Geometry of Algebraic Curves

Lectures delivered by Joe Harris
Notes by Akhil Mathew

Fall 2011, Harvard

Contents

Introduction

Joe Harris taught a course (Math 287y) on the geometry of algebraic curves at Harvard in Fall 2011. These are my "live-TEXed" notes from the course.

Conventions are as follows: Each lecture gets its own "chapter," and appears in the table of contents with the date. Some lectures are marked "section," which means that they were taken at a recitation session. The recitation sessions were taught by Anand Deopurkar.

Of course, these notes are not a faithful representation of the course, either in the mathematics itself or in the quotes, jokes, and philosophical musings; in particular, the errors are my fault. By the same token, any virtues in the notes are to be credited to the lecturer and not the scribe.

Please email corrections to `amathew@college.harvard.edu`.

Lecture 1
9/2

§1 Introduction

The text for this course is volume 1 of Arborello-Cornalba-Griffiths-Harris, which is even more expensive nowadays.

We will be covering a subset of the book, and probably adding some additional topics, but this will be the basic source for most of the stuff we do. There will be weekly homeworks, and that will determine the grade for those of you taking the course for a grade. There will be no final.

There will be a weekly section with Anand Deopurkar.

The course doesn't meet on Mondays because of the "Basic Notions" seminar. On those weeks that there won't be seminars, we might meet then. Every third or fourth week, the lecture will be cut short on Wednesdays. However, the course will basically meet on a three-hour-per-week basis.

§2 Topics

Here is what we are going to talk about. We are going to talk about compact Riemann surfaces, and a compact Riemann surface is the same thing as a smooth projective algebraic curve (over $\mathbb{C}$). That in turn is really the same thing as a smooth projective curve over *any* algebraically closed field of characteristic zero. By abuse of notation, we will use $\mathbb{C}$ to denote any such field as well.

The fact that these are the same thing—that is, that a compact Riemann surface is an algebraic curve—is nontrivial. It requires work to show that such an object even admits a nontrivial meromorphic function. Note:

1.1 Proposition. *There are compact complex manifolds of* $\dim \geq 2$ *that do not admit nonconstant meromorphic functions.*

The miracle in dimension one is that there are nonconstant meromorphic functions, and enough to embed the manifold in projective space.

There are a number of beautiful topics that we are not going to cover. We will not talk about *singular* algebraic curves, in general. We will encounter them (e.g. when we consider maps of curves to projective space, the image might be singular), but they will not be the focus of study. When we have a singular curve C in projective space, we will treat C its *normalization*, i.e. as an image of a smooth projective curves. We also will not talk about *open* (i.e. noncompact) Riemann surfaces. Another even-more large-scale and interesting topic is the theory of *families* of curves, and how the isomorphism class of a curve changes as we vary the coefficients of the defining equations (e.g. moduli of curves). Finally, we are not going to talk about curves over fields that are not algebraically closed. There has been a huge amount of work on algebraic curves over $\mathbb{R}$, but we won't discuss them.

§3 Basics

Today, we shall set the notation and conventions. Algebraic curves is one of the oldest subjects in modern mathematics, as it was one of the first things people did once they learned about polynomials. It has developed over time a multiplicity of language and symbols, and we will run through it.

Let X be a smooth projective algebraic curve over $\mathbb{C}$. There are many ways of defining the **genus** of X, e.g. via the Hilbert polynomial, the Euler characteristic (via coherent cohomology), and so on. We are just going to take the naive point of view.

1.2 Definition. The **genus** of X is the topological genus (as a surface).

We can also use:

1. $g(X) = 1 - \chi(\mathcal{O}_X)$.
2. $1 - \frac{1}{2}\chi_{top}(X)$.
3. $\frac{1}{2}\deg K_X + 1$ (for K_X the canonical divisor, see below).

Given a geometric object, one wishes to define the functions on it. On a compact Riemann surface, there are *no* nonconstant regular functions, by the maximal principle: every holomorphic function on X is constant. We need to allow poles, and to keep track of them we will need to introduce the language of divisors.

1.3 Definition. A **divisor** D on X is a formal finite linear combination of points $\sum n_i p_i, p_i \in X$ on the curve. We say that D is **effective** if all $n_i \geq 0$ (in which case we write $D \geq 0$), and we say that $\deg D = \sum n_i$ is the **degree** of D.

As a reality check, one should see that the family of effective divisors of a given degree d should be the dth symmetric power of the curve with itself, C_d; this is $C \times \cdots \times C$ (d times) modulo the symmetric group S_d.

The point of this is, if we have a meromorphic function $f : X \to \mathbb{P}^1$, we can associate to it a divisor measuring its zeros and poles.

