

Counting Bitangents with Stable Maps

DAVID AYALA AND RENZO CAVALIERI

*Department of Mathematics
Stanford University
450 Serra Mall, Bldg. 380
Stanford, CA 94305-2125*

and

*Department of Mathematics
University of Utah
155 South 1400 East, Room 233
Salt Lake City, UT 84112-0090*

Email: ayala@math.Stanford.EDU and renzo@math.utah.edu

Abstract

This paper is an elementary introduction to the theory of moduli spaces of curves and maps. As an application to enumerative geometry, we show how to count the number of bitangent lines to a projective plane curve of degree d by doing intersection theory on moduli spaces.

AMS Classification 14N35; 14H10, 14C17

Keywords Moduli Spaces, Rational Stable Curves, Rational Stable Maps, Bitangent Lines.

Introduction

Philosophy and Motivation

The most apparent goal of this paper is to answer the following enumerative question:

“What is the number $N_{\mathcal{B}}(d)$ of bitangent lines to a generic projective plane curve Z of degree d ?”

This is a very classical question, that has been successfully solved with fairly elementary methods (see for example [GH94], page 277). Here we propose to approach it from a very modern and “technological” angle: we think of lines in the projective plane as maps $\mu : \mathbb{P}^1 \rightarrow \mathbb{P}^2$ of degree 1. We mark two points p_1 and p_2 on \mathbb{P}^1 and keep track of their image via the map μ . We then construct the space of all such marked maps, and ask ourselves: can we understand the locus \mathcal{B} of all maps that are tangent to Z at the images of both p_1 and p_2 ? The answer fortunately is yes. The description of \mathcal{B} allows us to produce a closed formula for $N_{\mathcal{B}}$ in all degrees.

This brief description already reveals that there is something deep and interesting going on here, and that the journey is much more important than the destination itself. Our major goal is to introduce the reader to the rich and beautiful theory of Moduli Spaces in a hopefully “soft” way, with the final treat of seeing it concretely applied to solve our classical problem.

It is our intention for this paper to be a very readable expository work. We designed it to be accessible to a first year graduate student who is considering algebraic geometry as a specialty field. We emphasize geometric intuition and visualization above all, at the cost of silently glossing over some technical details here and there.

Outline of the Paper

This paper is divided into three sections, getting progressively more advanced.

The first section introduces some basic ideas and techniques in modern algebraic geometry, necessary to develop and understand the later two sections and is intended for the unexperienced reader. We quickly tread through the most basic ideas in intersection theory; we introduce the concept of families of algebro-geometric objects; we discuss the specific example of vector bundles, and give a working sketch of the theory of Chern classes. Finally, we describe two interesting constructions: the blow-up and jet bundles.

Entire books have been written on each of these topics, so we have no hope or pretense to be complete, or even accurate. Yet, we still think it valuable to present what lies in the back of a working mathematician's mind, in the firm belief that a solid geometric intuition is the best stairway to understand and motivate the technicalities and abstract generalizations needed to make algebraic geometry "honest".

The second section is the development of most of the theory. After a quick qualitative introduction to moduli spaces, we discuss our main characters: the moduli spaces of rational stable curves, and of rational stable maps. Intersection theory on the moduli spaces of stable maps, commonly referred to as Gromov-Witten Theory, is currently an extremely active area of research.

Finally, in the third section we apply all the theory developed so far to solve the bitangent problem. We explore in further detail the moduli spaces of rational stable maps of degree 1 to \mathbb{P}^2 , with one and two marked points. By intersecting appropriate cycle classes on these spaces we extract one of the classical *Plücker Formulas*, expressing the number of bitangents as a function of the degree d of the curve.

References

We suggest here some canonical references for the reader in search of more rigor and completeness. For intersection theory, [Ful98] is a fairly technical book, but definitely it has the last word on it. It also presents Chern classes from an algebraic point of view. A discussion of Chern classes from a geometric point of view can be found in [BT82].

A good treatment of blow-ups can be found in any basic book in algebraic geometry, for example [GH94] or [Har77].

A very pleasant reference for jet bundles is [Vak98]. An extensive treatment of jet bundles is found in [Sau89].

Our presentation of moduli spaces follows the spirit of [KV99]; for anybody interested in getting serious, [HM98] is the way to go. Finally, a good introduction to ψ classes is [Koc01].

Acknowledgments

We first of all owe the inspiration for this work to Joachim Kock, who outlines this strategy for counting bitangents in his talk [Koc99], and is also responsible, with Israel Vainsencher, for the best elementary introduction to Gromov-Witten Theory we know of, [KV99]. We also thank Aaron Bertram, Herb Clemens, Tommaso de Fernex, Hugo Rossi and Ravi Vakil for their comments, suggestions and encouragement.

1 Preliminaries

1.1 Intersection Theory

It will be helpful, but not essential, that the reader be familiar with the Chow ring, $A^*(X)$, of an algebraic variety X . The ring¹ $A^*(X)$ is, in some loose sense, the algebraic counterpart of the cohomology ring $H^*(X)$, and it allows us to make precise in the algebraic category the intuitive concepts of oriented intersection in topology.

We think of elements of the group $A^n(X)$ as formal finite sums of codimension n closed subvarieties (cycles), modulo an equivalence relation called rational equivalence. $A^*(X) = \bigoplus_0^{\dim X} A^n(X)$ is a graded ring with product given by intersection.

Intersection is independent of the choice of representatives for the equivalence classes.

In topology, if we are interested in the cup product of two cohomology classes \mathbf{a} and \mathbf{b} , we can choose representatives a and b that are transverse to each other. We can assume this since transversality is a generic condition: if a and b are not transverse then we can perturb them ever so slightly and make them transverse while not changing their classes. This being the case, then $a \cap b$ represents the cup product class $\mathbf{a} \cup \mathbf{b}$.

¹This is probably our greatest sloppiness. In order for $A^*(X)$ to be a ring we need X to be smooth. Since the spaces we will actually work with satisfy these hypotheses, we do not feel too guilty.

In algebraic geometry, even though this idea must remain the backbone of our intuition, things are a bit trickier. We will soon see examples of cycles that are rigid, in the sense that their representative is unique, and hence “unwigglable”. Transversality then becomes an unattainable dream. Still, with the help of substantially sophisticated machinery (the interested reader can consult [Ful98]), we can define an algebraic version of intersection classes and a product that reduces to the “geometric” one when transversality can be achieved.

Throughout this paper, a bolded symbol will represent a class, the unbolded symbol a geometric representative. The intersection of two classes \mathbf{a} and \mathbf{b} will be denoted by \mathbf{ab} .

Example: the Chow Ring of Projective Space.

$$A^*(\mathbb{P}^n) = \frac{\mathbb{C}[\mathbf{H}]}{(\mathbf{H}^{n+1})},$$

where $\mathbf{H} \in A^1(\mathbb{P}^n)$ is the class of a hyperplane H .

1.2 Families and Bundles

One of the major leaps in modern algebraic geometry comes from the insight that, to fully understand algebraic varieties, we should not study them one by one, but understand how they organize themselves in families.

We are all familiar, maybe subconsciously, with the concept of a family. When, in high school, we dealt with “all parabolas of the form $y = ax^2$ ” or “all circles with center at the origin”, we had in hand prime examples of families of algebraic varieties.

The idea is quite simple: we have a parameter space, B , called the base of the family. For each point $b \in B$ we want an algebraic variety X_b with certain properties. Further, we want all such varieties to be organized together to form an algebraic variety \mathcal{E} , called the total space of the family.

A little more formally we could define a family of objects of type \mathcal{P} as a

morphism of algebraic varieties

$$\begin{array}{c} \mathcal{E} \\ \pi \downarrow \\ B \end{array}$$

where $\pi^{-1}(b)$ is an object of type \mathcal{P} .

A **section** of a family $\pi : \mathcal{E} \rightarrow B$ is a map $s : B \rightarrow \mathcal{E}$ such that $\pi \circ s : B \rightarrow B$ is the identity map. Often, the section s is written

$$\begin{array}{c} \mathcal{E} \\ \pi \downarrow \uparrow s \\ B. \end{array}$$

Notice that $s(b) \in \pi^{-1}(b)$.

Given a family $\pi : \mathcal{E} \rightarrow B$ and a map $f : M \rightarrow B$ we can construct a new family

$$\begin{array}{c} f^*\mathcal{E} \\ \downarrow \pi_f \\ M, \end{array}$$

called the **pull-back of π via f** :

$$f^*\mathcal{E} = \{(m, e) \in M \times \mathcal{E} \mid f(m) = \pi(e)\}.$$

Intuitively, the fibre of π_f over a point $m \in M$ will be the fibre of π over $f(m)$. An essential property of this construction is that it is natural, up to isomorphism.

1.2.1 Vector Bundles

A **vector bundle** of rank n is a family $\pi : \mathcal{E} \rightarrow B$ of vector spaces over \mathbb{C} of dimension n which is **locally trivial**². By locally trivial we mean that there is an open cover $\{U_\alpha\}$ of B such that $\pi^{-1}(U_\alpha) \cong U_\alpha \times \mathbb{C}^n$. Our vector bundle is uniquely determined by how these trivial pieces glue together.

²To be precise, more structure is needed: the clutching functions must take values in $GL(n, \mathbb{C})$.

A vector bundle of rank one is called a **line bundle**, as its fibers are (complex) lines.

Given two vector bundles

$$\begin{array}{ccc} \mathcal{E}_1 & & \mathcal{E}_2 \\ \pi_1 \downarrow & \text{and} & \pi_2 \downarrow \\ B & & B \end{array}$$

over the same base space, one can define their Whitney sum

$$\begin{array}{c} \mathcal{E}_1 \oplus \mathcal{E}_2 \\ \pi \downarrow \\ B, \end{array}$$

where a fibre $\pi^{-1}(b)$ is the direct sum of the vector spaces $\pi_1^{-1}(b) \oplus \pi_2^{-1}(b)$. It can be easily verified that this family satisfies the local triviality condition.

Similarly, one can define the tensor product $\mathcal{E}_1 \otimes \mathcal{E}_2$, the dual bundle \mathcal{E}^* , the wedge product $\bigwedge^p(\mathcal{E})$ and the bundle $Hom(\mathcal{E}_1, \mathcal{E}_2) = (\mathcal{E}_2 \otimes \mathcal{E}_1^*)$.

1.2.2 Characteristic Classes of Bundles

For every vector bundle there is a natural section $s_0 : B \rightarrow \mathcal{E}$ defined by

$$s_0(b) = (b, 0) \in \{b\} \times \mathbb{C}^n.$$

It is called the zero section, and it gives an embedding of B into \mathcal{E} .

A natural question to ask is if there exists another section $s : B \rightarrow \mathcal{E}$ which is disjoint from the zero section, i.e. $s(b) \neq s_0(b)$ for all $b \in B$. The **Euler class** of this vector bundle ($\mathbf{e}(\mathcal{E}) \in A^n(B)$) is defined to be the class of the self-intersection of the zero section: it measures obstructions for the above question to be answered affirmatively. This means that $\mathbf{e}(\mathcal{E}) = 0$ if and only if a never vanishing section exists. It easily follows from the Poincaré-Hopf theorem that for a manifold M , the following formula holds:

$$\mathbf{e}(TM) \cap [M] = \chi(M).$$

That is, the degree of the Euler class of the tangent bundle is the Euler characteristic.

