COUNTING BITANGENTS OF A SMOOTH PLANE CURVE VIA STABLE MAPS

Joachim Kock (Recife, Brasil)
March 1999

ABSTRACT. I would like to illustrate a new approach to tangency conditions in the setting of stable maps, with a very simple example: Using intersection theory in the space $\overline{M}_{0,2}(\mathbb{P}^2, 1)$, I compute the number of bitangents of a general smooth plane curve Z.

First I give a quick introduction to stable maps, and explain the psi class of the universal curve. Then I define the cycles of pointed tangency and pointed third order contact w.r.t. Z. These cycles correspond to the conormal variety of Z, and to the set of flexes of Z, respectively. Their classes are computed via relative jets on the universal curve. Finally, in the two-pointed space I compute the cycle of bitangents, roughly by intersecting the two cycles of pointed tangency corresponding to the two marks.

1 Stable maps and their moduli spaces

Reference for this section: [2].

Basics

Definition. — An n-pointed map (in genus 0) is the data of a curve C which is a tree of projective lines (meeting in simple nodes); n distinct non-singular marked points $p_1, \ldots, p_n$, and a map $\mu : C \to \mathbb{P}^r$. The map is *Kontsevich stable* if any component that is contracted by μ has at least three special points (marks or nodes). The map is said to be of degree d when the image curve is of degree d.

1.1 Remark. — The condition that all components are rational and that the curve is a tree is equivalent to saying that the arithmetic genus of C is zero. The stability condition is equivalent to saying that the curve has only finitely many automorphisms.

Theorem 1 — *There is a normal projective variety $\overline{M}_{0,n}(\mathbb{P}^r, d)$ which is a coarse moduli space parametrising (isomorphism classes of) Kontsevich stable n-pointed degree d maps to $\mathbb{P}^r$.*

The first zero subscript indicates genus zero, the n is the number of marks.

1.2 Dimension count. — There are $n-3$ ways to choose the curve with its markings. A map $\mathbb{P}^1 \to \mathbb{P}^r$ of degree d is given by specifying $r+1$ homogeneous polynomials of degree d. There are $(r+1)(d+1)$ ways of doing that, but then we have to subtract 1 since not all the polynomials are allowed to vanish simultaneously. So we find

$$\begin{aligned}\dim \overline{M}_{0,n}(\mathbb{P}^r, d) &= (r+1)(d+1) - 1 + n - 3 \\ &= rd + r + d + n - 3.\end{aligned}$$

Gromov-Witten invariants

We want to do intersection theory on this space, so the first thing to do is to define some divisors we can intersect.

1.3 Evaluation maps. — For each mark p_i, there is a natural *evaluation map*

$$\begin{aligned}\nu_i : \overline{M}_{0,n}(\mathbb{P}^r, d) &\longrightarrow \mathbb{P}^r \\ (C; p_1, \ldots, p_n; \mu) &\longmapsto \mu(p_i).\end{aligned}$$

If $H \subset \mathbb{P}^r$ is a hyperplane, the inverse image $\nu_i^* H$ is a divisor in $\overline{M}_{0,n}(\mathbb{P}^r, d)$; let $\boldsymbol{\eta}_i$ denote its cohomology class.

1.4 Gromov-Witten invariants. — The intersection numbers

$$I_{r,d}(c_1, \ldots, c_n) := \int \boldsymbol{\eta}_1^{c_1} \cup \cdots \cup \boldsymbol{\eta}_n^{c_n} \cap [\overline{M}_{0,n}(\mathbb{P}^r, d)]$$

are called *Gromov-Witten invariants.* If the codimensions c_i add up to the dimension of the space, then $I_{r,d}(c_1, \ldots, c_n)$ is the number of rational curves in $\mathbb{P}^r$ incident to linear spaces H_i of codimension c_i. For example $I_{3,3}(3,3,3,3,3,2,2) = 5$ is the number of twisted cubics through 5 points and incident to 2 lines, all in general position in $\mathbb{P}^3$.