1.4 Definition. Given $f : X \to \mathbb{P}^1$, we say that the **divisor** of X is the sum $\sum_{p \in X} \operatorname{ord}_p(f)p$. The positive terms come from the zeros, while the negative terms come from the poles.

The problem is to deal with interesting classes of functions. Holomorphic functions do it, while there are too many meromorphic functions to work with them all at once. Instead, we do the following:

1.5 Definition. Given a divisor $D = \sum n_p p$, we can look at functions which may have poles at each p, but with orders bounded by n_p. Namely, we look at the space $\mathcal{L}(D)$ of rational functions $f : X \to \mathbb{P}^1$ such that $\operatorname{ord}_p(f) \geq -n_p$ for all $p \in X$. This is equivalently the space of f such that $\operatorname{div}(f) + D \geq 0$.

We note (without proof) that $\mathcal{L}(D)$ is finite-dimensional. We write:

1.6 Definition.

1. $\ell(D) = \dim \mathcal{L}(D)$.

2. $r(D) = \ell(D) - 1$.

In the literature, both notations ℓ, r are used.

The **basic problem** is this: given D, find explicitly these vector spaces $\mathcal{L}(D)$, and in particular the dimension $\ell(D)$ and the number $r(D)$. This is a *completely solved problem,* and not just by general theorems like Riemann-Roch. If one is given an algebraic curves as a smooth projective curve (given by explicit equations), and an explicit divisor, there is an *algorithm* to determine the space $\mathcal{L}(D)$. We'll do that in a week or two.

The one thing to observe is that there's a certain redundancy here, in this problem. This is for the following reason: if $\mathcal{L}(D)$ is known, and E is another divisor that differs from D by $\mathrm{div}(f)$ for some global meromorphic $f : X \to \mathbb{P}^1$ (nonconstant), then we can determine $\mathcal{L}(E)$. Namely, there is an elementary isomorphism

$$\mathcal{L}(D) \simeq \mathcal{L}(E),$$

given by multiplying by the rational function f. So, in some sense, asking to describe $\mathcal{L}(D)$ is the *same* as asking to describe $\mathcal{L}(E)$. To solve this problem, we need only study divisors modulo this equivalence relation.

1.7 Definition. We say that D, E are **linearly equivalent** when there is a global meromorphic function f such that $D - E = \mathrm{div}(f)$. We write $D \sim E$.

In some sense, the fundamental object is the space of divisors *modulo* linear equivalence. Note that the *degree* of a global rational function is zero, so the degree is defined modulo linear equivalence. (This states that the number of poles is the same as the number of zeros, for a global meromorphic functions.)

We are going to realize the space of divisors modulo linear equivalence as a space. For now:

1.8 Definition. We call $\mathrm{Pic}^d(X)$ the space of divisors on X modulo linear equivalence.

So far, we're just talking about divisors in general on X. There is a particular one we should keep in mind, the *canonical* divisor.

1.9 Definition. If ω is a meromorphic 1-form (i.e., something that locally looks like $f(z)dz$ for a meromorphic function f), we can define the **order** at a point $p \in X$ (via the order of the coefficient function). In particular, we can define the **divisor** $\mathrm{div}(\omega)$ of a meromorphic 1-form ω.

Note that the ratio of two meromorphic 1-forms ω_1, ω_2 is a global meromorphic (or rational) function ω_2/ω_1. In particular, $\mathrm{div}(\omega_1) = \mathrm{div}(\omega_2)$.

The **canonical class** K_X of X is the class of the divisor of ω for any meromorphic 1-form ω.

We're now going to turn around and say everything again, in a more modern language.

There are a lot of circumstances in which we want to forget about the divisor D, and think only of linear equivalence. We would like a terminology that would let us only specify the divisor mod linear equivalence. This will be the language of *line bundles.*

1.10 Definition. Suppose $D = \sum n_p p$ is a divisor on X. Let $\mathcal{O}_X$ be the sheaf of regular functions on X; similarly, we define $\mathcal{O}_X(D)$ to be the sheaf of functions with zeros and poles prescribed by this divisor D. In other words, the sections of $\mathcal{O}_X(D)$ over an open subset U are the meromorphic functions $f : U \to \mathbb{P}^1$ such that $\mathrm{ord}_p(f) \geq n_p$ for $p \in U$.

This is a local version of the space $\mathcal{L}(D)$ (i.e. $\mathcal{L}(D) = \Gamma(X, \mathcal{O}_X(D))$, and is much larger.

The point is that this sheaf $\mathcal{O}_X(D)$ is *locally free* of rank one. In other words, it is a holomorphic line bundle. We will basically identify holomorphic line bundles with locally free sheaves of $\mathcal{O}_X$-modules of rank one; this is standard in algebraic geometry.