The Euler class of a vector bundle is the first and most important example of a whole family of “special” cohomology classes associated to a bundle, called the **Chern classes** of \mathcal{E} . The k -th Chern class of \mathcal{E} , denoted $\mathbf{c}_k(\mathcal{E})$, lives in $A^k(B)$. In the literature you can find a wealth of definitions for Chern classes, some more geometric, dealing with obstructions to finding a certain number of linearly independent sections of the bundle, some purely algebraic. Such formal definitions, as important as they are (because they assure us that we are talking about something that actually exists!), are not particularly illuminating. In concrete terms, what you really need to know is that Chern classes are cohomology classes associated to a vector bundle that satisfy a series of really nice properties, which we are about to recall.

Let \mathcal{E} be a vector bundle of rank n :

identity: by definition, $\mathbf{c}_0(\mathcal{E}) = 1$.

normalization: the n -th Chern class of \mathcal{E} is the Euler class:

$$\mathbf{c}_n(\mathcal{E}) = \mathbf{e}(\mathcal{E}).$$

vanishing: for all $k > n$, $\mathbf{c}_k(\mathcal{E}) = 0$.

pull-back: Chern classes commute with pull-backs:

$$f^* \mathbf{c}_k(\mathcal{E}) = \mathbf{c}_k(f^* \mathcal{E}).$$

tensor products: if L_1 and L_2 are line bundles,

$$c_1(L_1 \otimes L_2) = c_1(L_1) + c_1(L_2).$$

Whitney formula: for every extension of bundles

$$0 \rightarrow \mathcal{E}' \rightarrow \mathcal{E} \rightarrow \mathcal{E}'' \rightarrow 0,$$

the k -th Chern class of \mathcal{E} can be computed in terms of the Chern classes of \mathcal{E}' and \mathcal{E}'' , by the following formula:

$$\mathbf{c}_k(\mathcal{E}) = \sum_{i+j=k} \mathbf{c}_i(\mathcal{E}') \mathbf{c}_j(\mathcal{E}'').$$

Using the above properties it is immediate to see:

- (1) all the Chern classes of a trivial bundle vanish (except the 0-th, of course);
- (2) for a line bundle L , $\mathbf{c}_1(L^*) = -\mathbf{c}_1(L)$.

To show how to use these properties to work with Chern classes, we will now calculate the first Chern class of the tautological line bundle over \mathbb{P}^1 . The tautological line bundle is

$$\begin{array}{c} \mathcal{S} \\ \pi \downarrow \\ \mathbb{P}^1, \end{array}$$

where $\mathcal{S} = \{(p, l) \in \mathbb{C}^2 \times \mathbb{P}^1 \mid p \in l\}$. It is called tautological because the fiber over a point in \mathbb{P}^1 is the line that point represents.

Our tautological family fits into the short exact sequence of vector bundles over \mathbb{P}^1

$$\begin{array}{ccccccc} 0 & \rightarrow & \mathcal{S} & \rightarrow & \mathbb{C}^2 \times \mathbb{P}^1 & \rightarrow & \mathcal{Q} \rightarrow 0 \\ & & & \searrow & \downarrow & \swarrow & \\ & & & & \mathbb{P}^1 & & \end{array}$$

where \mathcal{Q} is the bundle whose fibre over a line $l \in \mathbb{P}^1$ is the quotient vector space \mathbb{C}^2/l . Notice that \mathcal{Q} is also a line bundle. From the above sequence, we have that

$$0 = \mathbf{c}_1(\mathbb{C}^2 \times \mathbb{P}^1) = \mathbf{c}_1(\mathcal{S}) + \mathbf{c}_1(\mathcal{Q}). \quad (1)$$

Since \mathbb{P}^1 is topologically a sphere, which has Euler characteristic 2, then

$$2 = \mathbf{c}_1(T\mathbb{P}^1) = \mathbf{c}_1(\mathcal{S}^*) + \mathbf{c}_1(\mathcal{Q}) = -\mathbf{c}_1(\mathcal{S}) + \mathbf{c}_1(\mathcal{Q}). \quad (2)$$

The second equality in 2 holds because $T\mathbb{P}^1$ is the line bundle $\text{Hom}(\mathcal{S}, \mathcal{Q}) = \mathcal{S}^* \otimes \mathcal{Q}$.

It now follows from (1) and (2) that $\mathbf{c}_1(\mathcal{S}) = -1$.

1.3 Blow-up

Let us begin by discussing the prototypical example of the blow up of a point on a surface: first off, the blow up is a local construction and so we need only understand the local picture.

Consider the map

$$\begin{aligned} \phi : \quad \mathbb{C}^2 &\rightarrow \mathbb{C} \\ (x, y) &\mapsto y/x. \end{aligned}$$

This is a rational map and is not defined on the line $\{x = 0\}$. We may try to fix this by modifying our target space to \mathbb{P}^1 . Still, ϕ cannot be defined at $\mathbf{0} := (0, 0)$. In fact, along any line l through the origin, the limit of ϕ at $\mathbf{0}$ is the slope of l . We would like to modify \mathbb{C}^2 to a smooth surface birational to it, where the map ϕ can be defined everywhere. We would like points outside $\mathbf{0}$ to remain “untouched” and $\mathbf{0}$ to be replaced by a \mathbb{P}^1 whose points represent all tangent directions at $\mathbf{0}$.

Here is how to do it: consider the graph of ϕ , $\Gamma_\phi \subset \mathbb{C}^2 \times \mathbb{P}^1$. We have the commutative diagram

$$\begin{array}{ccc} & \Gamma_\phi \subset \mathbb{C}^2 \times \mathbb{P}^1 & \\ (id, \phi) \nearrow & \swarrow \pi_1 & \downarrow \pi_2 \\ \mathbb{C}^2 \setminus \{0\} & \xrightarrow{\phi} & \mathbb{P}^1 \end{array}$$

The closure $\overline{\Gamma}_\phi$ is what we are looking for. It is birational to \mathbb{C}^2 ; the left projection $\pi_1|_{\overline{\Gamma}_\phi}$ is an isomorphism onto $\mathbb{C}^2 - \{0\}$ and $\pi_1^{-1}(0) = \mathbb{P}^1$. We define the blow-up of \mathbb{C}^2 at $\mathbf{0}$ as $Bl_0(\mathbb{C}^2) := \overline{\Gamma}_\phi = \Gamma_\phi \cup \mathbb{P}^1$; the projective line $\pi_1^{-1}(0)$ is called the exceptional divisor of the blow-up and denoted E . We have obtained a new (smooth!) space by replacing, in a particularly favorable way, a point with the projectivization of its normal bundle.

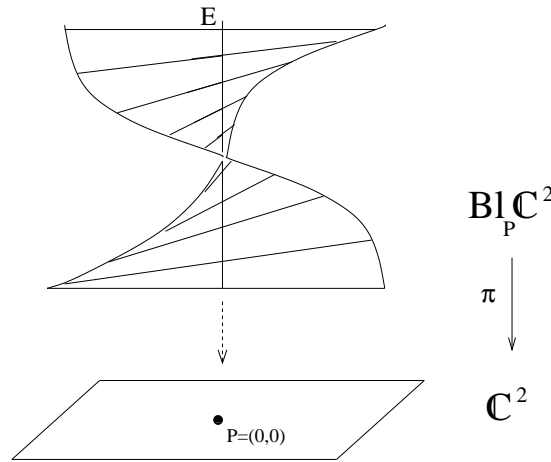


Figure 1: the blow-up of \mathbb{C}^2 at the origin.

In general, for $Y \subset X$ a closed subvariety of codimension $k \geq 0$, one can construct a new space $Bl_Y(X)$ such that:

- (1) $Bl_Y(X)$ is birational to X ;
- (2) points outside Y are untouched;
- (3) a point in Y is replaced by \mathbb{P}^{k-1} , representing the “normal” directions to Y at that point.

The total space of the blow-up of \mathbb{C}^2 at $\mathbf{0}$ admits a natural map to the exceptional divisor, consisting of projecting points along lines through the origin. This realizes $Bl_0(\mathbb{C}^2) \rightarrow E$ as a line bundle over \mathbb{P}^1 . This is the tautological bundle, which does not have any global sections besides the 0-section. It follows that the class \mathbf{E} of the exceptional divisor admits only one representative, namely E itself. It is therefore impossible to compute the self-intersection $\mathbf{E}\mathbf{E}$ by means of intersecting two transverse representatives of the class.

1.4 Jet Bundles

Let L be a line bundle over a variety, X . Then the local sections of this line bundle form a vector space. In fact, locally, such a section is just a complex valued function on some open set in X . We will now describe a new vector bundle over X whose fibre over $x \in X$ consists of all Taylor expansions of these sections centered about x and truncated after degree k . To see how the locally trivial charts of this bundle glue together is simply a matter of shifting the center of a Taylor expansion. We call this bundle the k th jet bundle of L and denote it by $J^k L$. In particular $J^0 L = L$.

Notice that the first jet bundle keeps track of all locally defined functions and differential forms and so there is an obvious surjection $J^1 L \rightarrow L$. This gives us the short exact sequence

$$0 \rightarrow L \otimes \Omega \rightarrow J^1 L \rightarrow L \rightarrow 0$$

which will be an essential tool later on.

The previous statement is a particular case of what can be considered the “fundamental theorem of jet bundles”.

Theorem 1 *For all $n \geq 0$, the sequence*

$$0 \rightarrow L \otimes Sym^n \Omega \rightarrow J^n L \rightarrow J^{n-1} L \rightarrow 0 \tag{3}$$

is exact.

For a slightly more rigorous and still enjoyable account of jet bundles, refer to [Vak98].

2 Moduli Spaces

2.1 A “High School” Example

What is the idea of a moduli space? A moduli space of geometric objects of a certain type is a space which “encodes” information about collections of geometric objects of a given type, in the sense that:

- (1) points in the moduli space correspond bijectively to the desired geometric objects;
- (2) the moduli space itself has an algebraic structure that respects how the objects can organize themselves in families.

For example, suppose that we would like to consider the space of all circles in the plane. Since a circle is uniquely the zero locus of a second degree polynomial of the form $(x - x_0)^2 + (y - y_0)^2 - r^2$, upon specifying the coordinates of the center and its radius, we have completely identified the circle. Thus, the space of all circles in the plane can be represented by $\mathcal{M} := \mathbb{R}^2 \times \mathbb{R}_+$. This is indeed much more than just a set-theoretic correspondence.