Theorem 2 (Kontsevich's First Reconstruction Theorem) — *All Gromov-Witten invariants can be determined recursively from the initial axiom: there is a unique line through two distinct points.*

In the case of $\mathbb{P}^2$, the recursion is particularly nice:

Theorem 3 (Kontsevich) — *Let N_d denote the number of irreducible rational plane curves passing through $3d-1$ general points. Then there is the following recursive relation between the numbers N_d:*

$$N_d + \sum_{\substack{d_A + d_B = d \\ d_A \geq 1, d_B \geq 1}} \binom{3d-4}{3d_A - 1} d_A^2 N_{d_A} \cdot N_{d_B} \cdot d_A d_B = \sum_{\substack{d_A + d_B = d \\ d_A \geq 1, d_B \geq 1}} \binom{3d-4}{3d_A - 2} d_A N_{d_A} \cdot d_B N_{d_B} \cdot d_A d_B$$

Since we know $N_1 = 1$, this formula gives all the numbers.

The proof relies on the fundamental fact that the divisors at the boundary are isomorphic to products of moduli spaces of lower dimension; this gives the recursion. Let us see:

The boundary and the recursive structure

1.5 The boundary. — Let $[n] = \{1, 2, \ldots, n\}$ be the labelling set. The boundary consists of maps having reducible source curve C.

The irreducible components of the boundary can be described as follows: For each partition of the marks $A \cup B = [n]$, and each partition of the degree $d_A + d_B = d$ (such that $\sharp A \geq 2$ if $d_A = 0$, and $\sharp B \geq 2$ if $d_B = 0$), there is an irreducible boundary divisor denoted $D(A, B; d_A, d_B)$. A general point of $D(A, B; d_A, d_B)$ represents a map μ from a two-component curve $C = C_A \cup C_B$ (with the A-points on C_A and the B-points on C_B), such that the restriction of μ to C_A is a map of degree d_A and the restriction of μ to C_B is of degree d_B. (The condition in parenthesis is just Kontsevich stability.)

1.6 Recursive structure. — We get in this way a natural isomorphism[1]

$$\overline{M}_{0,A\cup x}(\mathbb{P}^r, d_A) \times_{\mathbb{P}^r} \overline{M}_{0,B\cup x}(\mathbb{P}^r, d_B) \longrightarrow D(A, B; d_A, d_B).$$

The fibre product is taken along evaluation maps of the points x

$$\begin{array}{rcl} \nu_{x_A} : \overline{M}_{0,A\cup x}(\mathbb{P}^r, d_A) & \longrightarrow & \mathbb{P}^r \\ \nu_{x_B} : \overline{M}_{0,B\cup x}(\mathbb{P}^r, d_B) & \longrightarrow & \mathbb{P}^r, \end{array}$$

which is just the technical expression of the condition that in order to glue the maps from the two twigs, the two marked points x must map to the same point in $\mathbb{P}^r$.

1.7 Recursion. — In this way you can get along with passing intersection products to spaces of lower dimension, thus sending the problem on to recursion.

Tautological families and psi classes

1.8 Forgetful maps. — It is a fact that there are forgetful maps

$$\pi : \overline{M}_{0,n+1}(\mathbb{P}^r, d) \longrightarrow \overline{M}_{0,n}(\mathbb{P}^r, d)$$

which forgets the last mark. (Similarly you can forget other marks.) Off the boundary, this map is obviously well-defined. On the boundary it is more subtle: removing a mark from a twig of degree 0 can make that twig unstable, and in that case you have to stabilise, simply removing the unstable twig from the source curve. This is to say that the source curve of $\pi([\mu])$ can have fewer twigs than the original map $[\mu]$. But the image in $\mathbb{P}^r$ is not altered, because the removed twig was just contracting to a point anyway. This is best understood by drawing some pictures... It is not trivial to show that π is indeed a morphism...

1.9 Sections to forgetful maps. — The forgetful map $\pi : \overline{M}_{0,n+1}(\mathbb{P}^r, d) \longrightarrow \overline{M}_{0,n}(\mathbb{P}^r, d)$ comes with n canonical sections σ_i defined by repeating the i'th mark, and stabilising: This means that since we are not allowed to have two coincident marks, we introduce a

[1] It is in fact an isomorphism in most cases; only in some special cases, symmetry can account for some double point loci along closed subvarieties of D. In one case: $A = B = \emptyset$, $d_A = d_B$, the situation is so symmetric that the map becomes generically 2–1.

new twig of degree 0, attached to the source curve where the i'th mark was, and put mark i and mark $n+1$ on this new twig. Note that the new twig is stable since it has three special points: the attachment point, and the two marks. Also, since the configuration of three points on a $\mathbb{P}^1$ is uniquely determined up to automorphism, there is no choice to make with respect to the distribution of the marks on the new twig, so the map is well-defined. Note that the image of σ_i is exactly the boundary divisor characterised by having a degree 0 twig with just the two marks i and $n+1$.