Remark. If $D \sim E$, then $\mathcal{O}_X(D) \simeq \mathcal{O}_X(E)$ as line bundles. So, in some sense, we can realize the space of divisors modulo linear equivalence as the space of line bundles on X. That's how $\mathrm{Pic}^d(X)$ is typically defined.

There's one more thing we should define. There is a canonical divisor class, and the associated line bundle can be easily defined. Namely, we just have to take the holomorphic cotangent bundle T_X^*; this corresponds to the divisor associated to any global meromorphic 1-form.

We said at the outset that compact Riemann surfaces correspond to smooth projective curves. How do we go from one to another? How do we describe maps from X to projective space $\mathbb{P}^n$? This is where the notion of divisors and line bundles plays an essential role. We could write all this in the classical language of divisors, but we'll use the modern language of line bundles; you can think in the former way if you wish.

Let $\mathcal{L}$ be a line bundle on X. Suppose $\sigma_0, \dots, \sigma_r$ are global sections of $\mathcal{L}$, say linearly independent. (So if $\mathcal{L} = \mathcal{O}_X(D)$, then each σ_i corresponds to an element of $\mathcal{L}(D)$, i.e. a meromorphic function on X with appropriate zeros and poles.) Suppose that they have *no common zeros.* In that case, there is induced a map

$$X \to \mathbb{P}^r, \quad x \mapsto [\sigma_0(x), \dots, \sigma_r(x)].$$

Here the $\sigma_i(x)$ sure aren't numbers, but we can still make sense of this. Namely, the σ_i are not numbers, but they are elements of the fiber of $\mathcal{L}$ over p. Given an $r+1$-tuple of elements of a one-dimensional $\mathbb{C}$-vector space, we get a *uniquely determined* element of $\mathbb{P}^r$. Another way to see it is that a line bundle is locally trivial, and we can use a local trivialization to think of the $\sigma_i(x)$ as functions so that the $\sigma_i(x)$ are actual numbers. If you chose a different trivialization, you would get a different vector $[\sigma_0(x), \dots, \sigma_r(x)]$, but it would be the same *up to scalars.* In fact, if we changed the σ_i around, then we would just change the embedding by some automorphism group of $\mathbb{P}^n$.

So,

> Up to automorphisms of $\mathbb{P}^n$, the map $X \to \mathbb{P}^r$ is uniquely determined by the subspace of $H^0(\mathcal{L})$ spanned by the $\{\sigma_i\}$.

We thus get:

1.11 Proposition. *There is a correspondence between pairs $(\mathcal{L}, V)$ where $\mathcal{L}$ is a line bundle on X of degree d and $V \subset H^0(\mathcal{L})$ is an $r+1$-dimensional space of sections with* *no common zeros, and the space of nondegenerate maps of degree d, $X \to \mathbb{P}^r$ (up to automorphisms of $\mathbb{P}^r$).*

So if we're looking for maps to projective space, it's the same as looking for line bundles and sections. Here the word "nondegenerate" means that the image is not contained in a hyperplane. If the image is contained in a hyperplane, then the sections $\sigma_1, \ldots, \sigma_r$ used to define the map would satisfy a nontrivial linear relation.

1.12 Definition. A **linear series** of degree d and dimension r on X is a pair $(\mathcal{L}, V)$:

1. $\mathcal{L}$ is a holomorphic line bundle on X, of degree d (i.e. the associated divisor has degree d).
2. V is an $r+1$-dimensional space of sections, contained in $H^0(\mathcal{L})$.

We are no longer requiring that there be no common zeros. This is to make the space compact.

This will be denoted g^r_d: here, again, the d refers to the degree, r the dimension. Here "g" comes from the old word for a divisor, a "group."

A g^r_d can be thought of as a family of effective divisors. For each $\sigma \in V \subset H^0(\mathcal{L})$, we can associate the divisor of zeros $\operatorname{div}(\sigma)$. This means that $\mathbb{P}(V)$ corresponds to a family of effective divisors on X, parametrized by the projective space $\mathbb{P}(V)$. To recap, each nonzero section has a divisor (or divisor of zeros), and the section is defined up to rescaling by its divisor (because there are no nonconstant holomorphic functions on X). When you think of a linear series in this form, it is often denoted $\mathscr{D}$.

We can give a more intrinsic description of the map associated to a linear system. Given a linear system $(\mathcal{L}, V)$ on X without common zeros, we can describe the associated map $X \to \mathbb{P}^r$ as the map specifically sending X to the projectivization $\mathbb{P}(V^*)$, so that a point $p \in X$ is sent to the hyperplane of sections in V vanishing on p (and this hyperplane belongs to $\mathbb{P}(V^*)$).

