BRILL-NOETHER THEORY

TONY FENG

This article follows the paper of Griffiths and Harris, "On the variety of special linear systems on a general algebraic curve."

1. Introduction

Brill-Noether theory is concerned with a fundamental but (as it turns out) subtle question:

> What maps does a smooth projective curve C admit to projective space?

Let us be more precise. Of course, it is tautological that any projective curve can be embedded into *some* projective space. However, once we begin making demands on the embedding, we start to get some interesting answers.

For instance, can we make sure target projective space "small"? It is easy to show that not every curve can be embedded in $\mathbb{P}^2$. Conversely, *every* smooth projective curve can be embedded in $\mathbb{P}^3$.

Exercise 1.1. Prove this. [Hint: embed in some projective space. If there is a point not on the secant variety to the curve, then project from that point.]

What if we impose constraints on the *degree* of the image curve? Then the question becomes extremely hard; in fact it's still open (to my knowledge) precisely what degrees can occur for embeddings of smooth projective curves in $\mathbb{P}^3$.

The scope of Brill-Noether theory is the existence of *maps* (not necessarily embeddings) of curves to projective space. We will restrict our attention to a curve C of a given genus g, and ask how the answer changes as we vary the curve C in the moduli space $\mathcal{M}_g$. (For instance, some genus g curves admit very low degree maps even to $\mathbb{P}^1$, but most do not.)

Thus, a more precise question (which I will address in these talks) is: Given g and r, for what degrees d do *most* smooth projective curves of genus g admit a degree d map to $\mathbb{P}^r$? (Here by "most" curves I mean all curves corresponding to a (dense) open subset of $\mathcal{M}_g$.)

Since a map to projective space is the same data as a line bundle plus some linear system of global sections without a basepoint, one can rephrase this question in terms of the existence of line bundles with many global sections:

> For what g, d, r does a general smooth projective curve of genus g admit a line bundle with an $r+1$-dimensional space of global sections?

It's clear that if the answer is yes for some d, then it is yes for all larger d, as the degree of the map is simply the degree of the divisor. Therefore, the challenge is really to find the *lowest* d for which a general smooth projective curve of genus g admits a linear system of degree d and dimension $r+1$. (Such a linear system is called a g^r_d.)

Remark 1.2. It turns out that this is exactly the threshold when *all* curves possess such a line bundle.

Example 1.3. Let's consider the case $r = 1$. We seek linear systems with dimension 1 and degree d, or what is commonly known as a g^2_d. One says that a divisor in such a linear system "moves in a pencil." Note that the "trivial upper bound" coming from Riemann-Roch (which is not really trivial at all!) is $g+1$, since Riemann's inequality that for a divisor D of degree at least $g+1$, we have $h^0(D) \geq g+1+1-g = 2$.

$g = 0$. There is a degree 1 map $\mathbb{P}^1 \to \mathbb{P}^1$.

$g = 1$. There cannot be a degree 1 divisor on C. However, by Riemann-Roch any degree 2 divisor on C moves in a pencil: if $\deg D = 2$, then

$$h^0(D) = 2+1-1+h^0(K-D)-2.$$

You can see this directly if you know that every elliptic curve has a Weierstrass form

$$y^2 = x^3 + \dots$$

Then there is an obvious degree 2 map to $\mathbb{P}^1$, extending $(x, y) \mapsto x$. However, I want to highlight a different point of view which will be very useful for us.

What is the relation between a linear system and a map to $\mathbb{P}^1$? Usually we think of a map $C \xrightarrow{|D|} \mathbb{P}^1$ as being a function that associates to every point of C a point of $\mathbb{P}^1$. However, you can flip the script and thinking of it instead as a *family* of subvarieties of C parametrized by $\mathbb{P}^1$. What are the fibers of this family? Precisely the divisors of the linear system $|D|$!