1.10 Tautological families. — The collection of maps

$$
\begin{tikzcd}
\overline{M}_{0,n+1}(\mathbb{P}^r, d) \arrow[r, "\nu_{n+1}"] \arrow[d, shift left, "\pi"] & \mathbb{P}^r \\
\overline{M}_{0,n}(\mathbb{P}^r, d) \arrow[u, shift left, "\sigma_i"] \arrow[ur, "\nu_i"'] &
\end{tikzcd}
$$

is a tautological family over $\overline{M}_{0,n}(\mathbb{P}^r, d)$ in the following sense: Given a map $\mu : C \to \mathbb{P}^r$ in the space $\overline{M}_{0,n}(\mathbb{P}^r, d)$, the fibre $\pi^{-1}([\mu])$ is a curve canonically isomorphic to C (since the choices of where to put the new mark are exactly the points of C), and the sections σ_i single out n distinct marks on the fibre, corresponding exactly to the distribution of marks on C.

Under this identification, the restriction of ν_{n+1} to the fibre $\pi^{-1}([\mu])$ is exactly the map μ. By definition of the sections, the ν_{n+1}-image of the i'th mark of a fibre $\mu^1([\mu])$ is exactly $\mu(i)$; this explains the commutativity of the diagram.

1.11 The psi classes. — Let ω_π be the relative dualising sheaf of the tautological family π (see [3]). The i'th *cotangent line* of $\overline{M}_{0,n}(\mathbb{P}^r, d)$ is by definition the line bundle $\sigma_i^* \omega_\pi$. The i'th psi class is its first Chern class:

$$\boldsymbol{\psi}_i := c_1(\sigma_i^* \omega_\pi).$$

Definition. — Top products of eta classes $\boldsymbol{\eta}_i$ and psi classes $\boldsymbol{\psi}_i$ are called *gravitational descendants.*

This terminology comes from string theory. Reference for psi classes, gravitational descendants and the following theorem is [5].

Theorem 4 (Topological recursion) — *All gravitational descendants can be computed via recursions.*

2 Pointed tangency

Reference for this section: [4].

Let $Z \subset \mathbb{P}^2$ be a smooth curve, given as the zero scheme of some $f \in H^0(\mathbb{P}^2, \mathcal{O}(z))$. It is practical to perform the construction in a space with only one mark, so we place ourselves in the space $\overline{M}_{0,1}(\mathbb{P}^2, d)$. The mark of a map $\mu : C \to \mathbb{P}^r$ lands in Z if and only if $\mu^* f$ vanishes at the mark.

Definition. — A map $\mu : C \to \mathbb{P}^r$ with a single mark is *tangent* to Z at the mark if and only if $\mu^* f$ vanishes at the mark with multiplicity 2, i.e., when also the first derivative vanishes. Let $\Phi_1(Z)$ denote the cycle of maps tangent to Z at the mark.

2.1 Construction of $\Phi_1(Z)$. — Consider the tautological family

$$\begin{tikzcd} \overline{M}_{0,2}(\mathbb{P}^2, d) \arrow[r, "\nu_2"] \arrow[d, shift left, "\pi"] & \mathbb{P}^r \\ \overline{M}_{0,1}(\mathbb{P}^r, d) \arrow[u, shift left, "\sigma_1"] \arrow[ur, "\nu_1"'] & \end{tikzcd}$$