Consider the tautological family

$$\begin{array}{c} \mathcal{U} \\ \downarrow \pi \\ \mathcal{M} \end{array}$$

where $\mathcal{U} := \{((x_0, y_0), r, (x, y)) \mid (x - x_0)^2 + (y - y_0)^2 = r^2\} \subset \mathcal{M} \times \mathbb{R}^2$ and π is the projection onto the first factor. This family enjoys the following properties, that clarify the vague point 2 above:

- (1) for any family of circles in the plane $p : E \rightarrow B$, there is a map $m : B \rightarrow \mathcal{M}$ defined by $m(b) = p^{-1}(b)$;

- (2) to every map $m : B \rightarrow \mathcal{M}$ there uniquely corresponds a family of circles parametrized by B , i.e.

$$\begin{array}{c} m^*\mathcal{U} \\ p \downarrow \\ B \end{array}$$

such that the fibre $p^{-1}(b)$ is the circle $m(b)$.

This is the best that we could have hoped for. In this case we say that \mathcal{M} is a **fine moduli space** with \mathcal{U} as its **universal family**.

Often, due to the presence of automorphisms of the parametrized objects, it is impossible to achieve this perfect bijection between families of objects and morphisms to the moduli space. If only property 1 holds we call the moduli space **coarse**.

2.2 Moduli of n Points on \mathbb{P}^1

Let us now consider the moduli space $M_{0,n}$ of all isomorphism classes of n ordered distinct marked points $p_i \in \mathbb{P}^1$. The subscript 0 is to denote the genus of our curve \mathbb{P}^1 . Since the automorphism group $Aut(\mathbb{P}^1) = PSL_2(\mathbb{C})$ allows us to move any three points on \mathbb{P}^1 to the ordered triple $(0, 1, \infty)$, the space $M_{0,n}$ reduces to a single point for $n \leq 3$.

Going one step up, $M_{0,4} = \mathbb{P}^1 - \{0, 1, \infty\}$: given a quadruple (p_1, p_2, p_3, p_4) , we can always perform the unique automorphism of \mathbb{P}^1 sending (p_1, p_2, p_3) to $(0, 1, \infty)$; the isomorphism class of the quadruple is then determined by the image of the fourth point.

The general case is similar. Any n -tuple $\underline{p} = (p_1, \dots, p_n)$ is equivalent to a n -tuple of the form $(0, 1, \infty, \phi(p_4), \dots, \phi(p_n))$, where ϕ is the unique automorphism of \mathbb{P}^1 sending (p_1, p_2, p_3) to $(0, 1, \infty)$. This shows

$$M_{0,n} = \overbrace{M_{0,4} \times \dots \times M_{0,4}}^{n-3 \text{ times}} \setminus \{\text{all diagonals}\}.$$

If we define $U_n := M_{0,n} \times \mathbb{P}^1$, then the projection of U_n onto the first factor gives rise to a universal family

$$\begin{array}{c} U_n \\ \pi \downarrow \uparrow \sigma_i \\ M_{0,n} \end{array}$$

where the σ_i 's are the universal sections:

- $\sigma_i(\underline{p}) = (\underline{p}, \phi(p_i)) \in U_n$.

This family is tautological since the fibre over a moduli point, which is the class of a marked curve, is the marked curve itself.

With U_n as its universal family, $M_{0,n}$ becomes a fine moduli space for isomorphism classes of n ordered distinct marked points on \mathbb{P}^1 .

This is all fine except $M_{0,n}$ is not compact for $n \geq 4$. There are many reasons why compactness is an extremely desirable property for moduli spaces. As an extremely practical reason, proper (and if possible projective) varieties are much better behaved and understood than non compact ones. Also, a compact moduli space encodes information on how our objects can degenerate in families. For example, what happens when $p_1 \rightarrow p_2$ in $M_{0,4}$?

In general there are many ways to compactify a space. A “good” compactification $\overline{\mathcal{M}}$ of a moduli space \mathcal{M} should have the following properties:

- (1) $\overline{\mathcal{M}}$ should be itself a moduli space, parametrizing some natural generalization of the objects of \mathcal{M} .
- (2) $\overline{\mathcal{M}}$ should not be a horribly singular space.
- (3) the boundary $\overline{\mathcal{M}} \setminus \mathcal{M}$ should be a normal crossing divisor.
- (4) it should be possible to describe boundary strata combinatorially in terms of simpler objects. This point may appear mysterious, but it will be clarified by the examples of stable curves and stable maps.

In the case of rational n -pointed curves there is a definite winner among compactifications.

2.3 Moduli of Rational Stable Curves

We will discuss the simple example of $M_{0,4}$; this hopefully will, without submerging us in combinatorial technicalities, provide intuition on the ideas and techniques used to compactify the moduli spaces of n -pointed rational curves.

A natural first attempt would be to just allow the points to come together, i.e. enlarge the collection of objects that we are considering from \mathbb{P}^1 with

n ordered distinct marked points to \mathbb{P}^1 with n ordered, not necessarily distinct, marked points.

However, this will not quite work. For instance, consider the families

$$C_t = (0, 1, \infty, t) \quad \text{and} \quad D_t = (0, t^{-1}, \infty, 1).$$

For each $t \neq 0$, up to an automorphism of \mathbb{P}^1 , $C_t = D_t$, thus corresponding to the same point in $M_{0,4}$. But for $t = 0$, C_0 has $p_1 = p_4$ whereas D_0 has $p_2 = p_3$. These configurations are certainly not equivalent up to an automorphism of \mathbb{P}^1 and so should be considered as distinct points in our compactification of $M_{0,4}$. Thus, we have a family with two distinct limit points (in technical terms we say that the space is **nonseparated**). This is not good.

Our failed attempt was not completely worthless though since it allowed us to understand that we want the condition $p_1 = p_4$ to coincide with $p_2 = p_3$, and likewise for the other two possible disjoint pairs. On the one hand this is very promising: 3 is the number of points needed to compactify $\mathbb{P}^1 \setminus \{0, 1, \infty\}$ to \mathbb{P}^1 . On the other hand, it is now mysterious what modular interpretation to give to this compactification.

To do so, let us turn carefully to our universal family, illustrated in Figure 2. The natural first step is to fill in the three points on the base, to complete U_4 to $\mathbb{P}^1 \times \mathbb{P}^1$ and extend the sections by continuity.

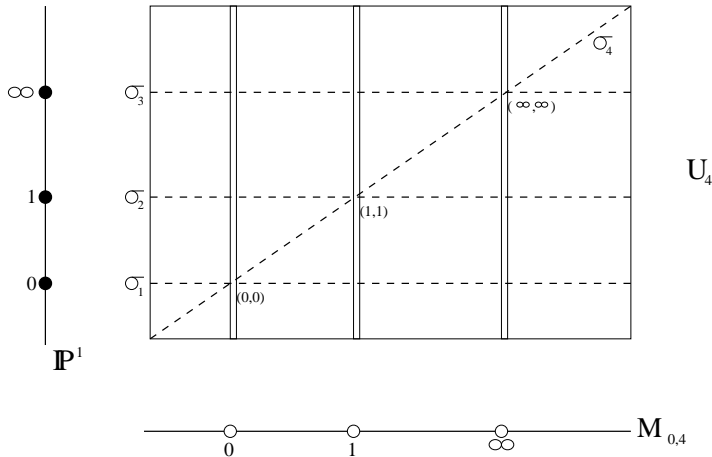


Figure 2: first attempt at compactifying U_4

We immediately notice a bothersome asymmetry in this picture: the point p_4 is the only one allowed to come together with all the other points: yet common sense, backed up by the explicit example just presented, suggests that there should be democracy among the four points. This fails where the diagonal section σ_4 intersects the three constant ones, i.e. at the three points $(0, 0)$, $(1, 1)$, (∞, ∞) . Let us blow-up $\mathbb{P}^1 \times \mathbb{P}^1$ at these three points. This will make all the sections disjoint, and still preserve the smoothness and projectivity of our universal family.

The fibres over the three exceptional points are $\mathbb{P}^1 \cup E_i$: nodal rational curves. These are the new objects that we have to allow in order to obtain a good compactification of $M_{0,4}$.

Let us finally put everything together, and state things carefully.

Definition 1 *A tree of projective lines is a connected curve with the following properties:*

- (1) *Each irreducible component is isomorphic to \mathbb{P}^1 .*
- (2) *The points of intersection of the components are ordinary double points.*
- (3) *There are no closed circuits, i.e., if a node is removed then the curve becomes disconnected.*

These three properties are equivalent to saying that the curve has arithmetic genus zero. Each irreducible component will be called a twig. We will often draw a marked tree as in fig 3, where each line represents a twig.

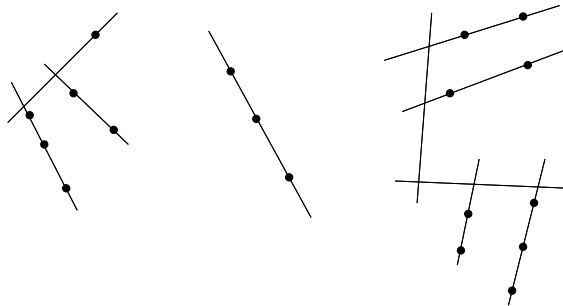


Figure 3: stable marked trees.

Definition 2 A marked tree is **stable** if every twig has at least three special points (marks or nodes).

This stability condition is equivalent to the existence of no nontrivial automorphisms of the tree that fix all of the marks.

Definition 3 $\overline{M}_{0,4} \cong \mathbb{P}^1$ is the moduli space of isomorphism classes of four pointed stable trees. It is a fine moduli space, with universal family $U_4 = Bl(\mathbb{P}^1 \times \mathbb{P}^1)$.

These results generalize to larger n .

Fact: The space $\overline{M}_{0,n}$ of n -pointed rational stable curves is a fine moduli space compactifying $M_{0,n}$. It is projective, and the universal family \overline{U}_n is obtained from U_n via a finite sequence of blow-ups. (In particular all the diagonals need to be blown up in an appropriate order) . For further details see [KV99] or [Knu83a], [Knu83b].

One of the exciting features of this theory is that all these spaces are related to one another by natural morphisms. Consider the map

$$\pi_i : \overline{M}_{0,n+1} \rightarrow \overline{M}_{0,n},$$

defined by forgetting the i th mark. It is obviously defined if the i th mark does not belong to a twig with only three special points. If it does belong to such a twig, then our resulting tree will no longer be stable. In this case, we must perform what is called contraction.

Contraction: We need to consider two cases:

- (1) The remaining two special points are both nodes. We make the tree again stable by contracting this twig so that the two nodes are now one (see Figure 4).
- (2) There is one other mark and one node on the twig in question. We make the tree stable by forgetting the twig and placing the mark where the node used to be (Figure5).

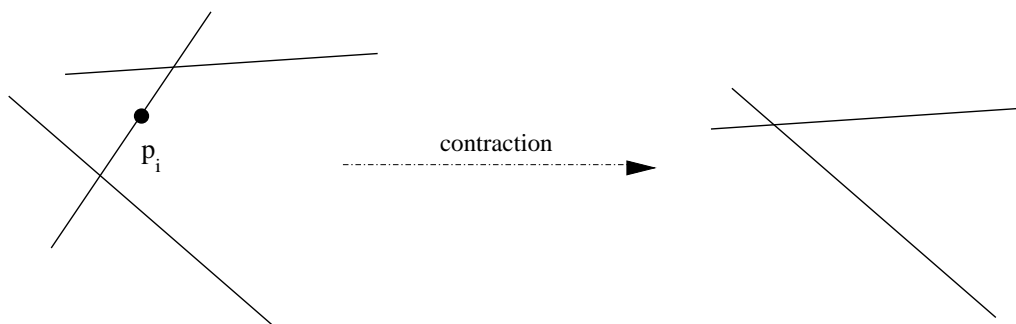


Figure 4: contracting a twig with only two nodes.

