TROPICAL ENUMERATION OF CURVES IN BLOWUPS OF $\mathbb{C}P^2$

BRETT PARKER

Abstract. We describe a method for recursively calculating Gromov–Witten invariants of all blowups of the projective plane. This recursive formula is different from the recursive formulas due to Göttsche and Pandharipande in the zero genus case, and Caporaso and Harris in the case of no blowups. We use tropical curves and a recursive computation of Gromov–Witten invariants relative a normal crossing divisor.

This paper computes Gromov–Witten invariants of blowups of the projective plan from recursively computable relative Gromov–Witten invariants. Recursions calculating some such Gromov–Witten invariants are already known. Göttsche and Pandharipande give a recursive formula for zero-genus Gromov–Witten invariants of arbitary blowups of the plane in [6]. Caporaso and Harris in [5] show that the Gromov–Witten invariants of $\mathbb{C}P^2$ relative to a line may be calculated recursively, giving a method for calculating Gromov–Witten invariants of $\mathbb{C}P^2$ of any genus. This is extended by Vakil in [20] to a recursive formula for Gromov–Witten invariants of $\mathbb{C}P^2$ blown up at a small number of points, and further extended by Shoval and Shustin, then Brugallè in [19, 3] to $\mathbb{C}P^2$ blown up at 7 and 8 points respectively. Our recursion gives a formula for (arbitary-genus) Gromov–Witten invariants[1] that is uniform in the number of blowups.

Behind our recursion are exploded manifolds, [12], and the tropical gluing formula from [18]. This paper should be readable to someone not familiar with exploded manifolds, however we lose some precision of language by forgoing any serious use of exploded manifolds.

In section 1, we explain what relative invariants we use, and explain the recursion that these invariants satisfy. Section 2 explains why this recursion follows from a simplified tropical gluing formula, Theorem 2.1. This simplified gluing formula is applied in Section 3 to reconstruct our absolute Gromov-Witten invariants from our relative invariants and some further relative invariants of $\mathbb{C}P^2$ blown up at one point, computed in section 4.

I have written a Mathematica program that computes these relative Gromov–Witten invariants, and compared the results to known Gromov–Witten invariants of $\mathbb{C}P^2$, and the results of [6, 3, 4], however further work is required to get these invariants in a closed form. For example, beyond checking up to degree 12, I have been unable to reprove the beautiful formulas of Bryan and Leung [4] which count curves of genus g passing through g points in $\mathbb{C}P^2$ blown up at 9 points. This formula also can be proved using topological recursion for certain descendant invariants and the symplectic sum formula for Gromov–Witten invariants as in [9], however the consequences in our model are less than obvious. The Mathematica

This work was supported by ARC grant DP1093094.

[1] We only count rigid curves with no constraints — the counts of curves constrained to pass through a collection of points may be recovered by blowing up these points, and counting rigid curves that intersect the new exceptional spheres once.

program and a link to a talk with lots of pictures is available on my website: http://maths-people.anu.edu.au/~parkerb/publications.html.

1. The relative Gromov–Witten invariants

Consider the following sequence of blowups of $\mathbb{C}P^2$. Choose a line N_0 in $\mathbb{C}P^2$, then blow up $\mathbb{C}P^2$ at a point on N_0, and label the exceptional sphere N_1. Then blowup at a point on N_1, and label the new exceptional sphere N_2. Let M be the complex manifold obtained by making n blowups in this fashion, and let N be the normal crossing divisor consisting of the union of (the strict transforms of) $N_0, \dots, N_n$. If toric blowups are used, the following is a moment map picture of (M, N) and the corresponding toric fan.

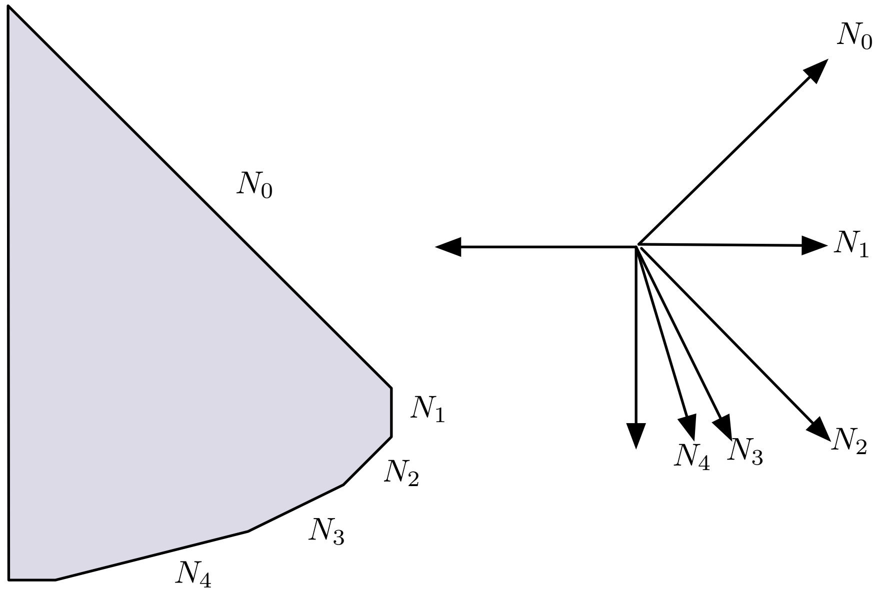

Integral vectors in the non-negative span of $(1,1)$ and $(1, 1-n)$ correspond to contact data of curves in M with N as follows: A vector $(d, d-id)$ corresponds to a point where a curve is required to have contact order d with N_i, and a vector, $a(1,1-i)+b(1,1-i-1)$ where a and b are positive integers, corresponds to a point sent to $N_i \cap N_{i+1}$, where the curve is required to have contact order a to N_i and b to N_{i+1}.

Let Γ be a finite set of integral vectors in the non-negative span of $(1,1)$ and $(1,-n+1)$. For reasons that will become apparent latter, we will identify such a Γ with a 'connected rigid tropical curve in the span of $(1,1)$ and $(1,-n+1)$'. For example, the following is a picture of the rigid tropical curve corresponding to $\Gamma = \{(1,0),(1,-2)\}$.

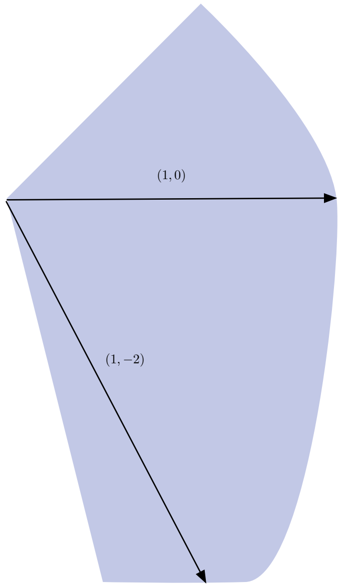

The homology class represented by a holomorphic curve in M is determined by its contact data Γ — for example the degree (or intersection with a line not passing through any of the blown up points) is $\sum_{(a,b)\in\Gamma} a$. For the virtual dimension of the space of curves of genus g and contact data Γ to be 0, we need that

$$g = 1 - \sum_{(a,b)\in\Gamma} (a+b+1) \ . \tag{1}$$

Let n_Γ be the virtual number of rigid curves in M with contact data Γ. In the case that a vector appears more than once in Γ, it is important to clarify that n_Γ counts curves with points labeled by the vectors in Γ. As a consequence, we sometimes must divide the corresponding count by the automorphisms of Γ. These numbers n_Γ are relative Gromov–Witten invariants defined either using exploded manifolds [12, 14, 17], log Gromov–Witten theory [1, 7], or Ionel's method for defining GW invariants relative normal crossing divisors [8]. These three approaches give the same invariants, as discussed in [13] and [15].

Arrange these relative Gromov–Witten invariants of (M, N) into a generating function

$$F_n := \sum_\Gamma \frac{n_\Gamma}{|\mathrm{Aut}\,\Gamma|}\Gamma$$

where the sum is over all connected rigid tropical curves Γ in the cone spanned by $(1,1)$ and $(1,-n+1)$.

Absolute Gromov–Witten invariants of $\mathbb{C}P^2$ blown up at n points may be recovered from F_n as follows. Consider the generating function

$$G_n := \sum n_{g,\beta} x^{g-1} q^\beta$$

where $n_{g,\beta}$ is the (virtual) number of rigid genus g curves in $\mathbb{C}P^2$ blown up at n points and representing the homology class β. To obtain G_n from F_n, we shall use a $\mathbb{R}$-linear map Ψ defined as follows. If $\Gamma = \{(1, 1-m_1), \dots, (1, 1-m_k)\}$, then

$$\Psi(\Gamma) := \prod_{i=1}^{k} \left(x^{m_i - 3} q^H \sum_{j_1 < \dots < j_{m_i}} q^{-\sum_{l=1}^{m_i} E_{j_l}} \right) \tag{2}$$

and $\Psi(\Gamma) = 0$ if Γ is not in the above form. In the above, H indicates the homology class represented by (the pullback of) a line, and E_i is the homology class of the ith exceptional sphere. Note that the sum on the right is the m_i-th elementary symmetric polynomial σ_{m_i} in the variables $\{q^{-E_1}, \dots, q^{-E_n}\}$, so the above formula may be written more succinctly as

$$\Psi(\{(1,1-m_1),\dots,(1,1-m_k)\}) = \prod_{i=1}^{k} x^{m_i-3} q^H \sigma_{m_i} \ .$$

Then

$$G_n = \Psi(F_n) + \sum_{i=1}^{n} x^{-1} q^{E_i} \ .$$

By allowing the number of blowups to approach infinity, we may write a formula that is uniform in the number of blowups. Let

$$F = \sum_{\Gamma} \frac{n_\Gamma}{|\operatorname{Aut}\Gamma|}\Gamma$$

now be a sum over all finite sets Γ of vectors in the form (a,b) where $a > 0$ and $b \le a$, and n_Γ is the corresponding relative Gromov–Witten invariant of (M,N) where the number n of blowups used[2] is enough that the vectors in Γ are in the non-negative span of $(1,1)$ and $(1,1-n)$. Then

$$G = \Psi(F) + \sum_i q^{E_i} \tag{3}$$

is a generating function representing the rigid curves in all possible blowups of $\mathbb{C}P^2$.

We shall now describe how to recursively compute F. Our recursion is simplified by using e^F, the generating function of possibly disconnected curves. Use the convention that $(n_1\Gamma_1)(n_2\Gamma_2) := (n_1 n_2)\Gamma_1 \coprod \Gamma_2$. Then write the generating function for possibly disconnected curves as

$$e^F = \sum_{\Gamma} \frac{n_\Gamma}{\operatorname{Aut}\Gamma}\Gamma$$

where the sum is now over possibly disconnected rigid tropical curves Γ, and when $\{\Gamma_i\}$ is some collection[3] of connected rigid tropical curves, $n_{\coprod \Gamma_i} = \prod_i n_{\Gamma_i}$. An alternate description of a possibly disconnected rigid tropical curve is as a finite set of integral vectors (a,b) so that $a > 0$ and $b \le a$ along with an equivalence relation — where we say some of these vectors are connected to each other. Again, we can read off the degree and Euler characteristic of the curves that n_Γ counts as before:

$$\text{Degree: } \deg\Gamma := \sum_{(a,b)\in\Gamma} a \tag{4}$$

$$\text{Euler Characteristic: } \chi(\Gamma) := \sum_{(a,b)\in\Gamma} (1+2a+2b)$$

The Euler characteristic above is the Euler characteristic of the surface with boundary obtained by taking the real oriented blowup of a curve at all contact points. So if a connected curve has genus g and k contact points, we say its Euler Characteristic is $2-2g-k$.