Take the pull-back of f along ν_2, and then its first jet relative to π, (i.e. we derive only in vertical direction). (See [3] for explanation of relative jets.) The vanishing of this section

$$\partial^1_\pi \nu_2^* f \in J^1_\pi \nu_2^* \mathcal{O}(z)$$

defines the locus of maps tangent to Z. We pull back along σ_1 to get only those maps tangent to Z at the mark. So we se that $\Phi_1(Z) \subset \overline{M}_{0,1}(\mathbb{P}^2, d)$ is the zero scheme of the section

$$\varphi := \sigma_1^* \partial^1_\pi (\nu_2^*(f))$$

of the bundle

$$E := \sigma_1^* J^1_\pi (\nu_2^* \mathcal{O}(z)).$$

Note that π is not a smooth map (because it has singular fibres), so a priori the sheaf $J^1_\pi(\nu_2^*\mathcal{O}(z))$ needs not be locally free. But π *is* smooth in a neighbourhood of the image of σ; indeed, sections in families of stable maps never pass through the singularities of the fibres, because marks are not allowed to fall on nodes. This shows that E is really a vector bundle.

You can show that the section φ is regular (in the sense that it is locally defined by a regular sequence, see [1]), so the class of $\Phi_1(Z)$ is given as

$$[\Phi_1(Z)] = c_2(E) \cap [\overline{M}_{0,1}(\mathbb{P}^2, d)].$$

This class is easily computed via the fundamental jet bundle exact sequence (pulled back along σ_1)

$$0 \longrightarrow \sigma_1^*\big(\nu_2^*\mathcal{O}(z) \otimes \omega_\pi\big) \longrightarrow \sigma_1^*\big(J^1_\pi \nu_2^* \mathcal{O}(z)\big) \longrightarrow \sigma_1^* \nu_2^* \mathcal{O}(z) \longrightarrow 0;$$

the last term is just $\nu_1^*\mathcal{O}(z)$, with first Chern class $z\boldsymbol{\eta}_1$, and the first term is also a line bundle, with first Chern class $z\boldsymbol{\eta}_1 + \boldsymbol{\psi}_1$.

This shows:

2.2 Proposition. — *In one-pointed space, the class of $\Phi_1(Z)$ is*

$$z\boldsymbol{\eta}_1(z\boldsymbol{\eta}_1 + \boldsymbol{\psi}_1).$$

In the presence of more marks, we must make the following reservation: let $\boldsymbol{\beta}_1$ denote the union of all boundary divisors whose general points parametrise curves with mark p_1 on a contracting twig:

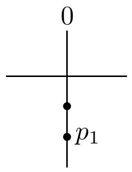

Now for maps in the intersection $\boldsymbol{\eta}_1(Z) \cdot \boldsymbol{\beta}_1$, the whole twig maps to Z. For such a map, $\mu^* f$ vanishes identically along the contracting twig and thus to any order. But we do not in general want to consider such maps tangent.

Now you can show that the simple correction for this locus works: The general formula for $\Phi_1(Z)$ is:

$$\Phi_1(Z) = z\boldsymbol{\eta}_1\big(z\boldsymbol{\eta}_1 + \boldsymbol{\psi}_1 - \boldsymbol{\beta}_1\big).$$

(Note that this formula also holds in the one-pointed space, because here $\boldsymbol{\beta}_1 = 0$. Indeed, without other marks the vertical component would be unstable, so that no such curves exist.)

2.3 Remark. — In [4], the classes $\boldsymbol{\psi}_i - \boldsymbol{\beta}_i$ are called modified psi classes and it is shown that they behave very well, and that there is a recursion allowing for the computation of top products of gravitational descendants twisted by modified psi classes.

2.4 Higher order contact? — Similarly, we would expect that the vector bundle $\sigma_1 J_\pi^{m-1} \nu_2^* \mathcal{O}(z)$ should measure m'th order contact at the mark p_1, but for $m > 2$, even in the one-pointed case, the section is not regular: its zero scheme has excess (see [1]) along the boundary cycle

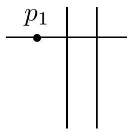

Our example: pointed lines

2.5 One-pointed lines. — The first space to study is $\overline{M}_{0,1}(\mathbb{P}^2, 1)$, the space of pointed lines. Let us denote the unique mark by the letter x. The space has dimension 3: two dimensions to determine the image line, and further one dimension to determine the marked point. Note that $\overline{M}_{0,1}(\mathbb{P}^2, 1)$ has no boundary. Indeed, the only degree partition is $1 + 0 = 1$, and there are not enough marks to put on the degree-0 twig to make it stable.