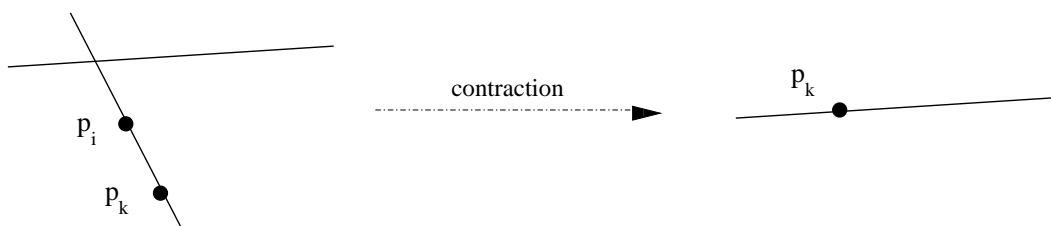


Figure 5: contracting a twig with one node and one mark.

We would like to construct a section σ_i of the family

$$\begin{array}{c} \overline{M}_{0,n+1} \\ \pi_k \downarrow \uparrow \sigma_i \\ \overline{M}_{0,n} \end{array}$$

by defining the k th mark to coincide with i th one. It should trouble you that in doing so we are not considering curves with distinct marked points, but we can get around this problem by “sprouting” a new twig so that the node is now where the i th mark was. The k th and the i th points now belong to this new twig.

This process of making stable a tree with two coinciding points is called **stabilization**.

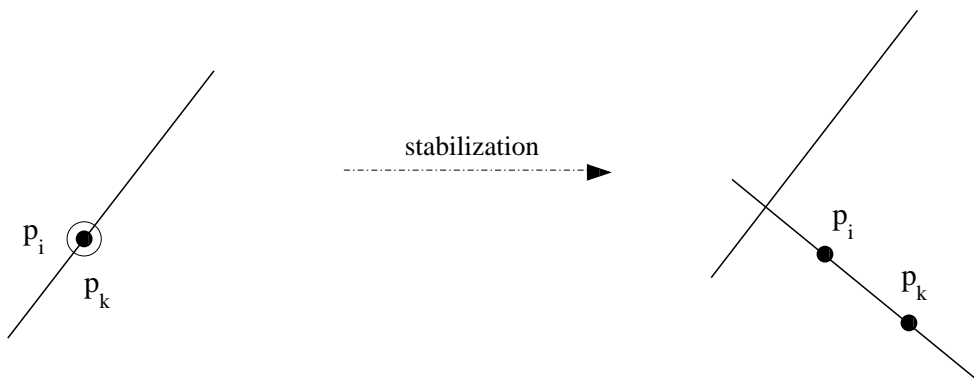


Figure 6: stabilization

Finally, we may now identify our universal family

$$\begin{array}{ccc} \overline{U}_n & & \overline{M}_{0,n+1} \\ \pi \downarrow & \text{with the family} & \pi_i \downarrow \\ \overline{M}_{0,n} & & \overline{M}_{0,n} \end{array}$$

as follows.

The fibre $\pi^{-1}([(C, p_1, \dots, p_n)]) \subset \overline{U}_n$ is the marked curve itself. So any point $p \in \overline{U}_n$ belonging to the fibre over C is actually a point on the stable n -pointed tree C , and may therefore be considered as an additional mark; stabilization may be necessary to ensure that our new $(n+1)$ -marked tree is stable. Vice-versa, given an $(n+1)$ pointed curve C' , we can think

of the $(n + 1)$ st point as being a point on the n -marked curve obtained by forgetting the last marked point (eventually contracting, if needed); this way C' corresponds to a point on the universal family \overline{U}_n . These constructions realize an isomorphism between \overline{U}_n and $\overline{M}_{0,n+1}$.

2.3.1 The boundary

We define the boundary to be the complement of $M_{0,n}$ in $\overline{M}_{0,n}$. It consists of all nodal stable curves.

Fact: the boundary is a union of irreducible components, corresponding to the different possible ways of arranging the marks on the various twigs; the codimension of a boundary component equals the number of nodes in the curves in that component. See [KV99] for more details.

The codimension 1 boundary strata of $\overline{M}_{0,n}$, called the **boundary divisors**, are in one-to-one correspondence with all ways of partitioning $[n] = A \cup B$ with the cardinality of both A and B strictly greater than 1.

A somewhat special class of boundary divisors consists of those with only two marked points on a twig. Together, these components are sometimes called the **soft boundary** and denoted by $D_{i,j}$. We can think of $D_{i,j}$ as the image of the i th section, σ_i , of the j th forgetful map, π_j (or vice-versa).

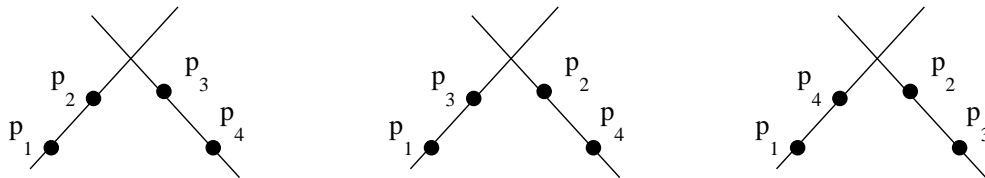


Figure 7: irreducible components of the boundary of $\overline{M}_{0,4}$

There is plenty more to be said about the spaces $\overline{M}_{0,n}$, their relationships, and their boundaries, but we will leave our treatment of $\overline{M}_{0,n}$ here, suggesting [KV99] as an excellent reference for beginners. In Figures 7 and 8 we draw all boundary strata for $\overline{M}_{0,4}$ and $\overline{M}_{0,5}$.

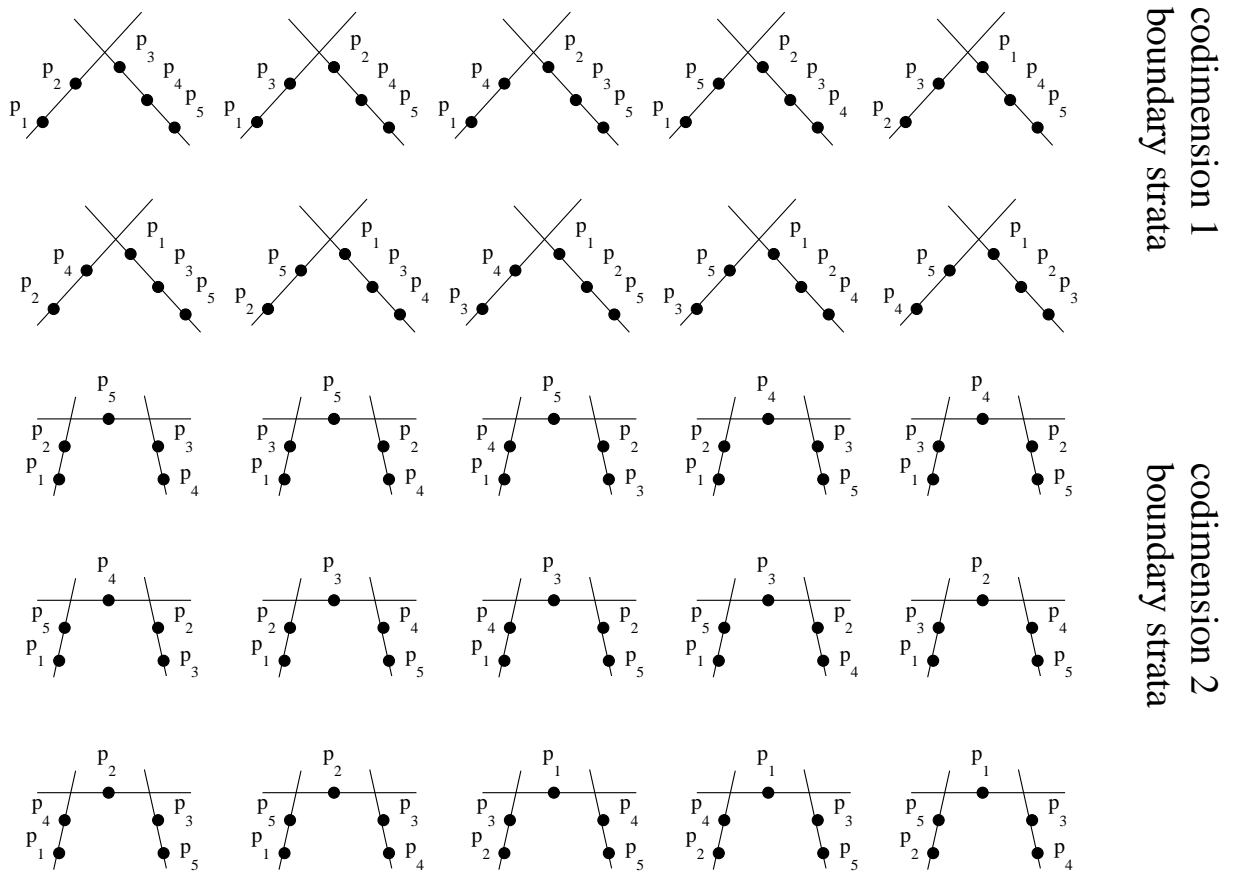


Figure 8: boundary cycles of $\overline{M}_{0,5}$

2.4 Moduli of Rational Stable Maps

Let us now move on to the moduli spaces of greatest interest for solving the bitangent problem. We would like to study, in general, rational curves in projective space. The characteristic property of an irreducible rational curve is that it can be parametrized by the projective line, \mathbb{P}^1 . For this reason, it is natural to study maps $\mu : \mathbb{P}^1 \rightarrow \mathbb{P}^r$.

When we talk about the **degree** of such a map we mean the degree of $\mu_*[\mathbb{P}^1]$ as in homology. Be careful, the degree of the map may be different from the degree of the image curve! For example the map

$$\begin{aligned} \mu : \mathbb{P}^1 &\rightarrow \mathbb{P}^2 \\ (x_0 : x_1) &\mapsto (x_0^2 : x_1^2 : 0) \end{aligned}$$

has degree two, but its image is a line.

Define $W(r, d)$ as the space of all maps from \mathbb{P}^1 to \mathbb{P}^r of degree d . A map in $W(r, d)$ is specified, up to a constant, by $r + 1$ binary forms of degree d that do not all vanish at any point. It can then be seen that $\dim W(r, d) = (r + 1)(d + 1) - 1$.

We also have the family

$$\begin{array}{ccc} W(r, d) \times \mathbb{P}^1 & \xrightarrow{\rho} & \mathbb{P}^r \\ \downarrow & & \\ W(r, d) & & \end{array}$$

where $\rho(\mu, x) = \mu(x)$. This family is tautological in the sense that the fibre over the map μ is the map $\rho|_{\{\mu\} \times \mathbb{P}^1} = \mu$. In fact, this is a universal family. Thus, $W(r, d)$ is a fine moduli space for maps $\mathbb{P}^1 \rightarrow \mathbb{P}^r$ of degree d .