Lastly (and this is really crucial), when the genus of X is ≥ 1, we have enough holomorphic differentials on X to get a map to projective space.

1.13 Proposition. *The dimension of the space $\mathcal{L}(K)$ of holomorphic 1-forms is exactly the genus g. Moreover, there are no common zeros of the canonical line bundle $\mathcal{L}(K)$, so there is a canonical map $X \to \mathbb{P}^{g-1}$ associated to the linear series of all 1-forms.*

We've left out a lot of details (and even a lot of definitions), but now we want to *do something.* Let's give a false proof of Riemann-Roch in the next five minutes.

Given a divisor D, we're (as before) interested in the space $\mathcal{L}(D)$ of rational functions with the desired poles. Suppose for simplicity D is a sum of distinct points, $D = p_1 + p_2 + \cdots + p_d$. The vector space $\mathcal{L}(D)$ consists of meromorphic functions that have at most a simple pole at the points p_i but are otherwise regular. How might you define such a function?

Choose a local coordinate z_i around p_i; given a function $f \in \mathcal{L}(D)$, we can write it near p_i as $\frac{a_i}{z_i} + f_0$ for some holomorphic f_0 (defined in a neighborhood of p_i), because f is only allowed to have simple poles. This *polar part* $\frac{a_i}{z_i}$ says a lot about f. In fact, f is determined *up to addition of scalars* by specifying these polar parts $\{a_i\}$, again because there are no nonconstant holomorphic functions on X. In other words, there is a natural map

$$\mathcal{L}(D) \to \mathbb{C}^d,$$

whose kernel consists of the constant functions. We get in particular,

$$\ell(D) \leq 1 + d.$$

If you want to describe $\mathcal{L}(D)$ (and in particular its dimension), you want to know the image. This raises the question:

> Given $a_1, \dots, a_d$, when is there a global meromorphic $f : X \to \mathbb{P}^1$ with polar part $\frac{a_i}{z_i}$ at p_i, and holomorphic elsewhere?

In other words, we need to find *constraints* on the $\{a_i\}$ for them to form a family of polar parts of a function in $\mathcal{L}(D)$. Here's the point: if $f \in \mathcal{L}(D)$, and ω is a *holomorphic 1-form*, we can consider the meromorphic differential $f\omega$. This has potentially simple poles at the $\{p_i\}$, but is holomorphic elsewhere. In particular, the sum of the residues is zero.

(Recall that the sum of the residues of a meromorphic differential on a compact Riemann surface is zero.)

So, if $\omega = g_i(z_i)dz_i$ locally, near p_i (using the local coordinate z_i around p_i), then the residue of $f\omega$ at p_i is $a_i g_i(p_i)$. It follows that if the $\{a_i\}$ arise as a system of polar parts, then we want $\sum a_i g_i(p_i) = 0$. Thus the image of the map $\mathcal{L}(D) \to \mathbb{C}^d$ is contained in the *orthogonal complement* of the space of holomorphic differentials (where each ω maps to $(g_i(p_i)) \in \mathbb{C}^d$ as before). We get a total of g linear conditions on the $\{a_i\}$, over a family of holomorphic differentials, suggesting that

$$\ell(D) \leq 1 + d - g$$

But this is **false**. It's possible that a differential might not give a serious relation on the $\{a_i\}$. For instance, a differential at the points $\{a_i\}$. The correct statement is

$$\ell(D) \leq 1 + d - (g - \ell(K - D)), \tag{1}$$

because the dimension of the image of $\mathcal{L}(D) \to \mathbb{C}^d$ is contained in the orthogonal complement of the image of the 1-forms in $\mathbb{C}^d$.

However, a 1-form is the dual of a vector field. So the degree of a 1-form is just minus the degree of the corresponding vector field, and that is the topological Euler characteristic. In particular, the degree of a 1-form is the opposite of the topological Euler characteristic, i.e. $2g - 2$.

Now we want to apply (1) to $K - D$, and we get

$$\ell(K - D) \leq 1 + (2g - 2 - d) - g + \ell(D),$$

where we have used the degree of K to get the degree of $K - D$ as $2g - 2 - d$. Now we add this to (1). We get $\ell(D) + \ell(K - D) \leq \ell(K - D) + \ell(D)$, and as a result we must have equalities in both inequalities. We "get":

1.14 Theorem (Riemann-Roch)**.** *For any divisor D of degree d, we have*

$$\ell(D) = d - g + 1 + \ell(K - D).$$

But, we've been cheating. The restriction in the argument to divisors of the form $p_1 + \cdots + p_d$ was ok, but there is a serious gap in this argument.