The generating function e^F is completely determined by the relations

$$e^F \overleftarrow{y}\, y = y\, \overrightarrow{y}\, e^F$$

where $y = (y_1, y_2)$ is any integral vector with $y_1 \le -1$ and $y_2 \ge y_1$, and the $\mathbb{R}$-linear operators $\overleftarrow{y}\, y$ and $y\, \overrightarrow{y}$ have a combinatorial definition, described shortly. In fact,

$$e^F \overleftarrow{y}\, y = \sum_{\Gamma \ni y} \frac{n_\Gamma}{\operatorname{Aut}\Gamma}\Gamma = y\, \overrightarrow{y}\, e^F \tag{5}$$

[2]So long as n is large enough for n_Γ to be defined, the value of n_Γ does not depend on n.

[3]We include the case $\Gamma = \varnothing$. There is a unique empty curve, so $n_\varnothing = 1$.

where now the sum is over possibly disconnected rigid tropical curves[4] Γ with a special edge labeled by y and called 'incoming', and n_Γ is a Gromov–Witten invariant corresponding to such a Γ. The target of the evaluation map at a point with contact data $-y$ is a 2-dimensional space. Use the convention that an incoming vector such as y in Γ corresponds to specifying that a point must have contact data $-y$, and the evaluation map at this point must be constrained to a specified point.[5] With this convention, n_Γ is the relative Gromov–Witten invariant[6] counting curves with data Γ. We shall see that $e^F\overleftarrow{y}\,y$ and $y\,\overrightarrow{y}\,e^F$ correspond to calculating n_Γ by specifying this constraint to be in two different positions.

An important ingredient for understanding the action of $\overleftarrow{y}\,y$ is the action of $\overleftarrow{\alpha}$ on a single vector v, given by

$$v\overleftarrow{\alpha} = \overleftarrow{\alpha}\,v + \max\{v \wedge \alpha, 0\}\overleftarrow{(v+\alpha)}$$

where we can think of the rightmost term as $\overleftarrow{\alpha}$ interacting with v, and the other term as α not interacting with v, as pictured below.

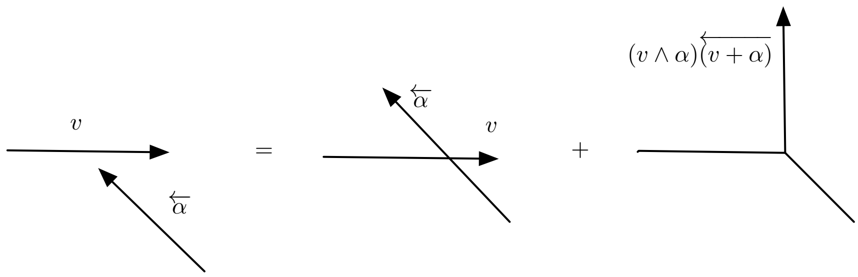

We use the above interactions to define an action of $\overleftarrow{y}\,y$ as follows:

- Choose an order on the vectors $\{v_i\}$ making up Γ so that v_i is to the right of v_j if $v_j \wedge v_i < 0$, where $(a,b)\wedge(c,d) := ad - bc$.
- Set
$$\Gamma\overleftarrow{y}\,y := v_1 \cdots v_n \overleftarrow{y}\,y$$
where the expression on the right indicates the ordered set $(v_1, \dots, v_n, \overleftarrow{y}, y)$ along with the equivalence relation connecting $\overleftarrow{y}$ only to y, and otherwise connecting the vectors v_i as in Γ.
- Move the vector decorated by the arrow to the left using the linear relations generated by
$$v_1 \cdots v_n \overleftarrow{y}\,y = v_1 \cdots v_{n-1} \overleftarrow{y}\,v_n y + \max\{v_n \wedge y, 0\} v_1 \cdots v_{n-1}\overleftarrow{(v_n + y)}y$$
where the first expression on the right has the same equivalence relation as above, and the second expression has the induced equivalence relation connecting everything formerly connected to v_n or $\overleftarrow{y}$ to $\overleftarrow{(v_n+y)}$. Similarly,
$$v_1 \cdots v_k \overleftarrow{\alpha}\,w \cdots = v_1 \cdots v_{k-1}\overleftarrow{\alpha}\,v_k w \cdots + \max\{v_k \wedge \alpha, 0\} v_1 \cdots v_{k-1}\overleftarrow{(v_k+\alpha)}v_k w \cdots$$
where the the second expression on the right uses the induced equivalence relation connecting everything formerly connected to v_k or $\overleftarrow{\alpha}$ to $\overleftarrow{(v_k+\alpha)}$.

[4]Equivalently, these Γ consist of a finite set containing y and integral vectors (a,b) with $a > 0$ and $b \le a$, along with an equivalence relation saying which vectors are connected.

[5]In the case that y is k times a primitive vector, the target of the evaluation map $\mathbf{E}$ is actually an exploded orbifold which is the quotient of an exploded manifold by a trivial $\mathbb{Z}_k$ action. In this case, constraining the moduli space to the specified point means taking the fiber product of the moduli space with a map of a point into $\mathbf{E}$. This multiplies n_Γ by k when compared with the other reasonable interpretation.

[6]We use [14, 17] to define our relative Gromov-Witten invariants — constraining contact data to a point corresponds to integrating the pullback of the Poincare dual to that point over the virtual fundamental class.

- define

$$\overleftarrow{(a,b)}w\cdots := \begin{cases} \{(a,b), w, \dots\} \text{ if } a > 0 \text{ and } b \le a \\ \{w, \dots\} \text{ if } (a,b) = (-1,0) \\ 0 \text{ otherwise} \end{cases}$$

A graphical illustration of two cases of this process is below: on the left, we see $\overleftarrow{y}$ interacting with two edges of Γ, and in the second picture, we see $\overleftarrow{y}$ interacting with one edge to produce $\overleftarrow{(-1,0)}$, which is then discarded.

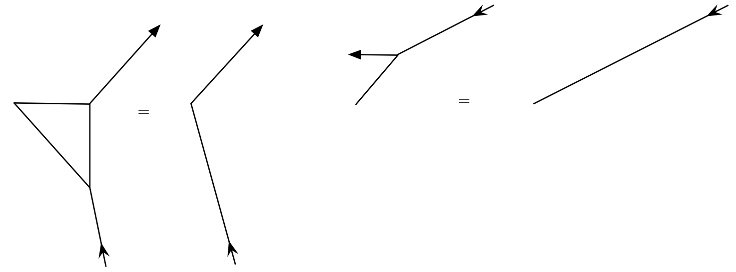

Similarly define $y\overrightarrow{y}\Gamma$ as follows:

- $$y\overrightarrow{y}\Gamma := y\overrightarrow{y}v_1\cdots v_n$$
- $$\cdots w\overrightarrow{\alpha}v_k v_{k+1}\cdots = \cdots w v_k\overrightarrow{\alpha}v_{k+1}\cdots + \max\{0, \alpha\wedge v_k\}\cdots w\overrightarrow{(\alpha+v_k)}v_{k+1}\cdots$$

where the equivalence relation of the first expression on the right is the same as that from $\cdots w\overrightarrow{\alpha}v_k v_{k+1}\cdots$, and the equivalence relation of the second expression on the right is obtained by connecting everything formerly connected to $\overrightarrow{\alpha}$ or v_k to $\overrightarrow{(\alpha+v_k)}$.

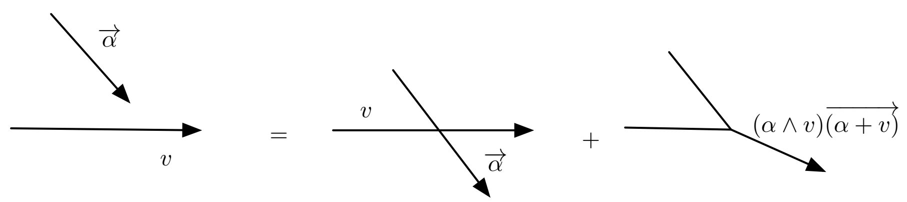

- $$\cdots w\overrightarrow{(a,b)} := \begin{cases} \{\dots, w, (a,b)\} \text{ if } a > 0 \text{ and } b \le a \\ \{\dots, w\} \text{ if } (a,b) = (0,-1) \\ 0 \text{ otherwise} \end{cases}$$

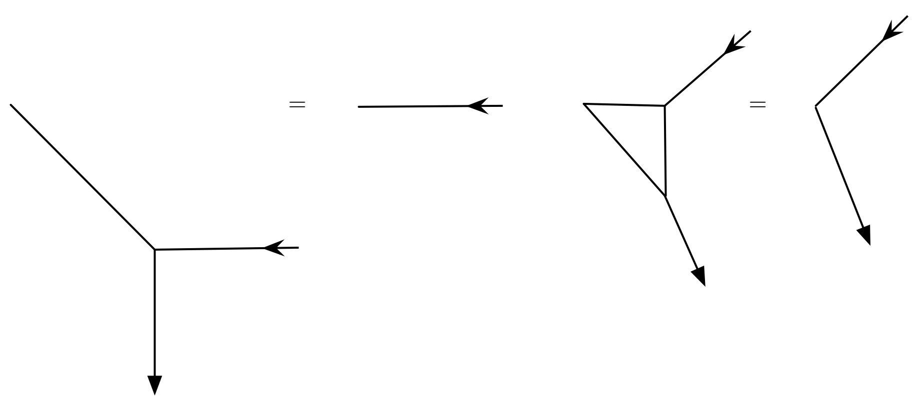

We shall argue below that these relations, $e^F\overleftarrow{y}\,y = y\,\overrightarrow{y}\,e^F$, recursively determine e^F, but let us first consider some examples. To begin with, all we know is that $e^F = 1\varnothing + \cdots$, nevertheless, we can use this information to compute $n_{(-1,0)}$ using $e^F\overleftarrow{(-1,0)}(-1,0)$. The only contributing term is $n_\varnothing\overleftarrow{(-1,0)}(-1,0)$ — as pictured below on the left — so $n_{(-1,0)} = 1$. Computing using $(-1,0)\overrightarrow{(-1,0)}e^F$, the only contributing term is

$$n_{(1,-1)}(-1,0)\overrightarrow{(-1,0)}(1,-1) = n_{(1,-1)}\{(-1,0)\}$$

as pictured below on the right, so $n_{(1,-1)} = n_{(-1,0)} = 1$, and $e^F = 1\varnothing + 1\{(1,-1)\} + \dots$.