Now, can we see any family of degree 2 divisors on C which move in a pencil, i.e. are parametrized by $\mathbb{P}^1$? Well, if you embed C in $\mathbb{P}^2$ as a cubic, you can consider the pencil of *lines* in $\mathbb{P}^2$ passing through a fixed $p \in C$. Any such line meets C in exactly two other points, and those collections of points precisely form the elements of a degree 2 linear system.

In particuliar, if we take $p = \infty$ then we are consider the pencil of lines through ∞. But in the Weierstrass model, those are precisely the vertical lines, which recovers the map

$(x, y) \mapsto x$ from before.

$g = 2$. By Riemann-Roch, every such curve is hypereliptic (admits a degree 2 map to $\mathbb{P}^1$).

$g = 3$. First off, can we take $d = 2$? This would imply that most (actually every) genus 3 curve is hyperelliptic, but it is easy to see that plane quartics have genus 3 and must be canonically embedded, and that these form an open subset of $\mathcal{M}_3$. Therefore, $d \geq 3$.

In fact, this observation gives a handle on general genus 3 curves as plane quartics under their canonical embeddings. We're looking for a degree 3 divisor. Again, we can take the linear system cut out by the system of lines passing through a fixed point $p \in C$.

$g = 4$. Now a general smooth projective curve of genus $g = 4$ is the complete intersection of a quadric and a cubic curve under its canonical embedding. Generally this quadric will be smooth ♠♠♠ TONY: [some initial calculations suggest it is always]. But a smooth quadric in $\mathbb{P}^3$ is isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$ (under the Segre embedding), hence has Picard group $\mathbb{Z} \times \mathbb{Z}$. Moreover, since the line bundle giving the Segre embedding is $\mathcal{O}(1, 1)$, the curve C represents the class $(3, 3)$ in $\text{Pic}(\mathbb{P}^1 \times \mathbb{P}^1)$. There is a visible linear system of degree 3 obtained from (say) the pullback of the line bundle $(0, 1)$, which is simply projection to the second coordinate.

This argument shows that $d \leq 3$. To see that $d > 2$, we could just apply Bertini's Theorem to argue that an intersection of a smooth cubic and smooth quadric surface in $\mathbb{P}^3$ is smooth and canonically embedded, and count dimensions to show that the space of such already fills out an open subset of $\mathcal{M}_4$. Instead, we will present a different argument that works in general.

If C has a g^1_d, then it can be presented as a degree d branched cover of $\mathbb{P}^1$. We simply compare the dimension of $\mathcal{M}_g$ (which is $3g - 3$) with the dimension N of all degree d branched covers of $\mathbb{P}^1$. Obviously, if every curve in $\mathcal{M}_g$ is going to have a g^1_d then we will need $3g - 3 \leq N$. In fact, we need something more. We have a map from branched covers to $\mathcal{M}_g$ which forgets the covering structure, but this is not finite. For instance, given any presentation of C as a branched cover of $\mathbb{P}^1$, we can apply an element of PSL_2 to obtain another branched cover. Therefore, the fibers of this map, when they are non-empty, *have dimension at least* 3, so we need $N \geq 3g$.

By Riemann-Hurwitz, a degree d branched cover of $\mathbb{P}^1$ has most

$$2g - 2 + 2d.$$

branch points. To specify the cover given the branch points, it suffices to specify a *monodromy representation* on the base, describing how the sheets are permuted under traversing a loop around a branch point. Thus, after specifying the branch points the cover is determined up to finite data.

Thus, we have a finite map associating a branched cover its branch points, which is finite. So we see that $N = 2g - 2 + 2d$, so what we need is

$$2g - 2 + 2d \geq 3g.$$

Thus we need $d \geq \frac{g}{2} + 1$. In particular, for $g = 4$ we need $d \geq 3$.

$g = 5$. By the dimension count we know that a general C is not trigonal, so $d > 3$. However, let's see this in another more geometric way.