However, $W(r, d)$ is not the moduli space that we would like to study. For one, it is not compact. For another, reparametrizations of the source curve are considered as different points in $W(r, d)$. To fix the latter problem, let us simply consider the space $M_{0,0}(\mathbb{P}^r, d) := W(r, d)/Aut(\mathbb{P}^1)$. For a detailed account on why this quotient is indeed a space, see chapter five of [HM98]

Another way to eliminate automorphisms is to consider n -pointed maps (maps $\mu : C \rightarrow \mathbb{P}^r$ with an n -marked source $C \simeq \mathbb{P}^1$). It should be no surprise that there is a fine moduli space $M_{0,n}(\mathbb{P}^r, d)$ for isomorphism

classes of n -pointed maps $\mathbb{P}^1 \rightarrow \mathbb{P}^r$ of degree d , namely $M_{0,n} \times W(r, d)$. But we still have not dealt with the non-compactness of this moduli space. The idea is to parallel the construction that led us to stable curves.

Definition 4 An ***n -pointed stable map*** is a map $\mu : C \rightarrow \mathbb{P}^r$, where:

- (1) C is a n -marked tree.
- (2) Every twig in C mapped to a point must have at least three special points on it.

Fact: Moduli spaces of n -pointed rational stable maps to \mathbb{P}^r of degree d (denoted $\overline{M}_{0,n}(\mathbb{P}^r, d)$) can be constructed; they compactify the moduli spaces of smooth maps. It is straightforward to verify that an n -pointed map is stable if and only if it has only a finite number of automorphisms. Unfortunately, there is no way to eliminate all nontrivial automorphism. Details can be found in [KV99].

Example: An element $\mu \in \overline{M}_{0,2}(\mathbb{P}^2, 2)$ that is the double cover of a line, marking the ramification points, admits a nontrivial automorphism exchanging the two covers. This allows us to construct a nontrivial family of maps μ_t that maps constantly to one point in the moduli space. Consider:

$$\begin{aligned} \mu_t : [0, 1]/\{0 = 1\} \times \mathbb{P}^1 &\longrightarrow \mathbb{P}^2 \\ (t, (x_0 : x_1)) &\longmapsto (0 : x_0^2 : e^{2\pi it} x_1^2). \end{aligned}$$

Because of this phenomenon there is no universal family associated to the spaces $\overline{M}_{0,n}(\mathbb{P}^r, d)$, and the corresponding moduli spaces are only coarse.

Since $M_{0,n}(\mathbb{P}^r, d)$ is dense in $\overline{M}_{0,n}(\mathbb{P}^r, d)$, the latter has dimension

$$(r + 1)(d + 1) - 1 + (n - 3) = rd + r + d + n - 3.$$

Example: in particular, $\overline{M}_{0,n}(\mathbb{P}^r, 0) = \overline{M}_{0,n} \times \mathbb{P}^r$.

There are some useful maps among the spaces $\overline{M}_{0,n}(\mathbb{P}^r, d)$. For instance, as with the spaces $\overline{M}_{0,n}$, we have the forgetful maps π_i defined by simply forgetting the i th mark and the sections, σ_j , of the family

$$\begin{array}{c} \overline{M}_{0,n+1}(\mathbb{P}^r, d) \\ \pi_i \downarrow \uparrow \sigma_j \\ \overline{M}_{0,n}(\mathbb{P}^r, d) \end{array}$$

defined by declaring the j th and the i th mark to coincide. Contraction and stabilization are performed to make these maps defined everywhere.

In addition, there are **evaluation maps**

$$\nu_i : \overline{M}_{0,n}(\mathbb{P}^r, d) \rightarrow \mathbb{P}^r$$

defined by $\nu_i(\mu) = \mu(p_i)$ where p_i is the i th mark on the source curve C .

The forgetful and evaluation morphisms allow us to identify $\overline{M}_{0,n+1}(\mathbb{P}^r, d)$ as a tautological family for $\overline{M}_{0,n}(\mathbb{P}^r, d)$:

$$\begin{array}{ccc} \overline{M}_{0,n+1}(\mathbb{P}^r, d) & \xrightarrow{\nu_{n+1}} & \mathbb{P}^r \\ \pi_{n+1} \downarrow & & \\ \overline{M}_{0,n}(\mathbb{P}^r, d) & & \end{array}$$

This way we can think of points of $\overline{M}_{0,n+1}(\mathbb{P}^r, d)$ either as $n + 1$ pointed maps or as points on an n -marked curve mapped to \mathbb{P}^r . Being comfortable with this identification comes in very handy when making computations.

2.4.1 The Boundary

The boundary consists of maps whose domains are reducible curves. In fact, its description is very similar to that of $\overline{M}_{0,n}$. Boundary strata are determined now not only by the combinatorial data of the arrangement of the marks, but also by the degree the maps restrict to on each twig.

Boundary divisors are in one to one correspondence with all ways of partitioning $[n] = A \cup B$ and $d = d_A + d_B$ such that:

- $\#A \geq 2$ if $d_A = 0$;
- $\#B \geq 2$ if $d_B = 0$.

2.5 Psi classes

Consider a family of curves admitting a section.

$$\begin{array}{ccc} & \mathcal{C} & \\ \pi \downarrow \uparrow \sigma & & \\ & B & \end{array}$$

We can define the **cotangent line bundle**, \mathbb{L}_σ , as the line bundle on B whose fibre at a point $b \in B$ is the cotangent space of $\mathcal{C}_b = \pi^{-1}(b)$ at the point $\sigma(b)$. We call the ψ **class** of the family the first Chern class of this line bundle. Observe that, for a trivial family of curves with a constant section, the ψ class vanishes.

This construction can be extended in a natural way to the moduli space $\overline{M}_{0,n}(\mathbb{P}^r, d)$.

Informally, we have a sheaf on the tautological family $\overline{M}_{0,n+1}(\mathbb{P}^r, d)$ whose local sections away from nodes are differential forms on the curves. We obtain it by considering 1-forms on the tautological family and quotienting by forms that are pulled back from the moduli space. This sheaf is called the **relative dualizing sheaf**³, and denoted by $\omega_{\pi_{n+1}}$.

Consider now the i th tautological section σ_i . If we restrict $\omega_{\pi_{n+1}}$ to this section, we obtain a line bundle on the moduli space whose fibres are naturally identified with the cotangent spaces of the curves at the i th marked point. Then we can define the class:

$$\psi_i := \mathbf{c}_1(\sigma_i^* \omega_{\pi_{n+1}}) \in A^1(\overline{M}_{0,n}(\mathbb{P}^r, d)).$$

The construction of ψ_i is natural in the sense that if we have a family of pointed stable maps, inducing a morphism to the moduli space, the ψ_i class of the family is the pull-back of the ψ_i class on the moduli space.

It may seem that there is no difference between the information carried by $\psi_i \in A^1(\overline{M}_{0,n+1}(\mathbb{P}^r, d))$ and $\psi_i \in A^1(\overline{M}_{0,n}(\mathbb{P}^r, d))$. We have a natural map between these two spaces, the tautological family $\pi_{n+1} : \overline{M}_{0,n+1}(\mathbb{P}^r, d) \rightarrow \overline{M}_{0,n}(\mathbb{P}^r, d)$. It may seem that $\mathbb{L}_{i,n+1} := \sigma_i^* \omega_{\pi_{n+1}}$ and $\pi_{n+1}^* \mathbb{L}_{i,n}$ are the same line bundle, thus yielding the relation $\psi_i = \pi_{n+1}^* \psi_i$. In reality, this is almost true. Surely these line bundles agree off the component $D_{i,n+1}$ of the soft boundary. From this consideration, we conclude

$$\mathbb{L}_{i,n+1} = \pi_{n+1}^* \mathbb{L}_{i,n} \otimes \mathcal{O}(mD_{i,n+1}) \tag{4}$$

³The word relative refers to the fact that we are quotienting by everything coming from downstairs. In other words, we are constructing a sheaf on the universal family of the moduli space by “gluing” together sheaves defined on the curves.

for some integer m .

Next, observe that $\mathbb{L}_{i,n+1}$ restricted to $D_{i,n+1}$ is a trivial line bundle: we are looking at curves with a twig having only three special points; the node, the i th and the $(n+1)$ st mark. By an automorphism of the twig, we can assume that the node is at 0 and the two marks are at 1 and ∞ . Therefore, this line bundle restricted to $D_{i,n+1}$ is the cotangent space at a single unchanging point of \mathbb{P}^1 . This implies

$$\begin{aligned} \mathcal{O}_{D_{i,n+1}} &= \mathbb{L}_{i,n+1}|_{D_{i,n+1}} = \\ &= (\pi_{n+1}^* \mathbb{L}_{i,n} \otimes \mathcal{O}(mD_{i,n+1}))|_{D_{i,n+1}} = \mathbb{L}_{i,n} \otimes \mathcal{O}(mD_{i,n+1})|_{D_{i,n+1}}. \end{aligned}$$

By the adjunction formula ([GH94], page 146), the line bundle $\mathcal{O}(D_{i,n+1})|_{D_{i,n+1}}$ is the normal bundle of the divisor $D_{i,n+1}$.

But the normal directions to a section in the moduli space are precisely the tangent directions to the fibres. Hence $\mathcal{O}(D_{i,n+1})|_{D_{i,n+1}}$ is the dual to the relative cotangent bundle $\mathbb{L}_{i,n}$. It follows that m must be 1.

Finally, by taking Chern classes in (3), we can deduce the fundamental relation:

$$\psi_i = \pi_{n+1}^* \psi_i + D_{i,n+1}. \quad (5)$$

The above pull-back relation can be used to describe explicitly ψ classes for moduli spaces of rational stable curves in terms of boundary strata. A closed formula can be found in [Koc01].

3 Counting Bitangents

3.1 The Strategy

We now have all the necessary machinery to tackle our problem of counting bitangents. Before we start digging deep into details and computations, let us outline our strategy.

Let $Z := \{f = 0\}$ be a projective plane curve of degree d :

- we consider the moduli space $\overline{M}_{0,2}(\mathbb{P}^2, 1)$, of two-pointed, rational stable maps of degree one;

- we construct a jet bundle on this space with the property that the zero set of a section of this bundle consists of stable maps having at least second order contact with Z at the image of the i th marked point; we name this cycle $\Phi_i(Z)$;
- we represent $\Phi_i(Z)$ in the Chow ring in terms of ψ classes and other natural classes;
- we step by step compute the intersection $\Phi_1(Z)\Phi_2(Z)$;
- we identify and clean up some garbage that lives in that intersection and corresponds to maps that are not bitangents;
- finally, we have counted two pointed maps that are tangent to Z at each mark; we just need to divide by 2, since we are not interested in the ordering of the marks.

Easy enough? Now let us start over slowly and do everything carefully.