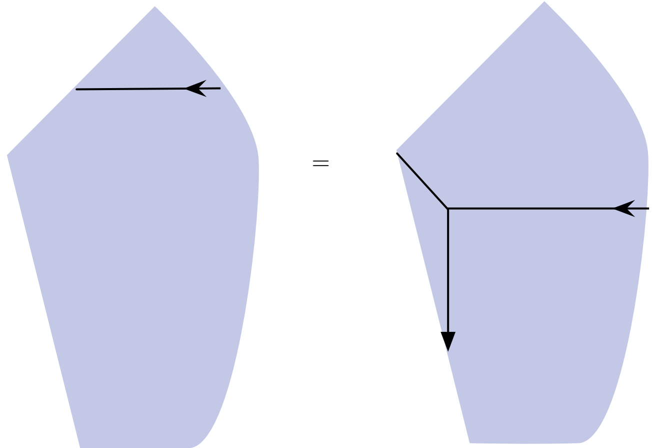

In fact, $n_{(1,-1)}$ is the only nonzero degree 1 invariant involved in e^F. To see that $n_{(1,-1-g)} = 0$ for all $g > 0$, note that computing $n_{(-1,g)}$ using $(-1,g)\overrightarrow{(-1,g)}e^F$ gives $n_{(-1,g)} = n_{(1,-1-g)}$ — as pictured below on the right — but computing using $e^F\overleftarrow{(-1,g)}(-1,g)$ gives $n_{(-1,g)} = 0$, because $(-1,g)$ will never be able to be turned into $(-1,0)$ by interacting with vectors (a,b) with $a > 0$, as indicated in the picture below on the left.

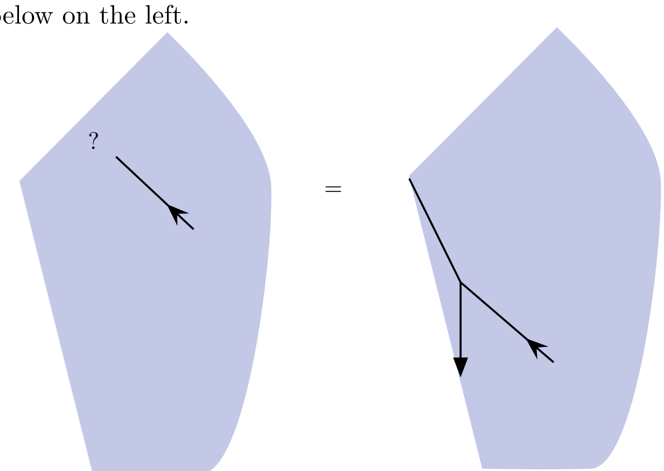

Similarly, $n_{(-2,1)} = 0$ because $(-2,1)$ can't be turned into $(-1,0)$ by interacting with vectors (a,b) so that $a > 0$ and $(a,b)\wedge(-2,1) > 0$. Calculating $n_{(-2,1)}$ the other way gives $n_{-2,1} = (n_{(1,-2)})^2/2 + 2n_{(2,-2)}$, so $n_{(2,-2)} = -\frac{1}{4}$. Pictorially, this calculation is as follows:

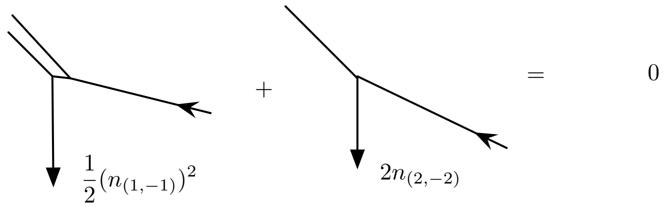

In fact, the only nonzero terms in e^F with a single vector are $n_{(k,-k)} = \frac{(-1)^{k+1}}{k^2}$. This may be verified by calculating $n_{(-k,b)} = 0$ for all $b > k$. A nice implication[7] of this is that if (a, b) and $(a+k, b-k)$ are appropriate incoming and outgoing vectors respectively, then

$$n_{\{(a,b),(a+k,b-k)\}} = \binom{|(a,b) \wedge (k,-k)|}{k}.$$

In particular, $n_{\{(-1,3),(1,1)\}} = 1$, so we can calculate $n_{\{(1,-4),(1,1)\}} = 1$. Pictorially, the computations involved are as follows:

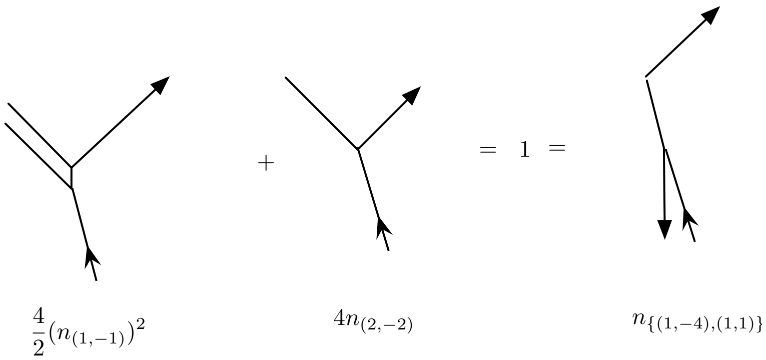

Let us now argue that equation (5) determines e^F recursively. Replace an edge (a, b) from Γ with an incoming edge $(-a, -1-b)$ to obtain Γ'. Then the coefficient of Γ' in $(-a,-1-b)\overrightarrow{(-a,-1-b)}e^F - e^F\overleftarrow{(-a,-1-b)}(-a,-1-b)$ is $a \times n_\Gamma / |\mathrm{Aut}\,\Gamma'|$ — obtained by[8] $\overrightarrow{(-a,-1-b)}$ interacting with (a,b) from Γ producing $a\overrightarrow{(0,-1)}$, which does not interact with any more edges — plus terms involving $n_{\Gamma''}$ where Γ'' has strictly greater Euler characteristic and not larger degree — obtained from $\overrightarrow{(-a,-1-b)}$ interacting with multiple edges of Γ'' to produce $\overrightarrow{(0,-1)}$, or $\overrightarrow{(-a,-1-b)}$ or $\overleftarrow{(-a,-1-b)}$ interacting with edges of Γ'' to produce either $\overrightarrow{(-1,0)}$ or one of the other edges of Γ'. As the number of Γ with degree bounded above and Euler characteristic bounded below is finite, this equation determines e^F entirely.

2. How recursive calculation of the relative invariants follows from a gluing formula.

This section explains how equation (5) follows from a gluing formula. Applying the explosion functor from [12] to (M, N) gives an exploded manifold $\mathbf{M}'$. One way of defining the relative Gromov–Witten invariants of (M, N) is as the Gromov–Witten invariants of $\mathbf{M}'$.

Each exploded manifold has a tropical part. The tropical part $\underline{\mathbf{M}'}$ of $\mathbf{M}'$ is the nonnegative span of $(1, 1)$ and $(1, 1-n)$ subdivided by the rays $(1, 1-k)$. For us, this subdivison is extraneous information which we can ignore, because $\mathbf{M}'$ is a refinement of an exploded manifold $\mathbf{M}$ with tropical part the nonnegative span

[7] See the derivation of equations (6) and (7) later in the paper.

[8] If there are k edges (a, b) in Γ, the k-fold choice of which of these edges to interact with is compensated for by $k = |\mathrm{Aut}\,\Gamma| / |\mathrm{Aut}\,\Gamma'|$.

of $(1,1)$ and $(1,1-n)$, and Gromov–Witten invariants do not change under the operation of refinement.

A refinement is a kind of blowup of an exploded manifold. In this case, $\mathbf{M}$ is the explosion of (M_0, N'), where M_0 is a toric space with the toric fan defined by the rays $(0,-1)$, $(-1,0)$, $(1,1)$, $(1,1-n)$, and N' is the divisor corresponding to $(1,1)$ and $(1,-n)$. (M,N) is a kind of toric blowup of (M_0,N'). In [11], it is proved that the virtual moduli space of curves in a refinement $\mathbf{M}'$ of an exploded manifold $\mathbf{M}$ is a refinement of the virtual moduli space of curves in $\mathbf{M}$. The upshot of this is that if we are counting the virtual number of curves in $\mathbf{M}'$ with some given constraints, the same count is achieved in $\mathbf{M}$. Similar invariance of log Gromov–Witten invariants under such blowups is proved in [2].

The tropical part of each holomorphic curve in $\mathbf{M}$ is a tropical curve in $\underline{\mathbf{M}}$. For us, such a tropical curve means a map of a complete metric graph into $\underline{\mathbf{M}}$ with integral derivative on its edges. In the interior of $\underline{\mathbf{M}}$, such tropical curves also obey the balancing condition familiar to tropical geometers: the sum of the derivatives of all edges leaving a vertex is 0.

For understanding equation (5), we regard n_Γ as a Gromov–Witten invariant of $\mathbf{M}$. The curves in $\mathbf{M}$ relevant to n_Γ have tropical part continuously deformable to Γ. In particular, the ends — or semi-infinite edges — of these tropical curves have derivatives equal to the vectors in Γ, and for the n_Γ from equation (5), there is a distinguished end with an extra constraint on its position, and with 'incoming' derivative equal to y. As we shall see, the computation of n_Γ using $e^F\overleftarrow{y}\,y$ corresponds to constraining this incoming edge above the ray $-y$, and using $y\,\overrightarrow{y}\,e^F$ corresponds to constraining this incoming edge below the ray $-y$.

After choosing where to constrain this incoming edge, the Gromov–Witten invariant n_Γ decomposes into contributions from tropical curves which are rigid when the incoming edge is constrained. If the incoming edge is constrained to lie on the ray $-y$, then the only such rigid curve is Γ itself. Otherwise, there are many possible curves γ, each telling how to combine 'relative' Gromov–Witten invariants corresponding to vertices to obtain a contribution to n_Γ.

For a vertex v of γ, the space that is used to define these relative Gromov–Witten invariants depends on the location of v in $\underline{\mathbf{M}}$. In particular:

(A) If v is sent to the interior of $\underline{\mathbf{M}}$, we use the Gromov–Witten invariants of the exploded manifold $\mathbf{T}^2$, or equivalently we may use the relative Gromov–Witten invariants of any compact two-complex-dimensional toric manifold relative to its toric boundary divisors.

The tropical curve γ_v in $\mathbb{R}^2$ obtained by extending the edges of γ adjacent to v to be infinite defines contact data for such a space. The contact data γ_v determines the homology class represented by the corresponding curves, and for topological reasons, the sum of the derivatives of edges of γ_v leaving v is 0. The virtual (complex) dimension of the space of genus g curves with (unconstrained) contact data γ_v is the valence of v plus $g-1$.

From this, we might hope to rigidify a genus 1 curve by constraining all edges of γ_v, but actually, generic constraints preclude any curves. This may be seen tropically: every balanced tropical curve in $\mathbb{R}^2$ has its infinite edges obeying a one-dimensional constraint caused by the balancing condition. Therefore, there is no tropical curve with infinite edges in the directions specified by γ_v constrained generically.

With genus 1 curves now excluded, the only case in which we can make such curves rigid by constraints on their edges is the case of genus 0, where the space of curves with contact data γ_v has virtual (complex) dimension the valence of v minus 1. Such curves can be made rigid by constraining all but one edge.