The point, however, is that everything derives from the condition that the sum of the residues of a meromorphic differential be zero.

§4 Homework

Here is the homework assignment: ACGH, chapter 1, "batch A," problems 1-5. This will be due next Friday. When we talk about a "curve," we mean the normalization of the compactification.

Lecture 2
9/7

We start with some material intended to help out with the problem set.

§1 Riemann surfaces associated to a polynomial

Consider a polynomial $f(x, y) \in \mathbb{C}[x, y]$. We want to think of this in the form $a_d(x)y^d + a_{d-1}(x)y^{d-1} + \cdots + a_1(x)y + a_0(x)$. Consider the equation that this should be set to zero. As we vary x, we want to think of there being d solutions of y. So y is a function implicitly defined by x, with d possible values; it varies holomorphically. A large part of the impetus for Riemann surfaces was to provide a proper foundation for such multi-valued functions. (They did not develop abstractly as *compact complex manifolds of degree one.* They came up as branched covers of $\mathbb{P}^1$, that occurred via multi-valued functions in this way.)

We're going to discuss this complex-analytically, to start with. Consider $X = \{(x, y) : f(x, y) = 0\} \subset \mathbb{C}^2$, the "complex plane" (this is very ambiguous terminology). Let X^0 be the smooth locus, so X^0 is a smooth submanifold of $\mathbb{C}^2$. We want to think of this as a branched cover of $\mathbb{P}^1$.

We have a projection map $\pi : X \to \mathbb{C}$ that sends a point to its x-coordinate. For $p \in X$, there are three possibilities:

1. π is smooth at p. That is, $\frac{\partial}{\partial y} f(p) \neq 0$, so X is smooth at p and π is a local isomorphism (in the complex analytic or étale topologies).

2. $\frac{\partial}{\partial y} f(p) = 0$, but $\frac{\partial}{\partial x} f(p) = 0$. Here X is smooth at p, but the projection map from X is not locally one-to-one. So π is locally m-to-one, and in terms of local coordinates, it looks like $z \mapsto z^m$. This is a branched cover of the disk.

   In this case, we say that p is a *branch point* of the map, and m the *ramification.*

3. Both derivatives vanish, in which case we have a singular point of the curve (in $X - X^0$). There is an algebraic answer, but here is the complex analytic answer.

However, it *doesn't matter.* No matter how bad the singularity is, we can still find a suitable neighborhood of the image point, such that over the *punctured disk*, the associated map is a covering space. That is, if $q = \pi(p)$, then there is a small disk Δ containing q such that $\pi^{-1}(\Delta^*) \to \Delta^*$ (for Δ^* the punctured disk) is a covering. This is because singular locus is a finite set.

So, any covering of a punctured disk Δ^*, say the unit disk in the complex plane, is given in the form $\Delta^* \stackrel{z \mapsto z^m}{\to} \Delta^*$. No matter how bad the singularity may look at first, be assured that if one eliminates the singular point and its pre-image, then you just get a disjoint union of punctured disks. In this case, one fills in each connected component (which is a punctured disk) by adding a point (to make these punctured disks disks). That's it. This is a procedure that gives a *new* complex manifold that projects onto the x-line.

So, after we make the transformation in the last item, then we get a *Riemann surface* projecting to the x-line. We have resolved the singularities.

2.1 Example. Consider the equation $y^2 - x^2$, projecting to the line. There is a singularity at the origin. If we take the pre-image over a punctured disk in the x-line, we get a union of two punctured disks. So the associated Riemann surface has two points above the origin, one for each punctured disk.

2.2 Example. Consider the curve defined by $y^2 - x^3$. We take the pre-image over the x-punctured-disc $\{x : 0 < |x| < \epsilon\}$. The curve looks analytically locally like one punctured disk, so there is one point over the origin in the associated Riemann surface.

2.3 Example. Consider $y^3 - x^2$; here the x-projection is locally three-to-one. The pre-image over a small punctured disk in the x-plane is connected, so one point lies over the origin in the Riemann surface.

So, right now we have associated a branched cover of $\mathbb{C}$ by a *Riemann surface* X to any polynomial equation $P(x, y)$, even if this is singular (by possibly adding points). We want, however, a compact Riemann surface in the end. We take the complex numbers $\mathbb{C}$ and complete it to $\mathbb{P}^1$. We want to extend the branched cover $X \to \mathbb{C}$ to a branched cover $\overline{X} \to \mathbb{P}^1$ such that $\overline{X}$ is a compact Riemann surface. To do this, take the complement of a large disk, so something of the form $U_R = \{x : |x| > R\}$. Let $\pi : X \to \mathbb{C}$ be the projection. Then $\pi^{-1}(U_R) \to U_R$ is a *covering space* for $R \gg 0$. So we get a covering space of the *punctured disk* U_R. Again, we can complete this to a compact Riemann surface $\overline{X}$ by adding one point in each connected component of $\pi^{-1}(U_R)$.