In fact the space is naturally isomorphic to the point-line incidence $I = \{(P, \ell) \in \mathbb{P}^2 \times \check{\mathbb{P}}^2 \mid P \in \ell\}$. Let $\boldsymbol{\iota}$ denote the divisor of maps passing a given point $P \in \mathbb{P}^2$; there is the following relation: $\boldsymbol{\psi}_x = \boldsymbol{\iota} - 2\boldsymbol{\eta}_x$. (See [6] for precise definition, and results on generators and relations of the divisor class group.

Let q and $\check{q}$ be the two projections from the incidence correspondence; let $L \subset \mathbb{P}^2$ be a line, and let $h = q^*h$ be the class of a line in $\mathbb{P}^2$. Then the divisor $\boldsymbol{\eta}_x(L)$ of maps whose mark goes to L, corresponds to q^*L, and the class $\boldsymbol{\eta}_x$ corresponds to h. Let $P \in \mathbb{P}^2$ be a point, let $P^* \subset \check{\mathbb{P}}^2$ be the corresponding pencil, and let $\check{h} = \check{q}^*\check{h}$ be its class. Then the divisor $\boldsymbol{\iota}(P)$ of maps passing through P, corresponds to $\check{q}^*P^*$, so the class $\boldsymbol{\iota}$ corresponds to $\check{h}$.

The relation $\boldsymbol{\eta}_x\boldsymbol{\iota} = \boldsymbol{\eta}_x^2 + \boldsymbol{\iota}^2$ holds in $A^*\overline{M}_{0,1}(\mathbb{P}^2, 1)$, and corresponds of course to the well-known relation $h\check{h} = h^2 + \check{h}^2$ for the incidence correspondence.

The cycle $\Phi_x(Z)$ is the set of maps tangent to Z at the mark x; that is, the set of pointed lines tangent to Z at the mark. So in terms of the incidence correspondence, it is exactly the *conormal variety* (graph of the Gauss map).

Let us compute the class, using $\boldsymbol{\psi}_x = \boldsymbol{\iota} - 2\boldsymbol{\eta}_x$:

$$\begin{aligned} \Phi_x(Z) &= z\boldsymbol{\eta}_x(z\boldsymbol{\eta}_x + \boldsymbol{\psi}_x) \\ &= z\boldsymbol{\eta}_x\boldsymbol{\iota} + z(z-2)\boldsymbol{\eta}_x^2. \end{aligned}$$

Using also the relation $\boldsymbol{\eta}_x\boldsymbol{\iota} = \boldsymbol{\eta}_x^2 + \boldsymbol{\iota}^2$, we get

$$\Phi_x(Z) = z\boldsymbol{\iota}^2 + z(z-1)\boldsymbol{\eta}_x^2$$

which is the well known formula $z\check{h}^2 + z(z-1)h^2$ for the class of the conormal variety of Z.

2.6 Flexes of Z. — Since $\overline{M}_{0,1}(\mathbb{P}^2, 1)$ has no boundary, the jet bundle argument actually works also for third order contact in this case: That is, the section of the bundle $\sigma^* J^2_\pi \nu_2^* \mathcal{O}(z)$ is actually regular and its zero scheme is exactly the cycle $\Phi_x^{(3)}$ of pointed lines having third order contact with Z at the mark x. This is a codimension-3 cycle, so it is a finite number of points.

The class is then given as

$$c_3(\sigma^* J^2_\pi \nu_2^* \mathcal{O}(z)) = z\boldsymbol{\eta}_x\big(z\boldsymbol{\eta}_x + \boldsymbol{\psi}_x\big)\big(z\boldsymbol{\eta}_x + 2\boldsymbol{\psi}_x\big).$$

Knowing that $\boldsymbol{\psi}_x = \boldsymbol{\iota} - 2\boldsymbol{\eta}_x$, and that $\boldsymbol{\iota}^2\boldsymbol{\eta}_x = \boldsymbol{\iota}\boldsymbol{\eta}_x^2 = 1$, the class expands to

$$\int \Phi_x^{(3)} = 3z(z-2),$$

which is just the number of flexes of Z. (Alternatively, trust the top intersection formulae $\boldsymbol{\psi}_x^2\boldsymbol{\eta}_x = -3$, $\boldsymbol{\psi}_x\boldsymbol{\eta}_x^2 = 1$, and $\boldsymbol{\eta}_x^3 = 0$.)