3.2 $\overline{M}_{0,1}(\mathbb{P}^2, 1)$

This space has dimension 3, and it is explicitly realized by the following incidence relation:

$$\overline{M}_{0,1}(\mathbb{P}^2, 1) = \{(p, l) \in \mathbb{P}^2 \times \check{\mathbb{P}}^2 \mid p \in l\} =: \mathcal{I} \subset \mathbb{P}^2 \times \check{\mathbb{P}}^2.$$

There are two projections of \mathcal{I} onto \mathbb{P}^2 and $\check{\mathbb{P}}^2$, that we will denote q and \check{q} . The latter makes \mathcal{I} into a tautological family of lines in \mathbb{P}^2 :

$$\begin{array}{ccc} \mathbb{P}^1 & \longrightarrow & \mathcal{I} \\ & & \downarrow \check{q} \\ & & \check{\mathbb{P}}^2. \end{array}$$

This family is tautological in the sense that the fibre over $l \in \check{\mathbb{P}}^2$ is l itself.

A fibre over $p \in \mathbb{P}^2$ under q is the pencil of lines in \mathbb{P}^2 passing through p and so this projection also gives rise to a \mathbb{P}^1 bundle over \mathbb{P}^2 . Observe that q is precisely the evaluation map ν_x .

Notation: We denote by x the unique mark in the space of one-pointed maps, and add the subscript x to any entity (class, map ...) related to it. We do so to keep track of the conceptual difference from the marked points on the two-pointed maps, which will be numbered 1 and 2.

Definition 5 We identify and name two natural divisors on $\overline{M}_{0,1}(\mathbb{P}^2, 1)$.

$\iota(p)$: Let us look at the hyperplane divisor $\check{p} \subset \check{\mathbb{P}}^2$ of all lines passing through a point p , and consider the cycle of its pull-back $\check{q}^*(\check{p}) := \iota(p)$.

$\eta_x(l)$: Similarly, we look at the hyperplane divisor $l \subset \mathbb{P}^2$ and define $\eta_x(l) := q^*(l)$ as its pull-back under q .

In general, define $\eta_x(Z) := q^*(Z) = \nu_x^*(Z)$ as the cycle of maps whose mark is sent into Z .

There is only one class of a line and only one of a point in $A^*(\mathbb{P}^2)$, hence $\eta_x := [\eta_x(l)]$ and $\iota := [\iota(p)]$ are independent of l and p respectively.

Since $\overline{M}_{0,1}(\mathbb{P}^2, 1)$ is a \mathbb{P}^1 bundle over \mathbb{P}^2 , it follows that ι and η_x , i.e. the pull-backs of hyperplane divisors in the base and in the fiber, generate $A^1(\overline{M}_{0,1}(\mathbb{P}^2, 1))$. It is therefore useful to know all intersections of the two classes.

It is a good exercise to construct a picture verifying each of the following statements about classes.

- $\eta_x = [\{(p', l') \mid p' \in l, l' \in \check{p}' \text{ with } l \text{ fixed}\}]$
- $\iota = [\{(p', l') \mid l' \in \check{p}, p' \in l' \text{ with } p \text{ fixed}\}]$
- $\eta_x^2 = [\{(p, l') \mid l' \in \check{p} \text{ with } p \text{ fixed}\}]$
- $\iota\eta_x = [\{(p', l') \mid p' \in l, l \in \check{p} \text{ with } p \text{ and } l \text{ fixed}\}]$
- $\iota^2 = [\{(p', l) \mid p' \in l \text{ with } l \text{ fixed}\}]$
- $\iota\eta_x^2 = [\{(p, l') \mid l \in \check{p} \cap \check{p}_0 \text{ with } p, p_0 \text{ fixed}\}] = \mathbf{pt}$ since there is exactly one line passing through two distinct points.
- $\eta_x\iota^2 = [\{(p', l) \mid p' \in l \cap l_0 \text{ with } l, l_0 \text{ fixed}\}] = \mathbf{pt}$ since there is exactly one point in the intersections of two distinct lines.
- $\eta_x^3 = [\{(p', l') \mid l' \in \check{p}_1 \cap \check{p}_2 \cap \check{p}_3, p' \in l' \text{ with } p_1, p_2, p_3 \text{ fixed}\}] = 0$ since in general three points do not lie on a common line.
- $\iota^3 = [\{(p', l') \mid p' \in l_1 \cap l_2 \cap l_3, p' \in l' \text{ with } l_1, l_2, l_3 \text{ fixed}\}] = 0$ since in general three lines do not share a common point.

The following two lemmas prove identities that will be crucial for our later computations.

Lemma 1

$$\eta_x \iota = \eta_x^2 + \iota^2. \quad (6)$$

PROOF: We construct a one-parameter family of cycles, parametrized by $[0, 1]$, with the left-hand side of our identity as one endpoint of this family and the right-hand side as the other.

To choose a representative of the class $\eta_x \iota$, one must specify a fixed point p and a fixed line l . Let us fix l once and for all and let p_t be a path in \mathbb{P}^2 such that $p_t \in l$ if and only if $t = 0$. Our one parameter family α_t is defined as follows:

$$\alpha_t = \{(p', l') \mid p' \in l, l' \in \check{p}_t\}.$$

Notice that $[\alpha_t] = \eta_x \iota$ for $t \neq 0$.

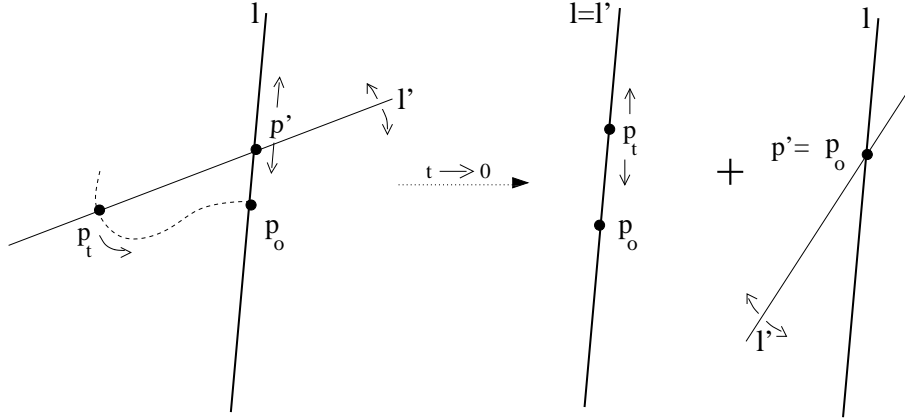


Figure 9: $\eta_x \iota = \eta_x^2 + \iota^2$.

We now examine what happens as $t \rightarrow 0$.

$p' \neq p_0$: necessarily, $l' = l$ and our resulting one dimensional class is parametrized solely by $p' \in l$ and is therefore ι^2 .

$p' = p_0$: Then l' is only required to be in \check{p}_0 and so such l' 's in \check{p}_0 parametrize our resulting one dimensional class. We have arrived at η_x^2 .

We now have that $\alpha_0 = \eta_x^2 + \iota^2$, which allows us to conclude (6).

Lemma 2

$$\psi_x = \iota - 2\eta_x. \tag{7}$$

PROOF: As $\psi_x \in A^1(\overline{M}_{0,1}(\mathbb{P}^2, 1))$, it is possible to express $\psi_x = a\iota + b\eta_x$ for some integers a and b . Let us determine these two integers.

a: intersecting ψ_x with η_x^2 we obtain

$$\psi_x \eta_x^2 = a\iota \eta_x^2 + b\eta_x^3 = a.$$

Consider $\sigma_x^* \omega_{\pi_2}$ restricted to $\eta_x^2 = \{(p, l') \mid p \text{ is fixed and } l' \in \check{p}\}$. It is the line bundle over η_x^2 whose fibre over a point $(p, l') \in \eta_x^2$ is the cotangent space of l' at the fixed point p .

It is worth convincing yourself that this is the dual of the tautological line bundle

$$\begin{array}{c} \mathcal{S} \\ \downarrow \pi \\ \mathbb{P}^1 = \{l \subset \mathbb{P}^2 \mid p \in l\} = \eta_x^2. \end{array}$$

We computed (section 1.2.2) that $c_1(\mathcal{S}) = -1$ and so $a = 1$.

b: similarly, b is the product $\psi_x \iota^2$. To find this intersection we must consider the line bundle $\sigma_x^* \omega_{\pi}$ restricted to $\iota^2 = \{p \in l \mid l \text{ is fixed}\}$. A fibre of this line bundle over a point $p \in l$ is the cotangent space of our fixed l at p . This is simply the cotangent bundle of l .

Since $l = \mathbb{P}^1 = S^2$ has Euler characteristic 2, then the degree of the first Chern class of the cotangent bundle is -2 . We thus have that $b = -2$.

3.3 $\overline{M}_{0,2}(\mathbb{P}^2, 1)$

First off note that $\dim \overline{M}_{0,2}(\mathbb{P}^2, 1) = 4$. The description of $\overline{M}_{0,2}(\mathbb{P}^2, 1)$ is slightly more complicated largely due to the existence of its only boundary divisor which we call β .

For a two-pointed map μ not in the boundary, $\mu(p_1) \neq \mu(p_2)$. Since we are considering maps of degree 1, i.e. isomorphisms of lines, $\mu(p_1)$ and $\mu(p_2)$ completely determine (up to reparametrization) our map μ . It follows that

$$\overline{M}_{0,2}(\mathbb{P}^2, 1) \setminus \beta = \mathbb{P}^2 \times \mathbb{P}^2 \setminus \Delta.$$

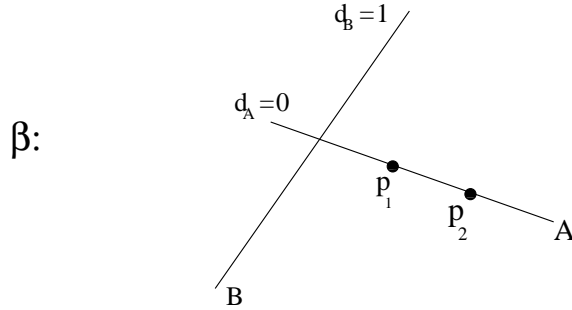


Figure 10: the boundary divisor in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$

On β , $\mu(p_1) = \mu(p_2)$: for any line l through p , there is a map in β that contracts the twig with the marks to p and maps the other twig isomorphically to l .

So for our description to be complete, we need to replace $(p, p) \in \Delta \subset \mathbb{P}^2 \times \mathbb{P}^2$ with a \mathbb{P}^1 worth of maps. We arrive at

$$\overline{M}_{0,2}(\mathbb{P}^2, 1) = Bl_{\Delta}(\mathbb{P}^2 \times \mathbb{P}^2).$$

Consider the tautological families

$$\begin{array}{c} \overline{M}_{0,2}(\mathbb{P}^2, 1) \\ \pi_1 \downarrow \downarrow \pi_2 \\ \overline{M}_{0,1}(\mathbb{P}^2, 1). \end{array}$$

Both families have a natural common section $\sigma := \sigma_1 = \sigma_2$. The image of σ is the unique boundary divisor β in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$.