In this way, we complete X to a compact Riemann surface.

2.4 Theorem. *To each equation $P(x, y) = 0$, we can associate a branched cover $\overline{X} \to \mathbb{P}^1$ where X is a compact Riemann surface, by the above procedure.*

The above construction gives a means of computing the number of points in the fiber above ∞, for instance: one has to consider the number of connected components in the pre-image of a large disk.

2.5 Example. Consider $y^2 = x^3 - 1$. Consider this as a two-sheeted cover of the x-line. The branch points occur at the cube roots of unity. Over the complement of a large disk in the x-line, the projection is a covering space. We only need to know whether the pre-image of said complement is connected.

It is. To see this, let's wander around a bit. As x wanders around a large circle and gets back to the same point, the argument increases by 2π, so the argument of $x^3 - 1$ increases by 6π if we are wandering around a large circle. Thus the argument of y increase by 3π, so y flips. It follows that the cover is connected, and one adds one point. (The idea is that analytically continuing a germ of a solution of y in terms of x around a circle gives the opposite solution.)

This is a *recipe.* Starting with a polynomial in two variables, you arrive at a compact Riemann surface. You can do the same thing, *purely algebraically.* For one thing, you need to make a *complete* variety (this is the algebraic way of saying "compact"), and we also want to make it smooth. Typically, what you would do algebraically is to start with the polynomial $P(x, y)$, and take the associated curve in $\mathbb{A}^2$. We can take the closure of this in $\mathbb{P}^2$, which is the hypersurface in $\mathbb{P}^2$ cut out by the homogeneization by $P(x, y)$. Then, one has to *normalize* the projective curve by blowing up to resolve singularities. If you want to carry this out algebraically, here is an alternative.

Warning: Don't do this. If you can take the analytic route, use that instead. You might have to blow up a curve multiple times to resolve a singularity (e.g. for a curve like $y^a - x^b$), while the analytic approach is very simple. Moreover, the operation of taking projective closure is not well-behaved. There is nothing to suggest that the curve necessarily wants to live in $\mathbb{P}^2$: taking the projective closure even of a smooth affine curve may give a non-smooth curve. Eye-balling polynomials and thinking about Riemann surfaces is easier.

§2 IOUs from last time: the degree of K_X, the Riemann-Hurwitz relation

There are several things owed from last lecture. If X is a compact Riemann surface of genus g, we want to claim:

2.6 Proposition. $\deg(K_X) = 2g - 2$ *for* K_X *the canonical divisor.*

This is going to come out of the Riemann-Hurwitz relation. Namely, let X, Y be compact Riemann surfaces. Let $g(X) = g, g(Y) = h$. Let $f : X \to Y$ be a nonconstant holomorphic map. By standard complex analysis, this means that the map f has finite degree d. So, over all but finitely many points of Y, the pre-image has cardinality d. There will be a finite number of ramification points in X, where f will fail to be a local isomorphism. (The images of the ramification points, which lie in Y, are called the *branch points.*)

2.7 Definition. Given $p \in X$ mapping to $q \in Y$, we can choose local coordinates near p, q such that f looks like $f : z \mapsto z^m$; that is, it looks like a standard m-fold cover of a disk by a disk. In this case, we say that the **ramification index** is $m - 1$. When $m \geq 2$, then there is **ramification** at p. We denote $m - 1$ by $\nu_p(f)$.

2.8 Definition. Let R be the **ramification divisor** on X. This is $\sum_{p\in X}\nu_p(f)p$. The **branch divisor** B is the image of R under f. This "image" is taken in the naive sense. So $B = \sum_{q\in Y}\left(\sum_{p\in f^{-1}(q)}\nu_p(f)\right)q$.

In particular, if we write $B = \sum n_q q$, then the cardinality of $f^{-1}(q)$ is $d - n_q$. When you have ramification, you lose points in the fiber (so there will be less than d).

Now we derive Riemann-Hurwitz. Consider X with all the points in the branch divisor removed. Take $q_1, \ldots, q_\delta$ be the points appearing in the branch locus. We consider $X' = X - f^{-1}(\{q_1, \ldots, q_\delta\})$, which maps by a *covering space map* to $Y' = Y - \{q_1, \ldots, q_\delta\}$. So $X' \to Y'$ is a d-sheeted covering space map. In particular, $\chi(X') = d\chi(Y')$. We know these Euler characteristics. So $\chi(Y') = \chi(Y) - \delta = 2 - 2h - \delta$. Similarly, $\chi(X') = \chi(X) - \sum_i (d - n_q)$ for $n_q = |f^{-1}(q)|$, and $\chi(X) = 2 - 2g$. We have just used the fact that removing points decreases the Euler characteristic accordingly.