3 Counting bitangents

3.1 The two-pointed space — Now we consider the space $\overline{M}_{0,2}(\mathbb{P}^2, 1)$, with markings p_1 and p_2. It has dimension 4, and there is a unique boundary divisor $\boldsymbol{\beta}$ corresponding to maps of type

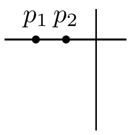

which again correspond to the set of image lines whose two marks are coincident.

The intersection ring is generated by the three divisor classes $\boldsymbol{\eta}_1$, $\boldsymbol{\eta}_2$ and $\boldsymbol{\beta}$.

3.2 The forgetful maps from $\overline{M}_{0,2}(\mathbb{P}^2, 1)$. — There are two canonical maps

$$\begin{tikzcd} \overline{M}_{0,2}(\mathbb{P}^2, 1) \arrow[d, shift right, "\pi_1"'] \arrow[d, shift left, "\pi_2"] \\ \overline{M}_{0,1}(\mathbb{P}^2, 1) \end{tikzcd}$$

that forget one of the two marks (and stabilise). The map π_1 forgets mark 2, whereas π_2 forgets mark 1. Note that $\boldsymbol{\eta}_1 = \pi_1^* \boldsymbol{\eta}_x$ and $\boldsymbol{\eta}_2 = \pi_2^* \boldsymbol{\eta}_x$. Similarly, $\Phi_1 = \pi_1^* \Phi_x$ and $\Phi_2 = \pi_2^* \Phi_x$.

3.3 Their section. — The maps π_1 and π_2 have a common section

$$\sigma : \overline{M}_{0,1}(\mathbb{P}^2, 1) \longrightarrow \overline{M}_{0,2}(\mathbb{P}^2, 1),$$

defined by repeating the mark, and stabilising. Stabilising means that since the new mark cannot coincide with the old, there must be a new contracting component to separate them. Here is a picture:

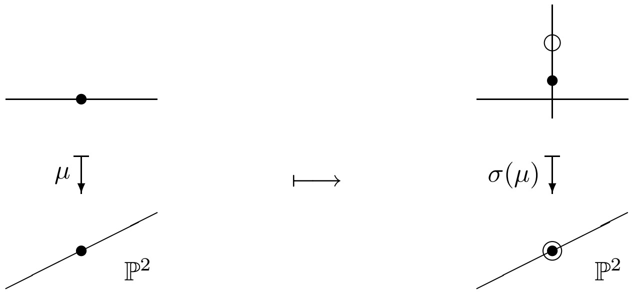

It follows that the image of σ is exactly $\boldsymbol{\beta}$, the unique boundary divisor of $\overline{M}_{0,2}(\mathbb{P}^2, 1)$.

Everything that is difficult or interesting takes place at the boundary. So it is worth having a look at the σ_*-image of all the cycles we studied down in $\overline{M}_{0,1}(\mathbb{P}^2, d)$.

The most important cycle in $\overline{M}_{0,1}(\mathbb{P}^2, d)$ is $\boldsymbol{\eta}_x$, and, since we are looking with respect to Z we are more interested in the cycle $z\boldsymbol{\eta}_x$ of maps whose mark goes to Z. Its image $\sigma_*(z\boldsymbol{\eta}_x)$, can also be described as $z\boldsymbol{\eta}_1\boldsymbol{\beta}$ or $z\boldsymbol{\eta}_2\boldsymbol{\beta}$. We will represent it graphically as

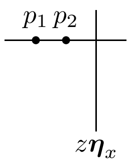

Next, we will need the cycle $\sigma_*(\Phi_x(Z))$, which will be depicted

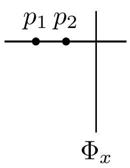

And finally the image of the flex cycle $\Phi_x^{(3)}(Z)$ of lines that are flexes to Z at the mark; we recall that this is a finite set of points in $\overline{M}_{0,1}(\mathbb{P}^2, 1)$. Its image $\sigma_* \Phi_x^{(3)}(Z)$ will be drawn as