(B) If v is sent to a codimension 1 boundary of $\underline{\mathbf{M}}$, we use the Gromov–Witten invariants of $\mathbf{T}$ times the explosion of $(\mathbb{C}P^1, 0)$, or equivalently, we may use the relative Gromov–Witten invariants of $(\mathbb{C}P^1)^2$ relative to 3 of its 4 toric boundary divisors.

In this case, the tropical curve γ_v, obtained by infinitely extending the edges attached to v, translates into contact data for this model in two different ways depending if v is on the top or bottom boundary of $\underline{\mathbf{M}}$: simple contact with the middle divisor corresponds to the vector $(1,0)$ if v is on the upper boundary of $\underline{\mathbf{M}}$, and $(0,1)$ if v is on the lower boundary. Simple contact with one of the other two divisors corresponds to the vectors $\pm(1,1)$ if v is on the top boundary, and $\pm(1,1-n)$ if v is on the bottom boundary.

Again, the homology class represented by a curve in such a model is determined by its contact data, and for topological reasons, a balancing condition is obeyed at such vertices: the sum of the derivatives of edges leaving v is equal to a nonnegative multiple of $(1,0)$ if v is at the top boundary, or a nonnegative multiple of $(0,1)$ if v is at the bottom boundary.

The virtual dimension of genus g curves with contact data γ_v is the valence of v plus $g-1$ plus d, where the sum of the derivatives of edges of γ_v leaving v is $(0,d)$ or $(d,0)$ (depending on which boundary v is on).

The only interesting case when we can make such curves rigid by constraints on their edges is when the genus is 0 and the sum of the derivatives of edges of γ_v is $(1,0)$ or $(0,1)$ respectively. Then such curves may be rigidified by constraining all edges of γ_v.

(C) At vertices sent to the corner of $\underline{\mathbf{M}}$, we use the Gromov–Witten invariants of $\mathbf{M}$, or alternately the relative Gromov–Witten invariants of the manifold with normal crossing divisor (M,N) described in the previous section. Our main interest is the moduli spaces of curves in $\underline{\mathbf{M}}$ that are rigid without being constrained. These rigid moduli spaces have degree and genus determined by the contact data γ_v as described in equations (1) and (4), however there are also non-rigid moduli spaces with the same contact data but lower genus.

We shall argue below that the only tropical curves γ contributing to n_Γ satisfy the following conditions:

If we remove from γ all vertices sent to the corner of $\underline{\mathbf{M}}$ and all edges attached to such vertices, we obtain a connected linear graph. This connected linear graph has the constrained edge at one end, bivalent vertices sent to the interior of $\underline{\mathbf{M}}$ in the middle, and either a univalent vertex sent to a boundary of $\underline{\mathbf{M}}$ or another infinite edge at the other end. In addition, the following conditions hold.

(A) Every vertex of γ sent to the interior of $\underline{\mathbf{M}}$ is at least trivalent, and as mentioned above, has exactly two edges not attached to a vertex sent to the corner of $\underline{\mathbf{M}}$. The tropical balancing condition is satisfied at these vertices.

(B) As mentioned above, there is at most one vertex sent to the one-dimensional boundaries of $\underline{\mathbf{M}}$, and such a vertex must have exactly one edge not attached to a vertex sent to the corner of $\underline{\mathbf{M}}$. If the vertex is sent to the upper boundary, the sum of the derivatives of edges entering this vertex is $(-1,0)$. If the vertex is sent to the lower boundary, the sum of the derivatives of edges entering this vertex is $(0,-1)$. If the constrained edge of γ is above the ray $-y$, this vertex must go to the top boundary of $\underline{\mathbf{M}}$, and if the constrained edge is below the ray $-y$, this vertex must go to the bottom boundary of $\underline{\mathbf{M}}$.

(C) At vertices v of γ sent to the corner of $\underline{\mathbf{M}}$, the configuration γ_v of edges leaving such a vertex is one of the configurations with nonzero coefficient n_{γ_v} in F. In this case, the relative Gromov–Witten invariants used counts curves with genus specified by γ_v using equation (1).

We shall now argue that the tropical curves described above are the only contributors to n_Γ. A contributing tropical curve γ must be rigid after restricting the constrained edge.[9] As scaling $\underline{\mathbf{M}}$ acts on tropical curves, and can be thought of as acting separately on different components of γ minus the inverse image of the corner of $\underline{\mathbf{M}}$, it follows that γ minus the inverse image of the corner of $\underline{\mathbf{M}}$ consists of some number of rays emanating from the corner of $\underline{\mathbf{M}}$, and one connected 'interesting' component that includes the constrained edge. The need for γ to be rigid and the balancing condition at vertices in the interior of $\underline{\mathbf{M}}$ implies that vertices in the interior of $\underline{\mathbf{M}}$ must be at least trivalent, and have at least two edges not attached to the corner of $\underline{\mathbf{M}}$, similarly, we may discount vertices on a boundary of $\underline{\mathbf{M}}$ with all edges constrained to the boundary.

With the above constraints on the tropical curve γ, the only way to get the virtual dimension of curves with tropical part γ to be 0 is for γ to satisfy the conditions above. In particular, the (complex) virtual dimension is equal to the sum of the virtual dimensions of the moduli spaces corresponding to vertices minus the number of internal edges, minus the number of constrained edges. To satisfy the balancing condition, each vertex in the interior must have at least 2 edges not attached to the corner of $\underline{\mathbf{M}}$. The only way to get virtual dimension 0 is for the virtual dimension corresponding to the corner vertices to be 0, and for the interesting component of γ minus all edges attached to the corner of $\underline{\mathbf{M}}$ to be a linear graph, with one end the constrained edge, and the other end either an unconstrained infinite edge, or a vertex sent to the boundary of $\underline{\mathbf{M}}$. Such a configuration corresponds to an unconstrained moduli space of virtual dimension 1. The balancing condition implies that the constrained edge and the other end of this linear graph must both be on the same side of the ray $-y$.

This completes the explanation of why the only tropical curves that contribute to n_Γ satisfy the above conditions.

As proved in [18], the contribution of each tropical curve γ to $n_\Gamma/|\operatorname{Aut}\Gamma|$ is determined by taking a fiber product of relative Gromov–Witten invariants corresponding to vertices of γ over spaces corresponding to internal edges of γ, and dividing the result by automorphisms of γ. For this formula, the relative Gromov–Witten invariants for a vertex v have contact data labeled by γ_v, the tropical curve obtained by extending all edges leaving v to be infinite. To describe the gluing formula precisely, all moduli spaces and the spaces over which we take fiber products must be described in terms of exploded manifolds. For fiber products of exploded manifolds to be reflected correctly using cohomology, it is necessary to use refined cohomology, defined in [10], to define the correct relative Gromov–Witten invariants. To avoid refined cohomology, we shall use the following simplified gluing formula that only applies in rather special cases:

Theorem 2.1 (Simplified gluing formula). *Suppose that a tropical curve γ has no edges with both ends attached to a single vertex, and that the following algorithm terminates with all edges of γ labeled rigid:*

- *Label any constrained infinite edge of γ 'rigid', and orient it to be incoming.*

[9] This is because the virtual fundamental class is empty wherever the real dimension of the moduli space of tropical curves exceeds the complex virtual dimension of the moduli space of curves.

- *Suppose that a vertex v satisfies the following: the expected dimension of curves with contact data γ_v with rigid edges incoming (or constrained), and all other edges outgoing (or unconstrained) is 0. Then label the remaining edges leaving v as rigid, and orient them away from v. Then repeat this step until all edges are labeled rigid.*

If the above algorithm terminates with all edges of γ labeled rigid, the contribution of γ to Gromov–Witten invariants (in this case $n_\Gamma / \lvert \operatorname{Aut} \Gamma \rvert$) is

$$\frac{1}{\lvert \operatorname{Aut} \gamma \rvert} \prod_v n_{\gamma_v}$$

where n_{γ_v} is the relative Gromov–Witten invariant that counts curves with contact data specified by γ_v with edges oriented (hence constrained or unconstrained) as specified by the above algorithm.

This theorem follows from the gluing formula proved in [18]: our conditions ensure that calculation of this contribution of γ using equation (1) of [18] can be performed using only top-dimensional forms, which simplifies calculations because top-dimensional refined cohomology is 1-dimensional, and equal to top-dimensional usual cohomology.

The above simplified gluing formula works for the curves γ contributing to n_Γ — all edges attached to the corner of $\underline{\mathbf{M}}$ are oriented to leave the corner, and all other edges are oriented away from the constrained edge.

(A) If v is a vertex sent to the interior of $\underline{\mathbf{M}}$, n_{γ_v} indicates the corresponding zero-genus Gromov–Witten of the exploded manifold $\mathbf{T}^2$, or alternately the zero-genus relative Gromov–Witten invariant of any compact two-complex-dimensional toric manifold relative to its toric boundary strata, with contact data given by γ_v. In the case that v is trivalent and we choose the standard complex structure, the (virtual and actual) moduli space is acted on freely and transitively[10] by the $(\mathbb{C}^*)^2$–action on the toric manifold, implying that $n_{\gamma_v} = \lvert \alpha \wedge \beta \rvert$ where α and β are the two incoming edges.

In the case that γ_v has several edges coming from the corner of $\underline{\mathbf{M}}$ (these edges must therefore be parallel), we may compute the relative Gromov–Witten invariant n_{γ_v} by constraining these edges to be at different locations tropically so that there is only one possible tropical curve that contributes to the count, and this tropical curve is trivalent. Then using our simplified gluing formula to glue the trivalent invariants gives

$$n_{\gamma_v} = \prod_i \lvert \alpha \wedge \beta_i \rvert$$

where the β_i are the derivatives of γ_v on the edges from the corner of $\underline{\mathbf{M}}$ and α is the derivative of γ_v on one of the other two edges of γ_v.

(B) If v is a vertex sent to a boundary of $\underline{\mathbf{M}}$, n_{γ_v} indicates the relative Gromov–Witten invariant of $\mathbf{T}$ times the explosion of $(\mathbb{C}P^1, 0)$, or alternatively $\mathbb{C}P^1 \times \mathbb{C}P^1$ relative to 3 of its 4 toric boundary divisors, where all edges of n_{γ_v} are labelled incoming (or constrained), and γ_v is interpreted as contact data so that (constrained) contact with a point on the middle divisor corresponds to $(-1,0)$ or $(0,-1)$ if v is respectively on the top or bottom boundary of $\underline{\mathbf{M}}$. In this case, the above formula for n_{γ_v} also holds so long as the sum of vectors in γ_v adds up to $(-1,0)$ or $(0,-1)$ respectively. The case that v is univalent is

[10]Of course, $(\mathbb{C}^*)^2$ only acts transitively on the noncompact moduli space of curves with no components contained in toric boundary strata — in the case of the exploded manifold $\mathbf{T}^2$, this $(\mathbb{C}^*)^2$–action is replaced by a $\mathbf{T}^2$–action which is free and transitive on the (already compact) moduli space.

readily calculated directly to give a unique curve. In cases with more incoming edges, we may constrain these edges so that the only contributing tropical curve has one univalent vertex and otherwise trivalent vertices for which the above formula follows from our simplified gluing formula.

