\text{nonsing at } p}$ consists of quasimaps whose s doesn't vanish at p.

How can we compactify this? Well it's clear from the picture!

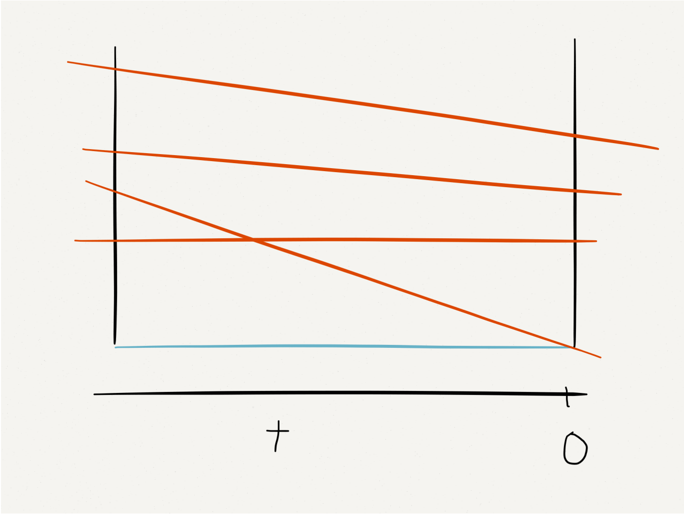

Algebraic geometers are angry people: when they see something they don't like, they just blow up. In the figure above, we have a family of curves with (basepoint) divisors avoiding p, with at least one point limiting to p over the central fiber of the family. When we blow up, we get an exceptional divisor, and the blue section and the red don't intersect anymore. So we get *two* curves $C_0 = C$ and C_1.

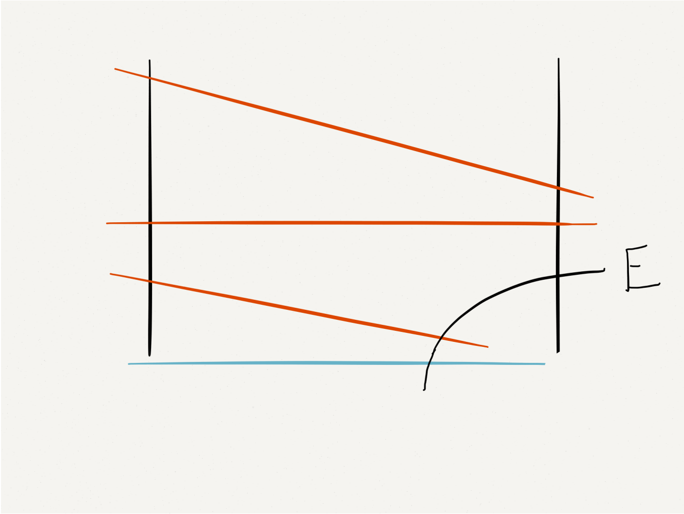

If you reparametrize this curve, it shouldn't affect the result. Any reparametrization has a differential at 0, which will scale the exceptional fiber, so we have an action of $\mathrm{Aut}(\mathbb{C}_t, 0)$ on the curve. For example, if there is one singular point on the exceptional fiber then there is no moduli of the "location" of the singular points (the moduli is $\mathbb{C}^\times/\mathbb{C}^\times$), but if there are more points then the moduli is larger.

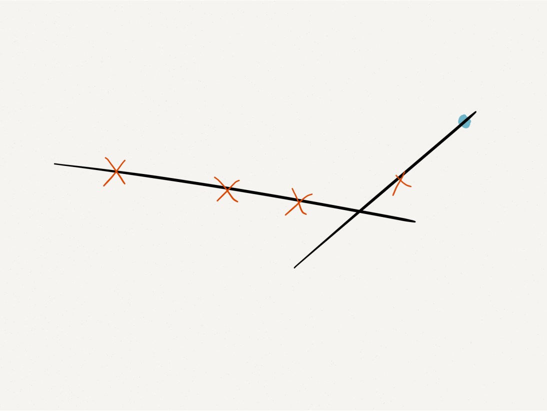

In general, you may have to blow up many times.

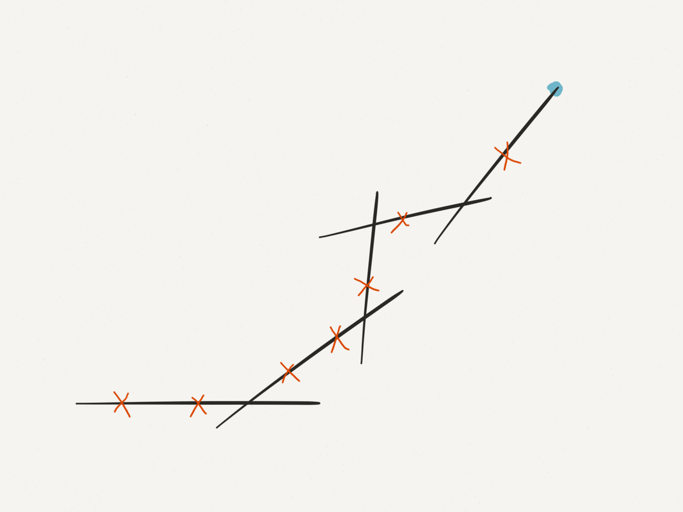

The singularities are away from the nodes and the marked point. There is a torus action by $(\mathbb{C}^\times)^{\#\text{exceptional divisors}}$, and the stabilizer subgroup is finite.

Capped vertex. The picture above is assembled out of two parts. One is the collection of exceptional components, and the other is the "original" C. Now we can attempt to use $\mathbb{C}_q^\times$ equivariant localization on this new compactification.

As before, all the singularities must be over 0. So the capped vertex is a combination of the "bare vertex," which comes from the pileup of singularities at $0 \in C$, multiplied by an operator $J := \text{ev}_*(\text{QMaps}_{\substack{\text{nonsing. at } p_1 \\ \text{relative } p_2}})$.

$$\text{capped vertex} = \text{bare vertex} \cdot \text{ev}_*(\text{QMaps}_{\substack{\text{nonsing. at } p_1 \\ \text{relative } p_2}}).$$

The operator J is a fundamental solution of a certain q-difference equation. The logic is the following. The object bare vertex contains a lot of information, while the capped virtex has very little information (you could make it better by replacing the virtual class with something more interesting like a tautological class). You get the interesting object from the silly object by applying a q-difference equation.

The remaining lectures will focus on understanding this q-difference equation from the perspective of geometric representation theory.