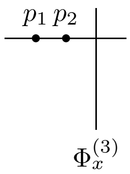

3.4 Remark. — The space is canonically isomorphic to the product $\mathbb{P}^2 \times \mathbb{P}^2$ blown up along the diagonal. Indeed, consider the product of the evaluation maps

$$\underline{\nu} : \overline{M}_{0,2}(\mathbb{P}^2, 1) \longrightarrow \mathbb{P}^2 \times \mathbb{P}^2$$

Now you can convince yourself that this map is just the blow-up of $\mathbb{P}^2 \times \mathbb{P}^2$ along the diagonal. Think like this: the information in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$ consists of a line with two marks on it, possibly coincident. If the marks are distinct then they already determine the line. So away from the boundary the map is an isomorphism. Two coincident points (i.e. a point on the diagonal) does not determine a line, but when blown up, each such point wins a $\mathbb{P}^1$ parametrising the possible tangent directions. And then it *does* determine a line.

Bitangents

3.5 More general definition. — Let $\Lambda(m_1 p_1 + m_2 p_2)$ denote the closure in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$ of the locus of maps $\mu : \mathbb{P}^1 \to \mathbb{P}^2$ such that $\mu^* Z \geq m_1 p_1 + m_2 p_2$. For example, $\Lambda(3p_1) = \Phi_1^{(3)}$ is the locus of maps having third order contact with Z at the mark p_1.

What we want to compute here, is the class of the locus $\Lambda(2p_1 + 2p_2)$, which is the locus of maps bitangent to Z (at the marks p_1 and p_2).

A first approximation to that locus is the intersection $\Phi_1^{(2)} \Phi_2^{(2)}$. However this intersection has excess (see [1]) along the locus of maps $\boldsymbol{\beta} \cdot \Phi_1 = \boldsymbol{\beta} \cdot \Phi_2$ (which is of codimension 3 instead of 4). This excess is very nasty because it is not a scheme-theoretic connected component of the intersection: it has embedded double points at the points that correspond to flexes. Therefore it is very difficult to apply residual intersection formulae.

I will sketch another method to separate the excess: The technique is to take the four divisors one by one and intersect, and in each step get a Λ-cycle, and some further components that we throw away.

The recursive argument

We place ourselves inside Φ_1. So all our maps are already guaranteed tangent to Z at p_1. Now intersect with the divisor $z\boldsymbol{\eta}_2$, i.e. force the second mark to land in Z. This can happen in two ways: *Either* the two marks on the source do *not* come together; then we are in $\Lambda(2p_1 + p_2)$. *Or*, the condition is satisfied but simply because the two marks are on a contracting twig; in this case we have run into $\boldsymbol{\beta} \cdot \Phi_1 = \sigma_* \Phi_x$.

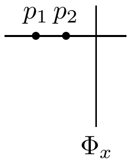

This gives a complete set-theoretic description of the intersection:

$$z\boldsymbol{\eta}_2 \cdot \Phi_1 = \Lambda(2p_1 + p_2) \cup \sigma_* \Phi_x.$$

We claim that the multiplicities are 1 and 2, respectively. So this gives the class of $\Lambda(2p_1 + p_2)$:

$$\Lambda(2p_1 + p_2) = z\boldsymbol{\eta}_2 \cdot \Phi_1 - 2\sigma_* \Phi_x.$$

Now we impose the last condition, namely $z\boldsymbol{\eta}_2 + \boldsymbol{\psi}_2 - \boldsymbol{\beta}$, which means second order vanishing of the jet section. (Note that this is a well-defined condition, since we are already inside $z\boldsymbol{\eta}_2$.) Again there are two components in the intersection: one is $\Lambda(2p_1 + 2p_2)$ and the other is some cycle at the boundary which we proceed to identify: We are already inside $\Lambda(2p_1 + p_1)$, so if we run into the boundary this means that p_2 approaches p_1 (and without its image leaving Z). When p_2 finally coincides with p_1, the line must have become an inflectional tangent. So we conclude that the second component in the intersection is the cycle $\sigma_* \Phi_x^{(3)}$:

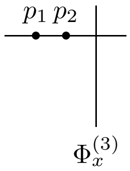

That is,

$$(z\boldsymbol{\eta}_2 + \boldsymbol{\psi}_2 - \boldsymbol{\beta}) \cdot \Lambda(2p_1 + p_2) = \Lambda(2p_1 + 2p_2) \cup \sigma_* \Phi_x^{(3)}.$$

Again we claim that the multiplicity of $\Lambda(2p_1 + 2p_2)$ is 1, and the multiplicity of $\sigma_* \Phi_x^{(3)}$ is 2.