We are left with:

2.9 Theorem (Riemann-Hurwitz)**.** *Notation as above,*

$$2 - 2g = d(2 - 2h) - \deg(B).$$

This is often written as

$$g - 1 = d(h - 1) - \frac{b}{2}.$$

One way to remember this, which generalizes to higher dimensions, is that the Euler characteristic $\chi(X)$ upstairs is exactly what it would be if $X \to Y$ were a covering space—that is, $d\chi(Y)$—with correction terms for the ramification. This in fact generalizes to higher dimensions, but that's for another course.

Here's another thing. Suppose $f : X \to Y$ is a map. Let ω be *any* meromorphic differential on Y. For simplicity, let's assume that $\operatorname{div}(\omega)$ is supported away from the branch points. So there are no zeros or poles of ω on these branch points. What's the divisor of $f^*\omega$? But if ω has a pole or a zero or downstairs, then pulling back gives it a pole or zero of the *same degree* upstairs except at the branch points. At the branch points, if f looks locally like $z \mapsto z^m$, then the differnetial ω looks locally like dz, so the pull-back $(m-1)z^{m-1}dz$ vanishes to degree $m - 1$. That is:

> The pull-back of the canonical divisor on Y is the canonical divisor on X plus the ramification divisor. That is, $\operatorname{div}(f^*\omega) = f^*\operatorname{div}\omega + R$ for R the ramification divisor.

Now let's compare this formula with the Riemann-Hurwitz relation. We find that the $2g - 2$ formula we want to prove can be proved via passing to a branched cover:

2.10 Proposition. *If the degree of the canonical divisor of Y is $2h - 2$, then that of X is $2g - 2$.*

Since any compact Riemann surface admits a map to $\mathbb{P}^1$, and since the $2g - 2$ formula is easily verified for $\mathbb{P}^1$, it follows that the formula holds for all X. That is, we are done with the claim we wanted about the degree of the canonical divisor.

§3 Maps to projective space

We're almost ready to begin the course. So far, we've been reviewing stuff that people are sort of expected to know.

We now want to talk about maps of Riemann surfaces to projective space. As a result, we need a characterization of when a map $X \to \mathbb{P}^n$ is an imbedding. We will derive a criterion. Recall that a *linear series* on X is a pair consisting of a line bundle $\mathcal{L}$ on X of degree d, and a vector subspace $V \subset H^0(\mathcal{L})$ of dimension $r+1$. We write this as a g^d_r for short.

2.11 Definition. If the sections in this subspace V have no common zeros, then we say that V is **base-point free.**

A base-point free linear series defines a regular map $X \to \mathbb{P}^r$. In general, a linear system will have common zeros ("base points").

Note that a global section $\sigma \in H^0(\mathcal{L})$ is determined, up to scalar, by its divisor (or divisor of zeros). The reason is simply that any two sections differ by a meromorphic function, and the only meromorphic functions on X with no zeros or poles are the constants. So, by associating to each section the divisor, we get a *family* of effective divisors D of degree d on X, parametrized by the projectivization $\mathbb{P}(V)$. In other words, a g^d_r corresponds to a family of effective divisors of degree d, parametrized by $\mathbb{P}^r$. Very often, when people talk about linear series, they mean a family of divisors in this sense, and then denote the family by a $\mathscr{D}$.

If $(\mathcal{L}, V)$ is a linear series on X without base points, then as before we get a map $\phi : X \to \mathbb{P}(V^\vee) = \mathbb{P}^r$. We can do this as follows. If E is any effective divisor, we let $V(-E)$ be the set of sections in V whose divisors are at least E. So, to give the map $X \to \mathbb{P}(V^\vee)$, we send each point $p \in X$ to the hyperplane $V(-p) \subset V$, and consider that as a line in $\mathbb{P}(V^\vee)$.

More concretely, if $\mathcal{L} = \mathcal{O}_X(D)$, then we can think of $H^0(\mathcal{L})$ as consisting of meromorphic functions with suitable restrictions on poles (or zeros). Then the map $X \to \mathbb{P}(H^0(\mathcal{L})^\vee)$ can be thought of by choosing a basis $f_0, \dots, f_r \in V$ of V of meromorphic functions, and then considering the homogeneous vector

$$x \in X \mapsto [f_0(x), \dots, f_r(x)].$$

This is a very concrete construction.