(C) If v is sent to the corner of $\underline{\mathbf{M}}$, n_{γ_v} is the relative Gromov–Witten invariant we described in the definition of F.

Let γ_0 indicate the disjoint union of γ_v for all vertices v of γ sent to the corner of $\underline{\mathbf{M}}$. If γ has its constrained edge above the ray $-y$, the coefficient of Γ in $n_{\gamma_0}/|\mathrm{Aut}\,\gamma_0|\gamma_0\overleftarrow{y}\,y$ is equal to the above contribution of γ to $n_\Gamma/|\mathrm{Aut}\,\Gamma|$, (and this coefficient is 0 if the constrained edge is below the ray $-y$). This observation is easy to prove after noting that $|\mathrm{Aut}\,\gamma_0|/|\mathrm{Aut}\,\gamma|$ is equal to the number of ways of choosing which edges from γ_0 should interact with the vector decorated with an arrow in order to be left with Γ. Similarly, if γ has its constrained edge below the ray $-y$, the above contribution is the coefficient of Γ in $n_{\gamma_0}/|\mathrm{Aut}\,\gamma_0|\,y\,\overrightarrow{y}\,\gamma_0$. Equation (5) follows, with the sum over all tropical curves with constrained edge below $-y$ replaced by terms in $y\,\overrightarrow{y}\,e^F$, and the sum over all tropical curves with constrained edge above $-y$ replaced by terms in $e^F\overleftarrow{y}\,y$.

3. Reconstructing the absolute Gromov–Witten invariants from the relative invariants

To relate our relative Gromov–Witten invariants to the absolute Gromov-Witten invariants of the n-fold blowup of $\mathbb{C}P^2$, we shall consider a degeneration of this manifold — we will then use our relative Gromov-Witten invariants as an essential ingredient in a tropical gluing formula reconstructing the Gromov–Witten invariants of our n-fold blowup. As a warmup, we will first describe a corresponding degeneration of $\mathbb{C}P^2$.

Consider the moment map of $\mathbb{C}P^2$ with the standard torus action. We can subdivide this moment map as below using rays in the directions $(-1,1)$, $(1,0)$, and $(k,-1)$ for all integers $k\in[-n+1,1]$, while ensuring that all the downward pointing rays intersect the lower edge of the moment map.

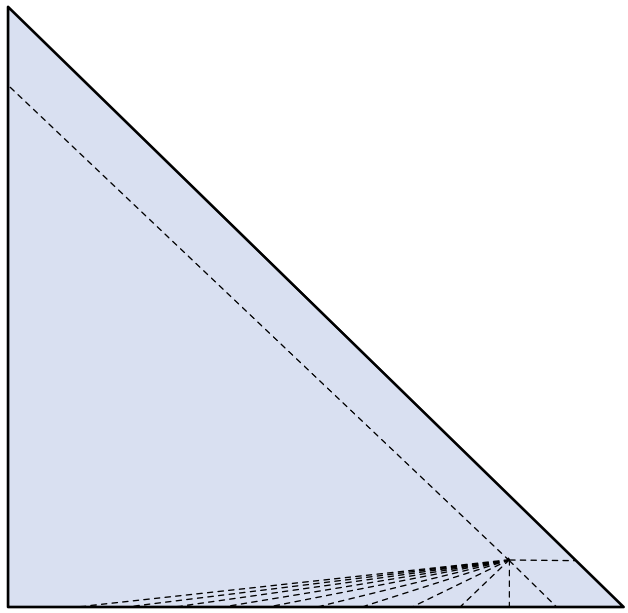

There exists a complex toric degeneration of $\mathbb{C}P^2$ with degenerate fiber consisting of a union of toric manifolds with moment maps the pieces of the subdivided triangle above, glued along boundary divisors as above. The particular toric degeneration chosen shall not be important for us, and it is not important that this degeneration is toric, only that it is log smooth.

One choice of degeneration is constructed as follows: Consider a dual polytope P to the above set of rays — in other words, consider a convex polygon P with edges orthogonal to the above rays. Choose P to have integral vertices. Then consider the toric partial compactification X of $(\mathbb{C}^*)^3$ given by the fan consisting of

- the cone over P, when P is placed in the plane with first coordinate 1,
- the cone formed by $(0,-1,0),(0,1,1)$ and the top right vertex of P,
- the cone formed $(0,1,1)$, $(0,0,-1)$ and the other righthand vertex of P,
- the cone formed by $(0,0,-1)$, $(0,-1,0)$ and top left vertex of P,
- the cone formed by the righthand face of P and $(0,1,1)$,
- the cone formed by the top left face of P and $(0,-1,0)$,
- and the cones formed by $(0,0,-1)$ and all the lower faces of P.

Projection of X to the first coordinate gives the required toric degeneration.

This degeneration $\pi : X \longrightarrow \mathbb{C}$ is log smooth when $\mathbb{C}$ is given the log structure from the divisor 0, and X is given the log structure from the divisor $\pi^{-1}(0)$. This divisor is the union of the toric divisors of X corresponding to all rays in the fan of X apart from $(0,-1,0),(0,1,1),(0,0,-1)$. To verify that X with this log structure is log smooth, note that the cones formed using two of the directions $(0,-1,0)$, $(0,1,1)$ and $(0,0,-1)$ may be transformed (using an invertible $\mathbb{Z}$-linear transformation) to standard quadrants, and that each time a cone is formed using a face of P and one of these directions, the configuration formed by the linear plane containing the face of P and the extra direction may be transformed to the standard configuration consisting of the plane spanned by the first two coordinates and $(0,0,1)$.

We may blow up X along n complex submanifolds intersecting the singular divisor transversely at n points distributed within the n triangles in the above subdivided moment map picture. By restricting the family π to some neighborhood D of $0 \in \mathbb{C}$, we may assume that these n complex submanifolds are transverse to all fibers of π, so the resulting blown-up family $\pi' : X' \longrightarrow D$ is also a log smooth family.

As explained in [12] or [13], we may apply the explosion functor to $\pi' : X' \longrightarrow D$ to obtain a smooth family of exploded manifolds.

$$\operatorname{Expl} \pi' : \operatorname{Expl} X' \longrightarrow \operatorname{Expl} D$$

Each exploded manifold has a tropical part consisting of a union of polytopes. The tropical part of $\operatorname{Expl} X'$ is the cone over P, the tropical part of $\operatorname{Expl} D$ is the half line, and the tropical part of $\operatorname{Expl} \pi'$ is the projection defining our toric degeneration. The definition of a smooth family in [12] contains a condition of being surjective on integral vectors — this condition is satisfied due to our choice of the corners of P having integer coordinates. It is easy to choose a symplectic form taming the complex structure of $\operatorname{Expl} X$ in the sense of [16]. After a choice of symplectic representation of our blowup, this gives a symplectic form on $\operatorname{Expl} X'$ taming the complex structure. As the tropical part of $\operatorname{Expl} X'$ may be embedded in a quadrant of $\mathbb{R}^3$, the results of [16] imply that Gromov compactness holds in our family $\operatorname{Expl} \pi'$, and we may define Gromov–Witten invariants as in [14, 17].

Some fibers of $\operatorname{Expl} \pi'$ are n-fold blowups of $\mathbb{C}P^2$, another fiber is an exploded manifold $\mathbf{B}$ that has tropical part P. We shall show how to calculate the Gromov–Witten invariants of $\mathbf{B}$. As the Gromov–Witten invariants of exploded manifolds do not change in connected families of exploded manifolds, [14, 17], the Gromov–Witten invariants of $\mathbf{B}$ correspond to the Gromov–Witten invariants of $\mathbb{C}P^2$ blown up at n points.

As explained in [18], Gromov–Witten invariants of $\mathbf{B}$ decompose into a sum of contributions from rigid tropical curves in the tropical part $\underline{\mathbf{B}}$ of $\mathbf{B}$. Below is a picture of $\underline{\mathbf{B}}$, and some tropical curves in $\underline{\mathbf{B}}$. We shall see that the left and righthand curves both contribute 1 to the Gromov–Witten count of curves and the middle picture does not contribute, because it is not rigid — actually this middle tropical curve deforms to the lefthand tropical curve. Confusingly, when using the non-generic complex structure on $\mathbf{B}$ from the toric model described above, there are no genuine holomorphic curves with tropical part given by the left and righthand pictures, but there are holomorphic curves with tropical parts such as those pictured in the middle that deform to the lefthand picture (and something similar happens for the righthand picture). When a generic complex structure on $\mathbf{B}$ is used, there are unique holomorphic curves in $\mathbf{B}$ with tropical parts the left and righthand pictures, but there does not exist any holomorphic curve with tropical part given by the middle picture.

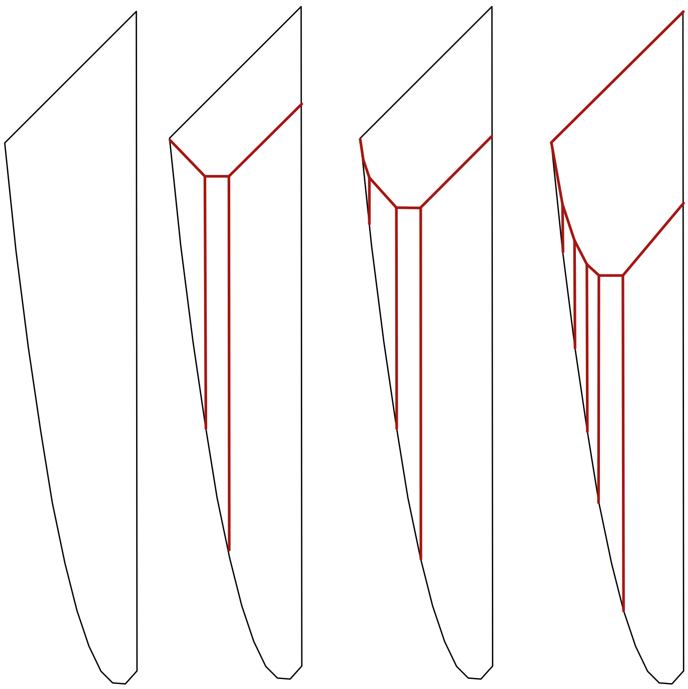

Use H to denote the homology class represented by the pullback of a line to the n-fold blowup of $\mathbb{C}P^2$, and let E_i be the homology class of the ith exceptional divisor. As proved in [10], our exploded manifold $\mathbf{B}$ has the same DeRham cohomology as $\mathbb{C}P^2$ blown up at n points, so the same classes make sense in $\mathbf{B}$ — we can choose a representative for H with tropical part the lefthand corner and top boundary of $\underline{\mathbf{B}}$, and a representative for each E_i with tropical part the ith corner at the bottom of $\underline{\mathbf{B}}$. With these choices, we can measure the homology class represented by curves using their intersection with our representatives, and talk of the individual contributions of vertices of tropical curves to the overall homology class. From the left, the first and second curves are rational curves in the class $H - E_3 - E_5$, and the last is a rational curve representing $2H - E_1 - E_2 - E_3 - E_4 - E_6$.

As proved in [18], the Gromov–Witten contribution of each rigid tropical curve in $\underline{\mathbf{B}}$ is determined by taking a fiber product of relative Gromov–Witten invariants corresponding to its vertices. The vertices of our tropical curves now come in the following types.