It's worth pointing out a reformulation of the Riemann-Roch Theorem which is sometimes called "Geometric Riemann-Roch." Let's take a look at the formula

$$h^0(D) = d + 1 - g + h^0(K - D).$$

Now, $H^0(K)$ can be interpreted as the linear system of divisors on C which are cut out by the vector space of hyperplanes under the canonical embedding. Therefore, $H^0(K - D)$ can be interpreted as the vector space of hyperplanes in the canonical embedding *that contain the divisor* D. The "expected dimension" of this space is $g - d$, since each point of D imposes one condition on the space of hyperplanes. However, if the points are in special position. So let's write $h^0(K - D) = g - d + i$, where i is the number of "extra linear relations" among the points of D. Substituting this into Riemann-Roch, we get

$$h^0(D) = d + 1 - g + g - d + i = i + 1.$$

Put differently, *the linear system* $|D|$ *has dimension equal to the number of extra linear relations among* D *under the canonical embedding*, i.e. the difference between $d - 1$ and the dimension of the span of D.

Let's go back and apply this to some of the earlier examples. To produce a g^1_3 on a canonical curve of genus 3, i.e. a plane quartic, we need to find a divisor of degree 3 on a plane quartic with $i(D) = 1$. This says that normally three points on the curve would span the plane, but we want them to span a space of one dimension less: a line. So the points of D should be collinear, which is indeed what we found.

What about a canonical curve C of genus $g = 4$? Again we want a divisor of degree 3 which lies on a line. And indeed, C lies on a quadric surface as a curve of type $(3, 3)$, so a line from a ruling of the surface intersects our curve in three points.

Finally, we turn our attention to $g = 5$. Suppose that C had divisor D of degree 3 with $i(D) = 1$, i.e. there should be 2 relations among 3 points in $\mathbb{P}^4$, i.e. they have to be collinear. It is easy to show that in the canonical embedding, a general C will be the complete intersection of three quadrics. If there were collinear points on C, then the line spanned by them would intersect C, and a fortiori the quadrics, in at least 3 points. Thus it would have to be contained in the quadrics, but then it would be in the intersection of all of them, which is supposed to be just the curve C. So this can't happen.

However, to produce a divisor with $d = 4$ we just have to find a plane containing 4 points. Thinking to the previous example, we are inspired to look for a plane in a quadric. This can be done by the standard incidence correspondence

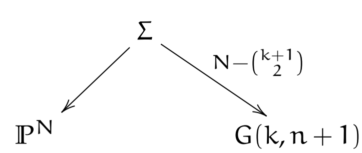

What we need is for $k(n + 1 - k) \geq \binom{k+1}{2}$, and indeed in the case of interest $n = 4, k = 3$, the left hand side is $3(2) = 6$ while the right hand side is 6. Thus, a general quadric threefold in $\mathbb{P}^4$ contains a finite number of planes. What is the intersection of such a plane with the curve? The same as the intersection of the plane with the other two quadrics, which cut out a degree 4 surface, so the answer is 4!

Remark 1.4. We could have seen this geometrically, without any calculation. Since C lies on two quadric hypersurfaces, it lies on the pencil of quadrics spanned by them, and some members (in fact, four) will be *singular*, i.e. cones over smooth quadrics in $\mathbb{P}^3$. But we know that a smooth quadric in $\mathbb{P}^3$ has two rulings of lines, so the cone over it has two rulings of 2-planes.

Let's compare the results with the lower bound obtained from the dimension argument.

g	d	$\lceil \frac{g}{2}+1 \rceil$
0	1	1
1	2	2
2	2	2
3	3	3
4	3	3
5	4	4

This reflects the principle that the answer is what will be predicted by dimension arguments. We'll shortly see how to get a general answer using dimension arguments.

2. The Brill-Noether Conjecture

Definition 2.1. We denote by $W^r_d(C)$ the subvariety of $J(C)$ parametrizing those line bundles whose space of global sections is at least $r+1$.

Definition 2.2. The *Brill-Noether* number is

$$\rho(g,r,d) := g-(r+1)(g-d+r) = h^0(K)-h^0(D)h^0(K-D).$$

Conjecture 2.3 (Brill-Noether)**.** *For all curves* C, $\dim W^r_d(C) \geq \rho(g,r,d)$. *Moreover, for a general curve* C *we have* $\dim W^r_d(C) = \rho(g,r,d)$.