Define $\eta_i(Z) := \nu_i^*(Z)$ It's easy to check the following identities:

- (1) $\sigma_*(\eta_x(Z)) = \beta\eta_1(Z) = \beta\eta_2(Z)$.
- (2) $\sigma^*(\eta_i(Z)) = \eta$.
- (3) $\pi_i^*\eta_x(Z) = \eta_i(Z)$.

3.4 Tangency Conditions

Let us define $\Phi_i(Z) \in A^*(\overline{M}_{0,n}(\mathbb{P}^2, 1))$ as the cycle of maps tangent to a plane curve $Z = \{f = 0\}$ at the image of the i th marked point. Formally,

$$\Phi_i(Z) = \{\mu \in \overline{M}_{0,n}(\mathbb{P}^2, 1) \mid \mu^* f \text{ vanishes at } p_i \text{ with multiplicity } \geq 2\}.$$

We now want to obtain an expression for $\Phi_x(Z) \in A^*(\overline{M}_{0,1}(\mathbb{P}^2, 1))$ in terms of η_x , β , and ψ_x .

Consider the tautological family

$$\begin{array}{ccc} \overline{M}_{0,2}(\mathbb{P}^2, 1) & \xrightarrow{\nu_2} & \mathbb{P}^2 \\ \sigma \uparrow \downarrow \pi_2 & \nearrow & \nu_x \\ \overline{M}_{0,1}(\mathbb{P}^2, 1). & & \end{array}$$

Let us pull-back the line bundle $\mathcal{O}(Z)$ via ν_2 , and consider the first jet bundle $J_{\pi_2}^1 \nu_2^* \mathcal{O}(Z)$ relative to π_2 . Relative here means that we quotient out by everything that can be pulled back from $\overline{M}_{0,1}(\mathbb{P}^2, 1)$. Let Z be defined by the vanishing of the polynomial f , and let us consider the zero locus of the section $\tau := \nu_2^* f + (\partial_{\pi_2}^1 \nu_2^* f) dt \in \Gamma(J_{\pi_2}^1 \nu_2^* \mathcal{O}(Z))$; what we obtain is the locus of maps in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$ such that the pull-back of f at the second marked point vanishes to second order. If we interpret $\overline{M}_{0,2}(\mathbb{P}^2, 1)$ as the universal family for $\overline{M}_{0,1}(\mathbb{P}^2, 1)$ ⁴, it follows that to obtain $\Phi_x(Z)$, the locus in $\overline{M}_{0,1}(\mathbb{P}^2, 1)$ of lines tangent to Z at the unique marked point, we need to pull-back via the section σ .

In formulas, this translates to

$$\Phi_x(Z) = \mathbf{e}(\sigma^* J_{\pi_2}^1 \nu_2^* \mathcal{O}(Z)).$$

Since the rank of the bundle in question is 2, the Euler class is the second Chern class.

To calculate $\mathbf{c}_2(\sigma^* J_{\pi_2}^1 \nu_2^* \mathcal{O}(Z))$ we use the following short exact sequence discussed in section 1.4.

$$0 \rightarrow \nu_2^* \mathcal{O}(Z) \otimes \omega_{\pi_2} \rightarrow J_{\pi_2}^1 \nu_2^* \mathcal{O}(Z) \rightarrow \nu_2^* \mathcal{O}(Z) \rightarrow 0.$$

Notice that the first and last terms of this sequence are line bundles. We then want to consider the pull-back along σ of this exact sequence. Using the Whitney formula, we now have that

$$\Phi_x(Z) = \mathbf{c}_1(\sigma^* \nu_2^* \mathcal{O}(Z)) \mathbf{c}_1(\sigma^* \nu_2^* \mathcal{O}(Z) \otimes \sigma^* \omega_{\pi_2}) = d\eta_x(d\eta_x + \psi_x)$$

⁴this is true because one pointed maps of degree one have no nontrivial automorphisms!

in $\overline{M}_{0,1}(\mathbb{P}^2, 1)$.

The last equality follows from the two facts:

- $[Z] = d\mathbf{H} \in A^1(\mathbb{P}^2)$, where \mathbf{H} is the hyperplane class generating $A^*(\mathbb{P}^2)$ and $d = \text{deg}f$;
- $\nu_x = \nu_2\sigma$, and η_x is by definition $\nu_x^*(\mathbf{H})$.

Now we want to consider the case when we have more than one marked point: let us say we want to compute $\Phi_1(Z)$ in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$. The obvious guess is $\Phi_1(Z) = d\eta_1(d\eta_1 + \psi_1)$. We need to be careful, though: for maps in $\beta\eta_1(Z) \subset d\eta_1(d\eta_1 + \psi_1)$, the whole twig containing the two marks is mapped to Z . So, for μ such a map, μ^*f vanishes identically along the contracting twig and thus to all orders. We do not want to consider these maps as tangents to Z .

Fact: This simple correction works. The formula in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$ is

$$\Phi_1(Z) = d\eta_1(d\eta_1 + \psi_1 - \beta). \quad (8)$$

Lastly, note that $\beta\Phi_1(Z) = \sigma_*\Phi_x(Z)$.

Remark: It would be nice to use higher order jet bundles to describe cycles of maps having higher order contact with our curve Z . Unfortunately in general it is quite difficult, as fairly complicated problems of excess intersection arise.

For our application, we only need to push our luck a little further: we need to describe the cycle $\Phi_x^{(3)}(Z)$ of inflection tangents to Z . Luckily, thanks to the fact that $\overline{M}_{0,1}(\mathbb{P}^2, 1)$ has no boundary, the argument carries through: $\Phi_x^{(3)}(Z)$ can be computed as the Euler class of the bundle $E' := \sigma^*J_{\pi_2}^2\nu_2^*\mathcal{O}(Z)$. Here, $\text{rank}E' = 3$. By the exact sequence (3) and the Whitney formula:

$$\Phi_x^{(3)}(Z) = \mathbf{c}_3(E') = \eta_x(Z)(\eta_x(Z) + \psi_x)(\eta_x(Z) + 2\psi_x).$$

Using our previous calculations,

$$\Phi_x^{(3)}(Z) = (3d^2 - 6d)\mathbf{pt}. \quad (9)$$

We will abuse notation from now on and leave off the class of a point in our calculations. When writing a dimension zero class we will simply write its integral over the fundamental class.

3.5 The Computation

Let us finally get down to business. We define in general

$$\Lambda_Z(m_1p_1 + m_2p_2) \subset \overline{M}_{0,2}(\mathbb{P}^2, 1)$$

as the cycle of maps μ , such that $\mu^*\mathbf{Z} \geq m_1p_1 + m_2p_2$. Note that $\Lambda_Z(2p_1) = \Phi_1(Z)$.

Our ultimate goal is to compute $\Lambda_Z(2p_1 + 2p_2)$, i.e. the class of maps in $\overline{M}_{0,2}(\mathbb{P}^2, 1)$ which are tangent to Z at both p_1 and p_2 .

Note that $\dim \Lambda_Z(2p_1 + 2p_2) = 0$ since we are imposing 4 independent conditions in a space of dimension 4. This tells us that our enumerative problem makes sense.

On first thought, one might suggest $\Lambda_Z(2p_1 + 2p_2) = \Phi_1(Z)\Phi_2(Z)$. After all, $\Phi_1(Z)$ is the class of all maps tangent to Z at p_1 and similarly for $\Phi_2(Z)$ at p_2 so their intersection seems to be what we are after. However, one must be careful: for example, $\beta\Phi_1(Z)$ is in this intersection and we don't want to consider such maps.

We will proceed in two steps and obtain $\Lambda_Z(2p_1 + 2p_2)$ from $\Phi_1(Z)\Phi_2(Z)$ by "throwing away the spurious maps".

Step 1: We consider the intersection $\Phi_1(Z)\eta_2(Z)$, consisting of all maps tangent to Z at p_1 that intersect Z at p_2 . What we get are two parts:

- $\Lambda_Z(2p_1 + p_2)$, in which p_1 and p_2 do not lie on the same twig of degree zero;
- $\beta\Phi_1(Z) = \sigma_*\Phi_x(Z)$, corresponding to maps with both marked points on a degree 0 twig.

Set theoretically,

$$\eta_2(Z)\Phi_1(Z) = \Lambda_Z(2p_1 + p_2) \cup \sigma_*\Phi_x(Z).$$

The correct multiplicities are 1 and 2 respectively, giving:

$$\Lambda_Z(2p_1 + p_2) = \eta_2(Z)\Phi_1(Z) - 2\sigma_*\Phi_x(Z).$$

Step 2: We now intersect $\Lambda_Z(2p_1 + p_2)$ with $\eta_2(Z) + \psi_2 - \beta$. This imposes second order vanishing at the second marked point.

Again, this intersection gives one part with multiplicity 1, which is not on the boundary β , and another part with multiplicity 2 in β . The first is what we are looking for: $\Lambda_Z(2p_1 + 2p_2)$. To find the other part, remember we are already working inside $\Lambda_Z(2p_1 + p_2)$, so to lie in β means $p_2 \rightarrow p_1$. We thus obtain $\beta\Lambda_Z(3p_1) = \beta\Phi_1^{(3)}$ with multiplicity 2.

Note: To explain these multiplicities rigorously is a subtle business which we will not go into here. The intuition behind these 2's is that the boundary contribution can be carried by either of the classes we are intersecting, and hence shows up twice in the intersection. The reader interested in how to carry out these computation can consult [Gat05].

Finally, we have identified

$$\begin{aligned}\Lambda_Z(2p_1 + 2p_2) &= -2\sigma_*\Phi_x^{(3)}(Z) + (z\eta_2 + \psi_2 - \beta)\Lambda_Z(2p_1 + p_2) \\ &= -2\sigma_*\Phi_x^{(3)} + \Phi_1(Z)\Phi_2(Z) - 2\sigma_*\Phi_x(Z)(z\eta_2 + \psi_2 - \beta).\end{aligned}\tag{10}$$

Now all that stands in our way of calculating bitangents are substitutions and computations.

Before we move on, take a second to remember or derive the following easy facts, that we will use in the forthcoming computations:

- (a) $\iota\eta_x^2 = \iota^2\eta_x = 1$
- (b) $\iota\eta_x = \eta_x^2 + \iota^2$
- (c) $\psi_x = \iota - 2\eta_x$
- (d) $\eta_i^3 = 0$
- (e) $\psi_1 = \pi_2^*\psi_x + \beta$
- (f) $\beta\psi_i = 0$
- (g) $\beta^2 = -\beta\pi_i^*\psi_x$
- (h) $\beta\eta_1 = \beta\eta_2$
- (i) $\sigma_*\Phi_x(Z) = \beta\Phi_1(Z) = \beta\Phi_2(Z)$
- (j) $\eta_1\pi_2^*\alpha = \pi_2^*(\alpha\eta_x)$ for any class $\alpha \in A^*(\overline{M}_{0,1}(\mathbb{P}^2, 1))$.