(A) The vertices v in the interior of $\underline{\mathbf{B}}$ use the Gromov–Witten invariants of $\mathbf{T}^2$, or alternately the relative Gromov–Witten invariants of any compact 2-complex-dimensional toric manifold relative to its toric boundary divisors. We described these invariants in the previous section.

In particular, the corresponding curves have genus 0, and homology class determined by the contact data from γ_v, the tropical curve obtained by extending all edges attached to v to be infinite. The virtual (complex) dimension of this moduli space is the valence of v minus 1. We shall only need to know the Gromov–Witten invariant in the case that γ_v has one unconstrained (outgoing) edge, and constrained (incoming) edges with derivatives α, and β_i where all the β_i are parallel. In this case,

$$n_{\gamma_v} = \prod_i |\alpha \wedge \beta_i|$$

as described in the last section.

(B) The vertices v sent to the one-dimensional boundaries of $\underline{\mathbf{B}}$ use the Gromov–Witten invariants of $\mathbf{T} \times \mathrm{Expl}(\mathbb{C}P^1, 0)$, or alternatively the relative Gromov–Witten invariants of $(\mathbb{C}P^1)^2$ relative 3 of its 4 toric boundary divisors. These were also described in the previous section, but now, constrained simple contact with the middle divisor corresponds to $(0, -1)$ for each bottom boundary of $\underline{\mathbf{B}}$, $(-1, 0)$ for the top boundary of $\underline{\mathbf{B}}$, and $(1, 1)$ for the righthand boundary of $\underline{\mathbf{B}}$. As described in the last section, the derivative of all edges entering v adds up to some multiple k of the respective vector $(0, -1)$, $(-1, 0)$, or $(1, 1)$, and the virtual dimension is equal to the valence of v plus $k - 1$.

As it turns out, the only such relative Gromov–Witten invariant we need has v on the righthand boundary, and γ_v one incoming edge with derivative $(1, 1)$. In this case, $n_{\gamma_v} = 1$.

(C) For vertices v sent to the lefthand corner of $\underline{\mathbf{B}}$, we use curves in the exploded manifold $\mathbf{M}$ described in the previous section, or alternately the relative Gromov–Witten invariants of the (M, N) described in the first section. The relevant Gromov–Witten invariants are those encoded by the generating function F_n.

(D) For vertices sent to either of the two righthand corners of $\underline{\mathbf{B}}$, we use the relative Gromov–Witten invariants of some toric manifold relative to two of its toric boundary divisors. In the picture of the subdivided moment map at the start of this section, these two toric manifolds are have moment map the two righthand cells, and we take invariants relative to the dotted boundary divisors. Although these two toric manifolds are different, what they have in common is that the complex dimension of the moduli space curves with unconstrained contact data γ_v is at least the valence of v. Our only hope of getting rigid curves is to count curves with every edge of γ_v incoming, or constrained. We may calculate these invariants tropically by constraining the edges of γ_v generically so that all contributing tropical curves avoid the corner, so these invariants are determined by the gluing formula and the invariants from (A) and (B) above.

In particular, we shall use that for these corners, $n_{(1,1)} = 1$.

(E) For vertices sent to one of the n (difficult to distinguish) bottom corners of $\underline{\mathbf{B}}$, we use the relative Gromov–Witten invariants of $\mathbb{C}P^2$ blown up at one point, relative to two lines, L_1, L_2. For the kth of these corners, (unconstrained) contact with L_1 corresponds to $(-1, n - k)$ and (unconstrained) contact with L_2 corresponds to $(1, k + 1 - n)$.

In this case, we determine the homology class of the relevant curves using their intersection with the exceptional sphere and the contact data γ_v. For topological reasons, the contact data γ_v satisfies the balancing condition that the sum of all derivatives at edges leaving v is $(0,d)$ for some nonnegative integer d. The virtual dimension is the valence of v plus the genus, plus $(d-1)$ minus the intersection with the exceptional sphere. In fact, d minus the intersection with the exceptional sphere is the intersection with the strict transform of a line passing through the blown up point, and must therefore be nonnegative. Therefore the virtual (complex) dimension is at least the valence of v minus 1.

In the next section, we show that, apart from the exceptional sphere (which has tropical part a single point), the rigid curves with unconstrained contact data have genus 0, contact data $(0,d)$, and intersect the exceptional sphere d times. We shall also see that the corresponding relative Gromov–Witten invariants are $n_{(0,d)} = \frac{(-1)^{d+1}}{d^2}$. For us, the useful consequence of this shall be that for this vertex,

$$n_{\{(a,b),(a,b+d)\}} = \binom{a}{d}$$

where the above counts curves with contact data an incoming (constrained) edge in direction (a,b), and an outgoing (unconstrained) edge in direction $(a,b+d)$. If we constrain the incoming edge (a,b) to be above the ray $-(a,b)$, the above relative Gromov–Witten invariant decomposes into contributions from tropical curves with some rigid edges coming up from the corner with derivatives summing to $(0,d)$. Importantly for us, in the case that $a=1$ and $d>1$, all these contributions cancel, even though the contributions of the individual tropical curves do not vanish.

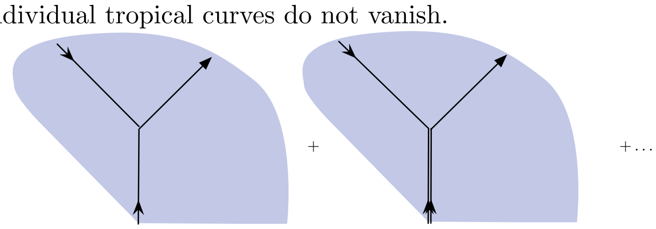

Similar cancellations happen when we compute the contribution of tropical curves with pieces looking like above: when we sum over all possibilities of an 'incoming' edge $(1,b)$ interacting with rigid edges leaving the bottom corner and leaving as $(1,b+d)$, all contributions cancel whenever $d>1$. This calculation holds regardless of the consideration of whether $(1,b)$ is actually travelling in a valid direction for being an 'incoming edge' to our corner.

With the above understood, we can identify which tropical curves γ contribute to counts of rigid curves in $\mathbf{B}$. First, there are the n exceptional spheres, which correspond to the rather uninteresting tropical curves consisting of a single point sent to one of the n bottom corners of $\underline{\mathbf{B}}$.

In the case that γ is not a single point, consider a connected component of γ minus the inverse image of the lefthand corner of $\underline{\mathbf{B}}$. Let us argue that this connected component must have a univalent vertex v on the closure of the righthand boundary of $\underline{\mathbf{B}}$ with an incoming edge of derivative $(1,1)$. Our balancing conditions imply that our connected component must have some vertex v on the closure of the righthand boundary of $\underline{\mathbf{B}}$ with derivatives of incoming edges adding up to (a,b), where $a>0$. The only way we can get rigid curves at such a vertex v by placing constraints on the incoming edges is if all edges are constrained, and their derivatives sum to $(1,1)$. As there are no types of vertices on the closure of the righthand boundary of $\underline{\mathbf{B}}$ that can correspond to rigid curves without constraining all edges, there is no way to obtain a constrained edge entering v with derivative

$(0,k)$, therefore our vertex v on the righthand edge must have a single incoming edge with derivative $(1,1)$.

Apart from vertices at the lefthand corner, any contributing space of curves with unconstrained contact data γ_v has complex virtual dimension at least the valence of v minus 1. Therefore, if one edge is connected to our vertex on the righthand boundary, (and hence is unconstrained) all other edges must be constrained. Similarly, all the other edges attached to a vertex on the other end of these edges must be constrained. It follows that

- Each connected component of γ minus vertices at the lefthand corner of $\underline{\mathbf{B}}$ must be a tree with a unique vertex on the closure of the righthand boundary of $\underline{\mathbf{B}}$.
- The only way to obtain nonzero Gromov–Witten invariants is to use the relative invariants at each vertex specified by orienting edges to point towards the vertices on the closure of the righthand boundary of $\underline{\mathbf{B}}$.
- There are no vertices on the (interior of the) top or bottom boundaries of $\underline{\mathbf{B}}$, as to satisfy the above dimension constraint, these vertices would need to have edges all contained in the boundary of $\underline{\mathbf{B}}$, and therefore would not be part of a rigid tropical curve. Similarly, there are no bivalent vertices in the interior of $\underline{\mathbf{B}}$.
- The balancing condition implies that each connected component of γ minus vertices at the lefthand corner of $\underline{\mathbf{B}}$ must have a connected bivalent subgraph with all edges having derivative in the form $(1,k)$, and that all other edges must have derivative in the form $(0,d)$. In particular, this connected component must have a unique edge coming from the lefthand corner of $\underline{\mathbf{B}}$, and the derivative of this edge must be $(1,k)$ for some k.
- The edges in direction $(0,d)$ must come from the bottom corners of $\underline{\mathbf{B}}$. If a vertex of v of γ has an edge entering with derivative $(1,k)$ and an edge exiting with derivative $(1,k+d)$ where $d>1$, then the contribution of γ to Gromov–Witten invariants will be 0 if v is at one of the bottom corners of $\underline{\mathbf{B}}$, or cancel with the contributions of other curves that are the same apart from the other edges entering v.
- For γ to be connected, it must have a unique vertex v_0 at the lefthand corner of $\underline{\mathbf{B}}$.
- Using the simplified gluing formula from Theorem 2.1 and the orientation of γ above, the contribution of γ to Gromov–Witten invariants is
$$\frac{1}{|\operatorname{Aut}\gamma|}\prod_v n_{\gamma_v} \ .$$
Apart from the unique vertex v_0, whose invariant $n_{\gamma_{v_0}}$ is equal to $|\operatorname{Aut}\gamma_{v_0}|$ times the coefficient of γ_{v_0} in F, all other $n_{\gamma_v}=1$ for tropical curves γ whose contribution to Gromov–Witten invariants do not cancel.

To summarize, apart from the exceptional curves — which correspond to particularly uninteresting tropical curves consisting of a single point mapping to one of the bottom vertices of $\underline{\mathbf{B}}$ — the only tropical curves that contribute to Gromov–Witten invariants of $\mathbf{B}$ are those that leave the left corner of $\underline{\mathbf{B}}$ with edges having derivative $\{(1,1-m_1),\dots,(1,1-m_k)\}=\Gamma$ with a corresponding nonzero coefficient $n_\Gamma/|\operatorname{Aut}\Gamma|$ in F. Each of these edges must go on to interact with m_i edges with derivative $(0,1)$ coming from m_i different bottom corners of $\underline{\mathbf{B}}$, before hitting the closure of the righthand boundary of $\underline{\mathbf{B}}$ with derivative $(1,1)$. The virtual number of curves corresponding to each such tropical curve is n_Γ. Their degree and genus is the degree and genus of Γ (as defined in equation (4)), and their intersection with E_i is the number of edges leaving the ith bottom corner of $\underline{\mathbf{B}}$.