Example 2.4. If $r=1$, then this is saying that pencils of degree d exist only if $g \geq 2(g-d+2)$, or $d \geq \frac{g}{2}+1$.

Where does this come from? Since we are interested in general curves, we may assume that C is canonically embedded. Then by the geometric form of Riemann-Roch, we are interested in when we can find points $p_1,\ldots,p_d \in C$ which are in a highly special position under the canonical embedding.

In concrete terms, we can pick a basis $\omega_1 = f_1\,dz,\ldots,\omega_g = f_g\,dz$ for $H^0(C,\Omega)$. Then under the canonical embedding the points have coordinates

$$M(p_1,\ldots,p_d) = \begin{pmatrix} f_1(p_1) & f_1(p_2) & \ldots & f_1(p_d) \\ f_2(p_1) & f_2(p_2) & \ldots & f_2(p_d) \\ \vdots & \vdots & \vdots & \vdots \\ f_g(p_1) & f_g(p_2) & \ldots & f_g(p_d) \end{pmatrix}$$

Geometric Riemann-Roch tells us that $\dim|D|$ is the difference between the dimension of the span of d points in general position and the dimension of the span of the d points comprising D, which is $d - \operatorname{rank} M(p_1,\ldots,p_d)$. Letting $p_1,\ldots,p_d$, we can think of this as a map

$$C^{(d)} \to \mathrm{Mat}(g \times d)$$

where $C^{(d)}$ is the dth symmetric power of C, and we are interested in the locus of d-tuples of points whose images are matrices of rank at most $d-r$. Let's call this locus C^r_d. It maps to the Jacobian, contracting tuples of points that represent equivalent linear series, and hence is a $\mathbb{P}^r$-fibration over W^r_d. Since codimension doesn't increase under pullback, the codimension of this locus is at most the codimension of the space of rank at most $d-r$ in the space of all $g \times d$ matrices. This is "just" a linear algebra problem!

View such a matrix as a map $\mathbb{C}^g \to \mathbb{C}^d$. To such a map we can associate its image, which is generically of dimension $d-r$, so the space of images is of dimension $\dim G(d-r,d) = r(d-r)$. Fixing the image, the space of homomorphisms is $g(d-r)$. So the dimension is $(g+r)(d-r)$. Therefore, the codimension is $\boxed{rg - rd + r^2.}$

We have found that the codimension of C^r_d in $C^{(d)}$ to be at most $r(g-d+r)$. Therefore, the dimension of C^r_d is at least $d - r(g-d+r)$, and $\dim W^r_d$ should then be at least $d - r(g-d+r) - r = g - (r+1)(g-d+r)$, which is precisely $\rho(g,d,r)$!

This is the original argument of Brill and Noether. There is one problem here, which is that we never precluded the possibility that C^r_d was actually *empty*. (The statement "codimension doesn't increase" has this caveat built into it; the empty set has any dimension!)

Remark 2.5. Another subtlety is that one doesn't know a priori that $\dim C^r_d = \dim W^r_d + r$. It *could* be the case that every degree d divisor with $r+1$ sections *automatically* has $r+2$ sections.

Other than that issue, this argument is surprisingly complete. It is not hard to globalize it to show that $\dim W^r_d \geq \rho$ as long as W^r_d is non-empty. Let me mention an elegant alternate way of phrasing this argument. Let $\mathcal{L}$ be a line bundle and $p_1, \ldots, p_n$ points on C. Then we have an exact sequence of sheaves

$$0 \to \mathcal{L} \to \mathcal{L}(p_1 + \ldots + p_n) \to \bigoplus_i \mathcal{O}_{p_i} \to 0.$$

Therefore, we get a left-exact sequence on global sections

$$0 \to H^0(C,\mathcal{L}) \to H^0(C,\mathcal{L}(p_1+\ldots+p_n)) \to \bigoplus_i \mathbb{C}$$

Therefore, we can interpret

$$H^0(C,\mathcal{L}) = \ker\left(H^0(C,\mathcal{L}(p_1+\ldots+p_n)) \to \bigoplus_i \mathbb{C}\right).$$

If n is sufficiently large, then we know $h^0(C,\mathcal{L}(p_1+\ldots+p_n))$ precisely from Riemann-Roch, and this gives another realization of $H^0(C,\mathcal{L})$ as the locus of linear maps with "small" rank.