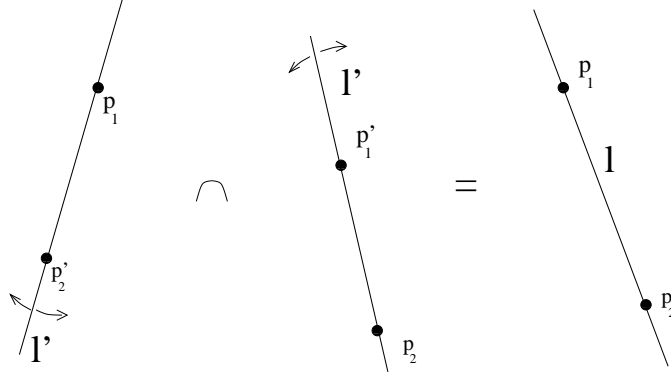


Figure 11: the intersection $\eta_1^2\eta_2^2$.

Let us expand the first term in (9), $\sigma_*\Phi_x^{(3)}$. We have already found $\Phi_x^{(3)}(Z) = (3d^2 - 6d)$ in $\overline{M}_{0,1}(\mathbb{P}^2, 1)$. Since pushing forward preserves dimension and σ_i is injective, thus sending the class of a point to the class of a point, then our first term becomes

$$\sigma_*\Phi_x^{(3)} = 3d^2 - 6d.$$

Now for the second and third terms.

$$\begin{aligned} \Phi_1(Z)\Phi_2(Z) &= d\eta_1(d\eta_1 + \psi_1 - \beta)d\eta_2(d\eta_2 + \psi_2 - \beta) \\ &= d^4(\eta_1^2\eta_2^2) + d^3(\eta_1^2\eta_2\psi_2 + \eta_1\eta_2^2\psi_1) + d^2(\eta_1\eta_2\psi_1\psi_2 + \eta_1^2\beta^2). \end{aligned}$$

$$\begin{aligned} \sigma_*\Phi_x(Z)(d\eta_2 + \psi_2 - \beta) &= \beta((d\eta_x)(d\eta_1 + \psi_1 - \beta))(d\eta_2 + \psi_2 - \beta) \\ &= -2d^2(\eta_1^2\beta^2) + d(\eta_1\beta^3). \end{aligned}$$

We must now compute each intersection in the above expressions⁵.

- $\eta_1^2\eta_2^2 = 1$, as seen from the figure 11 illustrating the fact that there is exactly one line passing through two prescribed points.
- $\eta_1^2\beta^2 \stackrel{(f)}{=} \eta_1^2\beta(-\pi_2^*\psi_x) \stackrel{(i)}{=} -\beta\pi_2^*(\eta_x^2\psi_x) \stackrel{(b+c)}{=} -\beta\pi_2^*(a) = -\beta[\text{fibre}]$.
But β is the image of a section, hence it intersects all fibres transversely. Thus, $\eta_1^2\beta^2 = -1$.
- $\eta_1\beta^3 \stackrel{(f)}{=} -\eta_1\beta^2\pi_2^*\psi_x \stackrel{(f)}{=} \eta_1\beta(\pi_2^*\psi_x)^2 = \eta_1\beta(\pi_2^*\psi_x^2) \stackrel{(i)}{=} \beta\pi_2^*(\psi_x^2\eta_x) \stackrel{(a+b+c)}{=} \beta\pi_2^*(-3) = -3$.

⁵The little numbers over the equal signs refer to the identities from page 35 that are used at each step.

- $\eta_1^2 \eta_2 \psi_2 \stackrel{(d)}{=} \eta_1^2 \eta_2 (\pi_1^* \psi_x + \beta) \stackrel{(i)}{=} \eta_1^2 \pi_1^* (\eta_x \psi_x) \stackrel{(a+b)}{=} \eta_1^2 \pi_1^* (\iota^2 - \eta_x^2)$.
 Notice first of all that $\eta_1^2 \pi_1^* \eta_x^2 = \eta_1^2 \eta_2^2 = 1$, as shown in figure 11.
 Next, we claim that $\eta_1^2 \pi_1^* \iota^2 = 0$. In fact η_1^2 is the class of all two-pointed lines passing through a fixed point, where the first mark is at the fixed point while the second mark is free to move.

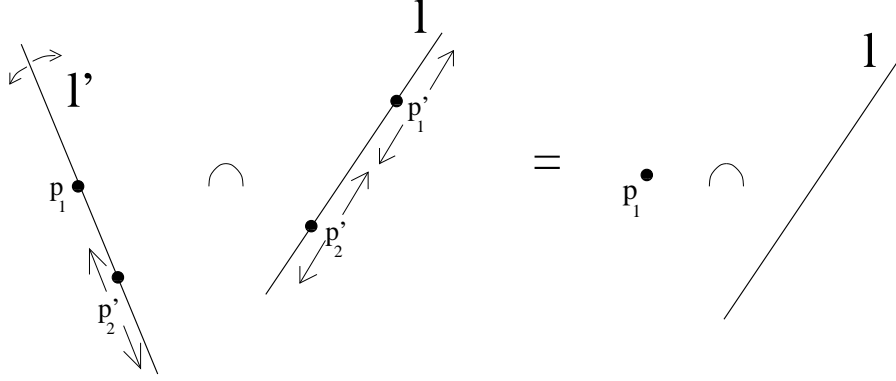


Figure 12: the intersection $\eta_1^2 \pi_1^*(\iota^2)$.

Intersecting with $\pi_1^*(\iota^2)$ means to require that our fixed point intersects our fixed line transversely, which is to say that they don't intersect at all. This is illustrated in figure 12.

Thus, $\eta_1^2 \eta_2 \psi_2 = -1$; by symmetry, we also have $\eta_2^2 \eta_1 \psi_1 = -1$.

- Finally,

$$\begin{aligned}
 \eta_1 \eta_2 \psi_1 \psi_2 &\stackrel{(d)}{=} \eta_1 (\pi_2^* \psi_x + \beta) \eta_2 (\pi_1^* \psi_x + \beta) \\
 &\stackrel{(i)}{=} (\pi_2^* (\psi_x \eta_x) + \beta \eta_1) (\pi_1^* (\psi_x \eta_x) + \beta \eta_2) \\
 &\stackrel{(a+b)}{=} \pi_1^* (\iota^2 - \eta_x^2) \pi_2^* (\iota^2 - \eta_x^2) + \beta \eta_2 \pi_1^* (\iota^2 - \eta_x^2) + \beta \eta_1 \pi_2^* (\iota^2 - \eta_x^2) + \beta^2 \eta_1^2 \\
 &= \pi_1^* (\iota^2) \pi_2^* (\iota^2) - \pi_1^* (\iota^2) \pi_2^* (\eta_x^2) - \pi_1^* (\eta_x^2) \pi_2^* (\iota^2) + \pi_1^* (\eta_x^2) \pi_2^* (\eta_x^2) \\
 &\quad + \beta \pi_1^* (\iota^2 \eta_x) + \beta \pi_2^* (\iota^2 \eta_x) + \eta_x^2 \beta^2.
 \end{aligned}$$

Via our previous calculations, the last three terms can be easily seen as 1, 1, and -1 respectively. So we must now find the first four terms above. We do so by recalling pictures.

$\pi_1^*(\iota^2)$ is the class of a fixed line and all ordered pairs of points on it.

To intersect two such classes is to require that our line is fixed as *two* transverse lines, which is impossible. Thus $\pi_1^*(\iota^2)\pi_2^*(\iota^2) = 0$.

$\pi_1^*(\eta_x^2) = \eta_2^2$, and notice that now symmetry implies that $\eta_2^2\pi_2^*(\iota^2) = \eta_1^2\pi_1^*(\iota^2) = 0$, as shown in figure 12. Similarly, $\pi_2^*(\eta_x^2)\pi_1^*(\iota^2) = 0$.

Lastly, let us intersect $\pi_2^*(\eta_x^2)$ with $\pi_1^*(\eta_x^2)$. But it's clear that $\pi_2^*(\eta_x^2)\pi_1^*(\eta_x^2) = \eta_1^2\eta_2^2$, and we have shown in figure 11 that this intersection is 1.

We have then found

$$\eta_1\eta_2\psi_1\psi_2 = 0 + 0 + 0 + 1 + 1 + 1 - 1 = 2.$$

Putting this all together, we have now calculated

$$\begin{aligned} \Lambda_Z(2p_1 + 2p_2) &= -2(3d^2 - 6d) + (d^4 - 2d^3 + d^2) - 2(-2d^2 - 3d) \\ &= d^4 - 2d^3 - 9d^2 + 18d \\ &= d(d - 2)(d - 3)(d + 3). \end{aligned}$$

After the dust has settled, we now know that a generic plane curve Z of degree d has

$$N_B(d) = \frac{1}{2}d^4 - d^3 - \frac{9}{2}d^2 + 9d$$

bitangents. Remember we are dividing by 2 because we do not care about the order of the marked points.

Notice that for $d = 2$ and $d = 3$ we get that there are no bitangents as should be the case. For $d = 4$ we find 28 bitangents, the first interesting result.

References

- [BT82] Raoul Bott and Loring Tu. *Differential Forms in Algebraic Topology*. Springer, 1982.
- [Ful98] William Fulton. *Intersection Theory*. Springer, second edition, 1998.
- [Gat05] Andreas Gathmann. The number of plane conics that are five-fold tangent to a given curve. *Compos. Math.*, 141(2):487–501, 2005.

- [GH94] Phillip Griffiths and Joseph Harris. *Principles of Algebraic Geometry*. Wiley, 1994.
- [Har77] Robin Hartshorne. *Algebraic Geometry*. Springer, 1977.
- [HM98] Joseph Harris and Ian Morrison. *Moduli of Curves*. Springer, 1998.
- [Knu83a] Finn F. Knudsen. The projectivity of the moduli space of stable curves. II. The stacks $M_{g,n}$. *Math. Scand.*, 52(2):161–199, 1983.
- [Knu83b] Finn F. Knudsen. The projectivity of the moduli space of stable curves. III. The line bundles on $M_{g,n}$, and a proof of the projectivity of $\overline{M}_{g,n}$ in characteristic 0. *Math. Scand.*, 52(2):200–212, 1983.
- [Koc99] Joachim Kock. Counting bitangents of a smooth plane curve via stable maps. talk at GAeL, 1999.
- [Koc01] Joachim Kock. Notes on psi classes. <http://mat.uab.es/~kock/GW.html>, 2001.
- [KV99] Joachim Kock and Israel Vainsencher. *A fórmula de Kontsevich para curvas racionais planas*. 22º Colóquio Brasileiro de Matemática. [22nd Brazilian Mathematics Colloquium]. Instituto de Matemática Pura e Aplicada (IMPA), Rio de Janeiro, 1999.
- [Sau89] D. J. Saunders. *The geometry of jet bundles*, volume 142 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1989.
- [Vak98] Ravi Vakil. A beginner’s guide to jet bundles from the point of view of algebraic geometry. Notes, 1998.

Department of Mathematics
Stanford University
 450 Serra Mall, Bldg. 380
 Stanford, CA 94305-2125

and
Department of Mathematics
University of Utah
 155 South 1400 East, Room 233
 Salt Lake City, UT 84112-0090

Email: ayala@math.Stanford.EDU
renzo@math.utah.edu