Consider all tropical curves γ with γ_{v_0} isomorphic to $\Gamma = \{(1, 1-m_1), \dots, (1, 1-m_k)\}$. Such tropical curves contribute the following term to the generating function, G_n, encoding Gromov–Witten invariants of $\mathbf{B}$:

$$\frac{n_\Gamma}{|\mathrm{Aut}\,\Gamma|} \prod_{i=1}^{k} x^{m_i-3} q^H \sigma_{m_i}$$

where the exponent of x encodes genus -1, the exponent of q encodes homology class, and σ_m indicates the mth elementary symmetric function in the variables $q^{-E_1}, \dots, q^{-E_n}$. This implies equation (3), our formula for Gromov-Witten invariants of the n-fold blowup of $\mathbb{C}P^2$ in terms of the relative Gromov-Witten invariants encoded in F.

4. Relative Gromov–Witten invariants of $\mathbb{C}P^2$ blown up at one point

In this section, we calculate the relative Gromov–Witten invariants required in the previous section for the bottom corners of $\underline{\mathbf{B}}$. The required invariants are the relative Gromov–Witten invariants of $\mathbb{C}P^2$ blown up at 1 point, relative to two lines L_1, L_2.

One way to picture $\mathbb{C}P^2$ blown up at a point symplectically is as a singular Lagrangian torus fibration with base pictured below. The size of the little removed triangle represents the size of the symplectic ball removed to do a symplectic blowup; the remaining polytope should be regarded as glued along the two faces of this little removed triangle so that it has an integral-affine structure with a singularity at the top point of the little removed triangle. This singularity in the integral-affine structure reflects a focus-focus singularity in the Lagrangian torus fibration above it — in other words, the torus fiber pinches to become a sphere which intersects itself once above this point. There are also elliptic singularities along the 3 edges of this picture, as is usual for moment-map pictures. We are interested in 4 holomorphic spheres in this picture. Over the left and righthand boundaries are spheres L_1 and L_2 that are lines from $\mathbb{C}P^2$. Running down the glued-together edges of the little removed triangle is the exceptional sphere, and over the bottom boundary is a sphere L_3 that is the strict transform of a line passing through the point we blew up.

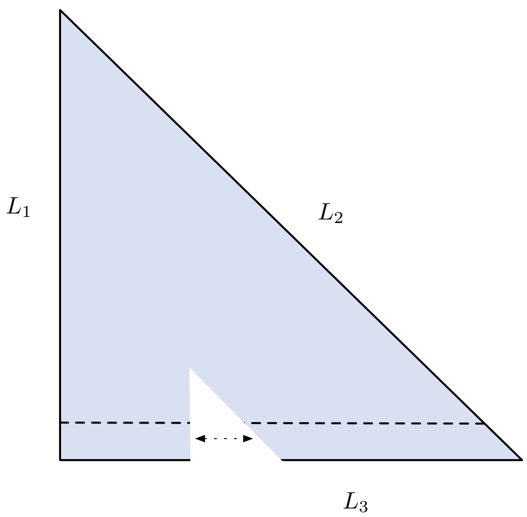

We need the relative Gromov–Witten invariants of this space relative to L_1 and L_2. These relative invariants are most effectively computed tropically by making a degeneration of this space into the two pieces above and below the dotted line. After making a symplectic cut along the dotted line, the bottom piece is $(\mathbb{C}P^1)^2$. We need — and have already encountered — the relative Gromov–Witten invariants of $(\mathbb{C}P^1)^2$ relative to 3 of its 4 toric boundary divisors, the left, right, and dotted top ones in the above picture. The top piece is again the blowup of $\mathbb{C}P^2$ at a point, but now we need its Gromov–Witten invariants relative to L_1, L_2, and L_3.

Let us work out the Gromov–Witten invariants of the above space relative L_1, L_2, and L_3. The tropical part of the explosion, $\mathbf{A}$, of the above space relative to $L_1 \cup L_2 \cup L_3$ is pictured below, with a tropical curve.

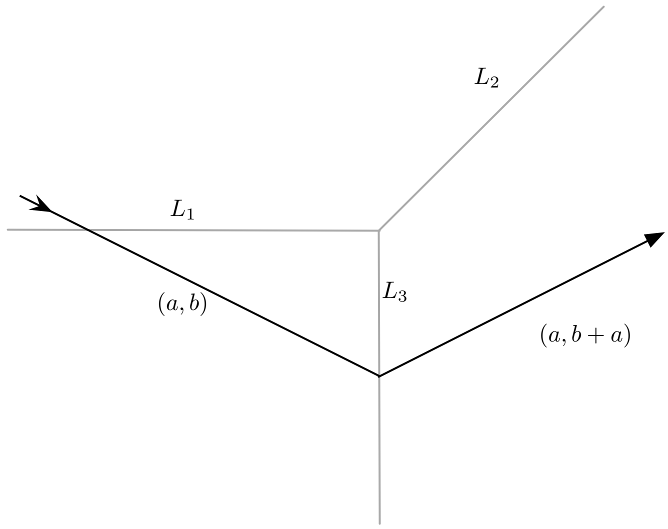

To translate tropical curves in the above picture to contact data with L_i, unconstrained contact with L_1, L_2 or L_3 corresponds to an end with outgoing derivative $(-1, 0)$, $(1, 1)$, or $(0, -1)$ respectively. The homology class of a curve is determined by its contact data with L_i, and the virtual dimension of a curve with contact data Γ is the number of infinite ends of Γ plus the genus of the curve minus 1. The contact data obeys a balancing condition: the sum of the derivative of all edges (oriented outgoing) is some multiple of $(0, 1)$. The case of empty contact data — or a tropical curve with no infinite edges — corresponds to a curve representing zero in homology, therefore we may ignore this case, because such curves of genus 1 are not stable. Therefore, for curves with unconstrained contact data to be rigid, they must be spheres with contact data consisting of one outgoing edge with derivative $(0, d)$.

To calculate the contribution of a tropical curve γ in $\underline{\mathbf{A}}$ to the Gromov–Witten invariants of $\mathbf{A}$, we need to know what relative Gromov–Witten invariants to associate to each vertex. A vertex v sent to the origin in $\underline{\mathbf{A}}$ uses the Gromov–Witten invariants of $\mathbf{A}$. A vertex sent to the interior of one of any of the two-dimensional strata in $\underline{\mathbf{A}}$ uses the Gromov–Witten invariants of $\mathbf{T}^2$, or equivalently the relative Gromov–Witten invariants of any two-complex-dimensional toric manifold relative its toric boundary divisor. These relative Gromov–Witten invariants were discussed in the previous sections. A vertex sent to the interior of any of the one-dimensional strata of $\underline{\mathbf{A}}$ uses the Gromov–Witten invariants of $\mathbf{T} \times \mathrm{Expl}(\mathbb{C}P^1, \{0, \infty\})$, or equivalently the relative Gromov–Witten invariants of any two-complex-dimensional toric manifold relative to its toric boundary divisors.

In other words, vertices at every point of $\underline{\mathbf{A}}$ apart from the origin use the Gromov–Witten invariants of the same space, $\mathbf{T}^2$. There is, however, a difference. For any vertex v of γ, recall that we produce a tropical curve γ_v in the plane by extending all edges leaving v to be infinite. Everywhere but the strata corresponding to L_3, we interpret γ_v directly as a tropical curve in the tropical part of

$\mathbf{T}^2$ to define contact data. On the strata of $\underline{\mathbf{A}}$ corresponding to L_3 however, we need to exchange every edge of γ_v with outgoing derivative (a,b) with an edge with outgoing derivative $(a, b - \max\{a,0\})$. This implies that tropical curves in $\underline{\mathbf{A}}$ obey the usual balancing condition if we give $\underline{\mathbf{A}}$ the singular integral-affine structure obtained by cutting the plane along the line corresponding to L_3, and gluing the left and righthand sides of this cut so that a vector (a,b) on the left side corresponds to the vector $(a, b+a)$ on the right. In particular, the tropical curve drawn in the picture above corresponds to a 'straight line'.

We can now easily compute some Gromov–Witten invariants of $\mathbf{A}$. Let $n_{v_1,\dots,v_n;w}$ indicate the Gromov–Witten invariant of $\mathbf{A}$ that counts zero genus curves with contact data consisting of incoming (constrained) edges with derivative v_i and an outgoing (unconstrained) edge with derivative w.

Consider computing $n_{(0,1);(1,1)}$. The top tropical curve pictured below is the unique curve contributing to $n_{(0,1);(1,1)}$ when the incoming edge is constrained in the upper half plane. The bottom tropical curve is the only curve that contributes when the incoming edge is constrained in the lower half plane.

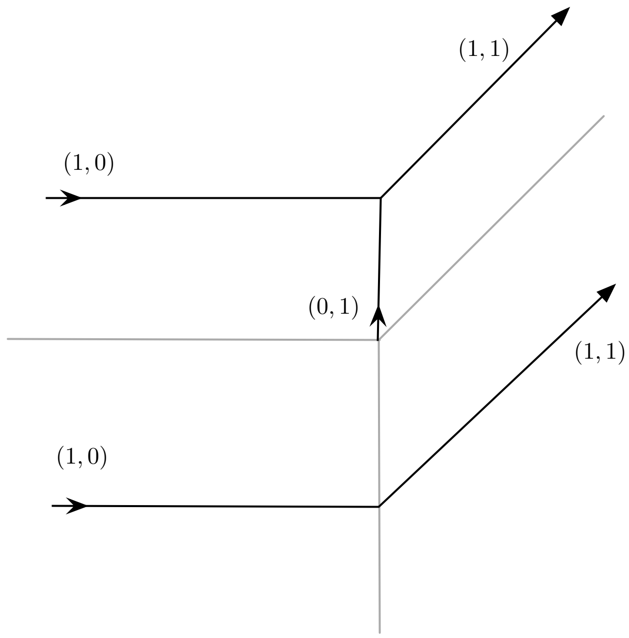

Therefore, using our simplified gluing formula from Theorem 2.1,

$$n_{(1,0);(1,1)} = 1 = n_{;(0,1)} \ .$$

For $a > 0$, we can equate the two ways pictured below of calculating $n_{(a,b),(0,1);(a,c)}$.

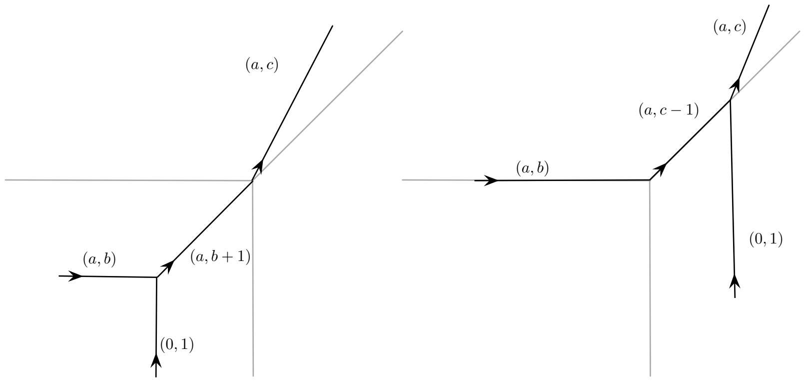

So

$$n_{(a,b+1);(a,c)} = n_{(a,b);(a,c-1)} \ .$$

In other words, $n_{(a,b);(a,b+d)}$ depends only on a and d.