3. Outline of the proof

Kleiman and Laksov established the inequality.

Theorem 3.1 (Kleiman-Laksov). *For all curves C, if $\rho(g,r,d) \geq 0$ then $W^r_d(C)$ is non-empty.*

As we mentioned above, this immediately implies that $\dim W^r_d(C) \geq \rho(g,r,d)$.

Remark 3.2. When $\rho(g,d,r) = -1$, we expect the curves possessing a g^r_d to have dimension $\dim \mathcal{M}_g - 1$, i.e. to form a *divisor* in the moduli space. This idea is exploited by Harris and Eisenbud to prove that $\mathcal{M}_g$ is of general type for large g.

The other inequality was (historically) more challenging. This may be surprisingly, because standard upper-semicontinuity results imply that it is enough to exhibit a *single* C for which $\dim W^r_d(C) \leq \rho(g,d,r)$. Let us explain.

It is famous fact that the smooth projective curves of genus g assemble into a *moduli space* $\mathcal{M}_g$ of dimension $3g-3$. Every such curve C has a *Jacobian variety* $J(C)$ parametrizing line bundles on C of degree d, and it is a theorem that these fit together into a relative Jacobian variety $\mathcal{J}_g$ over $\mathcal{M}_g$. (You might complain that this is better thought of as a torsor for the Jacobian rather than the Jacobian itself; this is correct but irrelevant for our purposes since we won't use the group structure in any way.) The fiber of $\mathcal{J}_g$ over $[C] \in \mathcal{M}_g$ is precisely $J(C)$.

These spaces $W^r_d(C)$ assemble into a relative space $\mathcal{W}^r_d \to \mathcal{M}_g$, and the work of Kleiman-Laksov/Brill-Noether shows that if $\rho \geq 0$ then this map is dominant, and $\dim \mathcal{W}^r_d \geq \dim \mathcal{M}_g + \rho(g, r, d)$. By upper-semicontinuity of fiber dimensions for proper maps *on the base*, in order to prove the equality *it suffices to exhibit a single curve* C *with* $\dim W^r_d(C) = \rho$. We call such curves *Brill-Noether general*.

This is not as easy as it sounds. To this day, nobody has ever written down explicitly such a smooth projective curve. The curves that are easy to write down and analyze - hyperelliptic (or more generally low gonality), complete intersections - are all Brill-Noether special.

Remark 3.3. A few years after the Griffiths-Harris proof, Lazarsfeld produced a construction of smooth Brill-Noether general curves. However, this construction takes place within general K3 surfaces, which we don't really have any explicit way of writing down either.

One feature that distinguishes this problem from classical results on linear series, such as Riemann-Roch or Clifford's theorem, is that it *only* holds for a general curve. A key trick in such situations is to introduce a *variational* element into the problem. More precisely, the strategy is:

(1) Write down a *singular* curve which is "Brill-Noether general" in some appropriate sense.
(2) Consider a family of smooth curves degenerating to this singular curve, and show that the property of Brill-Noether generality spreads out to the family.

We remark that the second step is not trivial. It *would* be trivial if one had a proper compactification of $\mathcal{W}^r_d$ over $\overline{\mathcal{M}_g}$. However, the Jacobian of non-compact curves need not be compact, and for non-proper maps upper-semicontinuity of fiber dimension on the base does *not* hold (example: an open embedding).