This gives the class of $\Lambda(2p_1 + 2p_2)$:

$$\begin{aligned} \int \Lambda(2p_1 + 2p_2) &= -2\sigma_* \Phi_x^{(3)} + (z\boldsymbol{\eta}_2 + \boldsymbol{\psi}_2 - \boldsymbol{\beta}) \cdot \Lambda(2p_1 + p_2) \\ &= -2\sigma_* \Phi_x^{(3)} + \Phi_2 \cdot \Phi_1 - (z\boldsymbol{\eta}_2 + \boldsymbol{\psi}_2 - \boldsymbol{\beta}) \cdot 2\sigma_* \Phi_x \end{aligned}$$

All these intersection products are known: you can substitute the Φ-classes by their class $z\boldsymbol{\eta}_i(z\boldsymbol{\eta}_i + \boldsymbol{\psi} - \boldsymbol{\beta})$ as we have computed above; and as to the σ-push-ups, you can compute them for example by going back down in the one-pointed space.

Computing like this you find that the integral of $\Lambda(2p_1 + 2p_2)$ is

$$z^4 - 2z^3 - 9z^2 + 18z.$$

This is twice the number of bitangents (as you can compute for example via the Plcker formulae, or via de Jonquires like arguments, cf. [3]). Of course, we have counted with named marks, and really we are not interested in which point has which name...

Appendix: Bitangents via the Plücker formula

— Just to recall how the count is simple to do in the classical way: Let $Z \subset \mathbb{P}^2$ be a curve of degree z with exactly δ nodes and κ cusps. Then the Plücker formula gives the degree z' of the dual curve Z':

$$z' = z(z-1) + 2\delta + 3\kappa.$$

Let respectively δ' and κ' denote the number of nodes and cusps of the dual curve. Then δ' is also the number of bitangents of Z, and κ' is the number of flexes of Z. If Z is general, it has no other special tangents in addition to bitangents and flexes, so the dual curve has only nodes and cusps as singularities. Since the dual curve of the dual curve is the original curve, the Plücker formula for the dual curve reads:

$$z = z'(z'-1) + 2\delta' + 3\kappa'.$$

Now suppose we already know the formula for the number of flexes,

$$\kappa' = 3z(z-2),$$

then we can plug into the formula to find δ', the number of bitangents:

$$\begin{aligned} \delta' &= \frac{z - z(z-1)\big(z(z-1)-1\big) - 3\cdot 3z(z-2)}{2} \\ &= \tfrac{1}{2}z^4 - z^3 - \tfrac{9}{2}z^2 + 9z. \end{aligned}$$

References

[1] Fulton, W. *Intersection Theory.* Springer-Verlag, New York, 1985.

[2] Fulton, W., and Pandharipande, R. Notes on stable maps and quantum cohomology. In *Algebraic Geometry, Santa Cruz 1995*, J. Kollár, R. Lazarsfeld and D. Morrison, Eds., vol. 62, II of *Proc. Symp. Pure. Math.*, pp. 45–96. (alg-geom/9608011).

[3] Gatto, L. The Recife notes on intersection theory over moduli spaces of curves. Recife Lecture Notes 1998, Preliminary version, 190pp., available at `http://www.dmat.ufpe.br/~jojo/Mat/RLN/letterio.ps`.

[4] Kock, J. Recursion for twisted descendants and characteristic numbers of rational curves. Preprint, math.AG/9902021.

[5] Pandharipande, R. Rational curves on hypersurfaces (after A. Givental). Preprint, math.AG/9806133.

[6] Pandharipande, R. Intersections of $\mathbb{Q}$-divisors on Kontsevich's moduli space $\overline{M}_{0,n}(\mathbb{P}^r, d)$ and enumerative geometry. *Trans. Amer. Math. Soc.* **351** (1999), 1481–1505. (alg-geom/9504004).