Calculating $n_{(a,b);(a,b+a)}$ by constraining the incoming edge below the ray $(-a,-b)$ gives that

$$n_{(a,b);(a,b+a)} = 1$$

and

$$n_{(a,b);(a,b+d)} = 0 \text{ if } d > |a| \ .$$

Similarly, calculating constraining the incoming edge above the ray $(-a,-b)$ gives that

$$n_{(a,b);(a,b+d)} = 0 \text{ if } d < 0 \ .$$

If $0 \le d \le a$, we can equate the two ways of calculating $n_{(a,a+1-d),(1,0);(a+1,a+1)}$ pictured below.

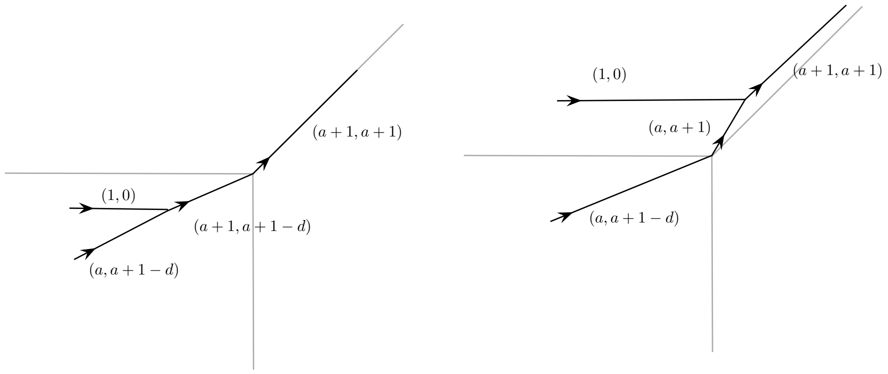

Therefore

$$(a+1-d)n_{(a+1,a+1-d);(a+1,a+1)} = (a+1)n_{(a,a+1-d);(a,a+1)} \ .$$

Starting with the case

$$n_{(a,b);(a,b+a)} = 1 = \binom{a}{a}$$

induction on $d \ge a$ using the above equation gives

$$n_{(a,b);(a,b+d)} = \binom{a}{d} \ . \tag{6}$$

We still need to compute $n_{;(0,d)}$. Consider $p(x) = n_{(x,0);(x,d)}$. For $x \ge 0$, we can compute $p(x)$ by restricting the incoming edge $(x,0)$ in the upper half plane. The gluing formula for $p(x)$ uses tropical curves in the form pictured below. In this diagram, the thick edge indicates some number of edges with upward pointing derivatives adding to $(0,d)$.

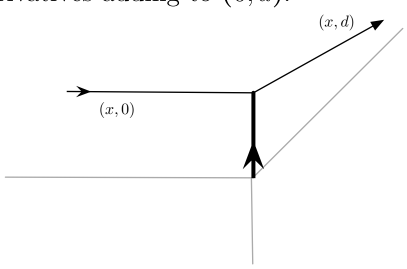

A diagram with k such upward pointing edges contributes some fixed multiple of x^k to $p(x)$, so $p(x)$ is a polynomial in x of degree at most d. The only possibility for a single such edge is when its derivative is $(0,d)$. This tropical curve contributes

$xdn_{;(0,d)}$ to $p(x)$, therefore the coefficient of the linear term in $p(x)$ is $dn_{;(0,d)}$. We already know that $p(x) = 0$ for $0 \leq x < d$, and $p(d) = 1$, therefore

$$p(x) = \frac{1}{d!}\prod_{i=0}^{d-1}(x-i) \ .$$

The coefficient of the linear term of $p(x)/d$ then gives us $n_{;(0,d)}$.

$$n_{;(0,d)} = \frac{1}{d!d}\prod_{i=1}^{d-1}(-i) = \frac{(-1)^{d-1}}{d^2} \tag{7}$$

Although we have only argued that the above equation holds for $d \geq 1$, a very similar argument also gives that $n_{;(0,d)} = (-1)^{d+1}/d^2$ for $d < 0$. Of course, the equations determining $n_{;(0,d)}$ are massively overdetermined. A combinatorially-talented thinker could deduce the above equations for $d > 0$ and the corresponding formula for $n_{(a,b);(a,b+d)}$ simply from knowing that $n_{;(0,1)} = 1$ and $n_{(1,0);(1,d)} = 0$ for $d > 1$.

The above are Gromov–Witten invariants of the blowup of $\mathbb{C}P^2$, relative to L_1, L_2 and L_3. We need the Gromov–Witten invariants relative to L_1 and L_2, which are the Gromov–Witten invariants of an exploded manifold $\mathbf{A}'$ with tropical part pictured below.

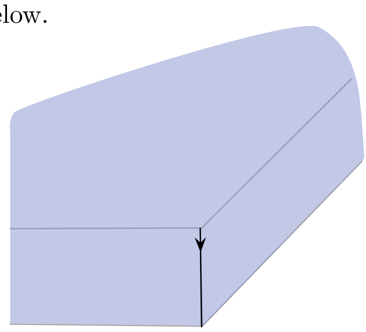

The tropical curve in the above picture is the tropical part of the exceptional sphere. For calculating the contribution of a tropical curve γ in $\underline{\mathbf{A}'}$ to Gromov–Witten invariants of $\mathbf{A}'$, vertices above the bottom boundaries of $\underline{\mathbf{A}'}$ contribute the same as the corresponding vertices in $\underline{\mathbf{A}}$. For any vertex v on the bottom of $\underline{\mathbf{A}'}$, we use our relative Gromov–Witten invariants of $(\mathbb{C}P^1)^2$ relative to 3 of its 4 toric boundary divisors. In each case, constrained contact with the middle divisor corresponds to $(0,-1)$, and $n_{(0,-1)} = 1$. Contact with the other divisors translates to tropical information differently depending on the location of v. The upshot of this is that to translate Gromov–Witten invariants of $\mathbf{A}$ to Gromov–Witten invariants of $\mathbf{A}'$, we allow ourselves to remove any outgoing edges with derivative $(0,-1)$, and otherwise restrict to the case that all infinite edges have (incoming) derivative in the span of $(1,0)$ and $(-1,-1)$. For example, the contribution of the above tropical curve is 1 (corresponding to the exceptional curve). Either of the curves pictured below may be used to calculate that 1 is the Gromov–Witten invariant of $\mathbf{A}'$ with contact data a constrained edge entering with derivative $(1,0)$ and one unconstrained edge exiting with derivative $(1,1)$. Moreover, this counts genus 0 curves that intersect the exceptional curve once.

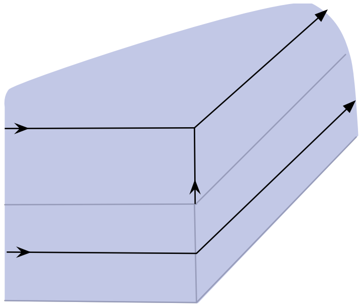

As $n_{;(0,d)} = (-1)^{d+1}/d^2$ implies the formula (6), whenever we see a part of a tropical curve γ looking like below — where the thick edge may be replaced by many edges with derivatives adding up to $(0,d)$ — the total effect of this part of γ is to multiply by $\binom{a}{d}$, and to affect the corresponding homology class by adding d to its intersection with the exceptional divisor. In the tropical part $\underline{\mathbf{B}}$ of the exploded manifold we used to represent $\mathbb{C}P^2$ blown up at n points, we may intuitively understand this as saying that a rigid edge with derivative (a,b) may interact with rigid edges coming up from the ith lower corner to leave with derivative $(a,b+d)$. After summing over all possibilities, this introduces a factor $\binom{a}{d}$ to the Gromov–Witten invariant, and corresponds to intersecting the ith exceptional divisor d times.

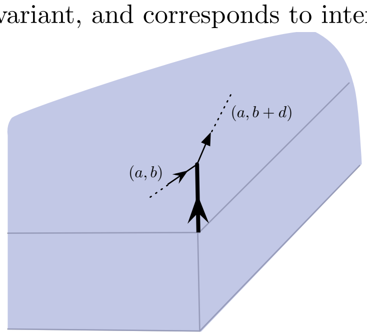

References

[1] Dan Abramovich and Qile Chen. Stable logarithmic maps to Deligne-Faltings pairs II. *The Asian Journal of Mathematics*, 18(3):465–488, 2014.

[2] Dan Abramovich and Jonathan Wise. Invariance in logarithmic Gromov-Witten theory. arxiv:1306.1222, 2013.

[3] Erwan Brugallé. Floor diagrams relative to a conic, and GW-W invariants of del Pezzo surfaces. *Adv. Math.*, 279:438–500, 2015.

[4] Jim Bryan and Naichung Conan Leung. The enumerative geometry of $K3$ surfaces and modular forms. *J. Amer. Math. Soc.*, 13(2):371–410 (electronic), 2000.

[5] Lucia Caporaso and Joe Harris. Counting plane curves of any genus. *Invent. Math.*, 131(2):345–392, 1998.

[6] L. Göttsche and R. Pandharipande. The quantum cohomology of blow-ups of $\mathbf{P}^2$ and enumerative geometry. *J. Differential Geom.*, 48(1):61–90, 1998.

[7] Mark Gross and Bernd Siebert. Logarithmic Gromov-Witten invariants. *J. Amer. Math. Soc.*, 26(2):451–510, 2013.

[8] Eleny-Nicoleta Ionel. GW invariants relative to normal crossing divisors. *Adv. Math.*, 281:40–141, 2015.

[9] Eleny-Nicoleta Ionel and Thomas H. Parker. The symplectic sum formula for Gromov-Witten invariants. *Ann. of Math. (2)*, 159(3):935–1025, 2004.

[10] Brett Parker. De Rham theory of exploded manifolds. arXiv:1003.1977, 2011.

[11] Brett Parker. Gromov-Witten invariants of exploded manifolds. arXiv:1102.0158, 2011.

[12] Brett Parker. Exploded manifolds. *Adv. Math.*, 229:3256–3319, 2012. arXiv:0910.4201.

[13] Brett Parker. Log geometry and exploded manifolds. *Abh. Math. Sem. Hamburg*, 82:43–81, 2012. arxiv:1108.3713.

[14] Brett Parker. Holomorphic curves in exploded manifolds: Kuranishi structure. arXiv:1301.4748, 2013.

[15] Brett Parker. On the value of thinking tropically to understand Ionel's GW invariants relative normal crossing divisors. arXiv:1407.3020, 2014.

[16] Brett Parker. Holomorphic curves in exploded manifolds: compactness. *Adv. Math.*, 283:377–457, 2015. arXiv:0911.2241.

[17] Brett Parker. Holomorphic curves in exploded manifolds: virtual fundamental class. arXiv:1512.05823, 2015.

[18] Brett Parker. Tropical gluing formula for Gromov-Witten invariants. arXiv:1703.05433, 2017.

[19] Mendy Shoval and Eugenii Shustin. On Gromov-Witten invariants of del Pezzo surfaces. *Internat. J. Math.*, 24(7):1350054, 44, 2013.

[20] Ravi Vakil. Counting curves on rational surfaces. *Manuscripta Math.*, 102(1):53–84, 2000.

E-mail address: `brettdparker@gmail.com`