Eisenbud and Harris discovered that there is a nice notion of linear systems on certain kinds singular curves which are called *compact type* (defined by the property that their Jacobian is compact, or more concretely that the dual graph is a *tree*). They used this theory to give somewhat slicker proofs of the theorem, while the original approach of Griffiths and Harris (which we follow) relies on ad hoc synthetic constructions. Recently, Osserman and tropical geometers have made some further progress on the front of defining reasonable notions of linear series on singular curves.

We will focus on the argument for the first step, which is based on ideas of Castelnuovo and Severi, and revived by Kleiman.

Castelnuovo. Castelnuovo was interested in the enumerative problem of counting $\deg W^r_d(C)$ for a general curve of C when $\rho(g, d, r) = 0$. In other words, when $\rho(g, d, r)$ we expect a general smooth projective curve of genus g to have a *finite number* of g^r_ds, and Castelnuovo wanted to count this number. To do so, he introduced the following geometric construction. Degenerate C to a g-nodal rational curve C_0, which we shall call a *Castelnuovo* *canonical curve.*

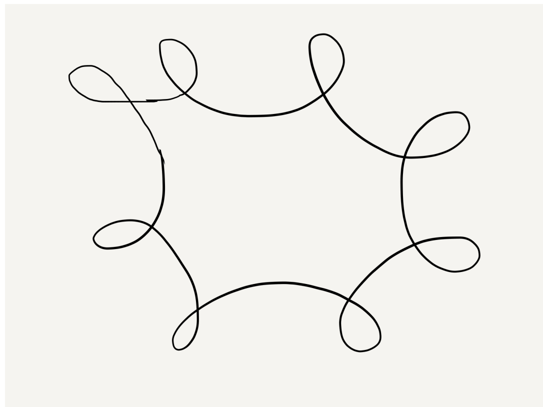

Castelnuovo's idea was that if d, g, r are such that $\rho(g, d, r) = \dim W^r_d(C_0) = 0$, then by the principle of conservation of number its degree will be the general answer. This fundamental principle of degeneration was the basis for Schubert's enumerative calculus, and a favorite trick of the 19th century algebraic geometers. The second talk will focus on the technique of degeneration.

Now, the normalization of C_0 is $\mathbb{P}^1$, so we can view a linear system on C_0 as a linear system on $\mathbb{P}^1$ factoring through the quotient map gluing g pairs of points $p_1, q_1, \ldots, p_g, q_g$.

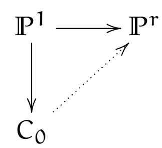

A linear system of degree d on $\mathbb{P}^1$ is simply a linear system of hyperplanes in the dth Veronese embedding $\mathbb{P}^1 \hookrightarrow \mathbb{P}^d$. To say that a linear system factors through the quotient map is to say that it takes the same values at p_1 and q_1, p_2 and q_2, etc. In other words, p_i and q_i *should always appear in the same hyperplane section*.

An r-dimensional linear system of hyperplanes is determined by its base locus, which has dimension $d - r - 1$. What is the condition that a hyperplane in this linear system passes through p_i if and only if it passes through q_i? This is equivalent to the base locus intersecting the *secant* between p_i and q_i, for then any hyperplane containing p_i contains two points on the secant, hence the whole secant.

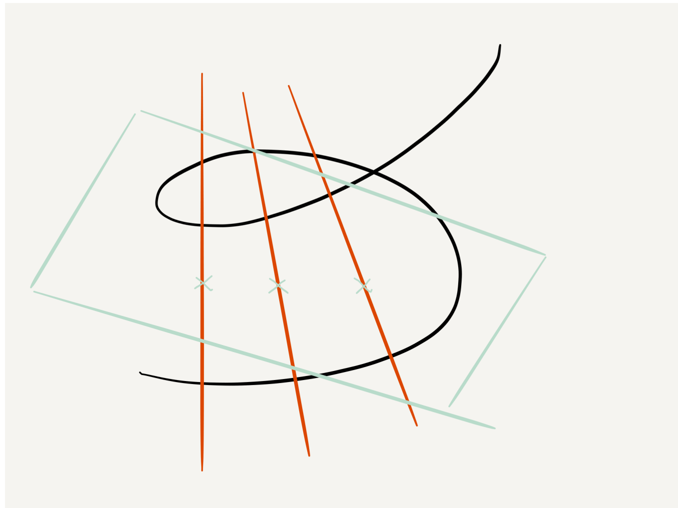

Therefore, the dimension of the space of linear systems of degree d and dimension r is the dimension of the space of $d-r-1$ planes intersecting every secant $\overline{p_i q_i}$. This is an intersection problem in $\mathbb{G}(d-r-1, d)$, which has dimension $(d-r)(r+1)$. The condition that a $d-r-1$ plane passes through a line is of codimension r, so the resulting space has expected dimension

$$(d-r)(r+1) - gr = g - (r+1)(g-d+r) = \rho(g, d, r).$$

Severi. Castelnuovo assumed the Brill-Noether Conjecture and used this idea to count (correctly!) the number of such $d-r-1$ planes when $\rho(g, d, r) = 0$. It wasn't until Severi (in the 1920s) that it pointed out, at least in public writing, that the "Brill-Noether theorem" had never been rigorously proven, and he suggested that Castelnuovo's construction could be used to prove it. Specifically, he suggested a two-step argument.

(1) Show that the expected dimension is correct for Castelnuovo's canonical curves.
(2) Show that this result would "spread out" to imply the theorem.

Severi could not rigorously establish either part. Eventually, Kleiman and Laksov, with the apparatus of modern algebraic geometry, resolved the second step, thus reducing the theorem to:

Conjecture 3.4 (Severi-Kleiman-Laksov)**.** *For g general secants to a rational normal curve in $\mathbb{P}^d$, the space of $d-r-1$ planes intersecting each secant is the expected dimension $\rho(g, d, r)$.*

By Schubert calculus one knows of course that the dimension of $d-r-1$ planes intersecting g *general lines* is $\rho(g, d, r)$. The challenge is to prove that this is *still* true if the lines are all secants to a rational normal curve, which is certainly not general.

Griffiths and Harris employ a degeneration argument, which I'll outline now and work through in detail next time. The idea is to specialize the secants. More specifically, we allow the point p_2 to limit to p_1, and then the point q_2 to limit to q_1.

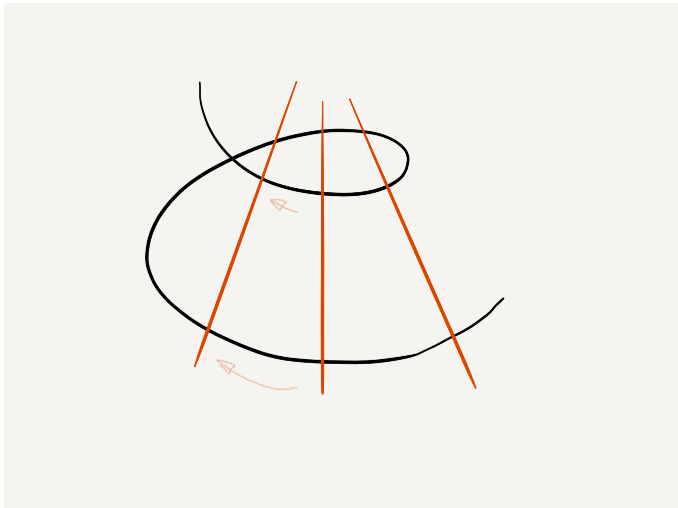

If one can prove that the intersection is transverse in this special situation, then it must be transverse in the general situation. The advantage of this degeneration is that it is easy to reduce the statement to one of smaller complexity, thus opening up an inductive argument. Indeed, as you can see we have reduced from considering $g+1$ secants to g secants. But there are some additional conditions that one picks up in the limit process: obviously not every plane passing through g secants can arise as a limit of planes passing through $g+1$ secants. The challenge is to identify these extra conditions and show that they are of the right codimension. This will be explored in detail in the second talk.