We make the following observation.

2.12 Proposition. *The map $\phi : X \to \mathbb{P}^r$ associated to a linear series is characterized, up to automorphisms of $\mathbb{P}^n$, by the property that $\phi^{-1}(H)$ are exactly the divisors in the linear series $\mathcal{D}$ on X.*

This is true either in algebraic geometry or in complex geometry, in *any dimension.* We don't just have to work with Riemann surfaces (though we can't use the same language of divisors).

We want a condition that the map $\phi : X \to \mathbb{P}^r$ associated to a linear series without base-points be an imbedding.

2.13 Proposition. *ϕ is an imbedding if and only if:*

1. *For all pairs $p, q \in X$, the subspace $V(-p-q)$ of sections (of the line bundle, contained in V) vanishing at both p, q has dimension $\dim V - 2$. This is equivalent to $V(-p) \neq V(-q)$, since both have codimension one in V.*

2. *$V(-2p)$ is properly contained in $V(-p)$ for each $p \in X$.*

Proof. The first condition is equivalent to ϕ being set-theoretically injective; this is a direct translation of what we defined. The second condition corresponds to ϕ being an *immersion.* It states that there is a function in V, coming from this map, that vanishes at p, but does not vanish to order two. ▲

We leave the next result as an exercise.

2.14 Proposition. *If $\mathcal{L}$ has degree $d \geq 2g+1$ and V is the* complete linear series *(of all sections of $\mathcal{L}$), then the associated map ϕ to projective space is an imbedding.*

So, if we have a large degree line bundle on a compact Riemann surface, then we get an imbedding in projective space. In particular, we *can* realize any compact Riemann surface as a smooth projective curve. We are going to spend a lot of time trying to *do better.* Say you can imbed a Riemann surface in some huge projective space by a line bundle of degree 10^{100}—this will be too complicated, and won't let us get a handle on the geometry of the curve. We want to know what the *smallest possible degree* of a map to projective space.

Here's another question we'd like to ask. What is the smallest degree of a nonconstant meromorphic function on a general Riemann surface? We can always express a Riemann surface of genus g as a branched cover of projective space of degree $2g+1$; however, we can do better. That'll be one of our main goals.

§4 Trefoils

In the last five minutes, we will do something random and fun. Consider the curve $X \subset \mathbb{C}^2$ given by $y^2 = x^3$. The real picture of this curve is familiar. What do the complex points look like? It's hard to draw the complex plane $\mathbb{C}^2$ on a chalkboard, and even harder to draw it while live-TeXing a course. But say there is a small ball B_ϵ around the origin of radius ϵ in $\mathbb{C}^2$, and consider the intersection of the curve X with $\partial B_\epsilon \simeq S^3$. For ϵ small, the curve X will intersect ∂B_ϵ in a *real* 1-manifold. This is compact, and so a union of S^1's—in fact, it is a link contained in $\partial B_\epsilon = S^3$. What is it?

Well, we are restricting to the case $|y|^2 = |x|^3$ and $|x|^2 + |y|^2 = \epsilon$. These two equations determine the absolute value of x, y. All we have to do is determine the arguments. All of this is to say that the intersection $X \cap \partial B_\epsilon$ is a trefoil, a torus knot, of type $(2,3)$. The locus of x, y in this sphere such that $|x|, |y|$ are fixed is a torus, parametrized by the arguments—the equation of the curve is a linear restriction on these arguments. So you can think of this very singular variety as something like a *cone* on this trefoil knot.

Lecture 3
9/9

The second homework is up. There will be a class 3-4 next Monday (technically, called the "Basic Notions Seminar"), and no class next Wednesday. There will be a class on Friday at an undisclosed location. There are three things to do today. The first is to finish the discussion we were having on characterizing when a map $X \to \mathbb{P}^n$ associated to a linear series is an imbedding.

§1 The criterion for very ampleness

Let us recall the basic theorem from last time:

3.1 Proposition. *If $(\mathcal{L}, V)$ is a linear system on the curve X, then the map $\phi : X \to \mathbb{P}(V^\vee)$ is an imbedding if and only if, for all $p, q \in X$, $\dim V(-p-q) = \dim V - 2$ (where $V(-p-q)$ is the subspace of sections vanishing on the divisor $p+q$).*

For $p \neq q$, this states that the map is injective on points. For $p = q$, this states that the map is injective on the tangent space.

3.2 Corollary. *If $\mathcal{L}$ has degree $d \geq 2g+2$ and V is the complete linear series $H^0(\mathcal{L})$, then $(\mathcal{L}, V)$ defines an imbedding of X in projective space.*

Proof. Let D be the associated divisor.

We can check the previous result by Riemann-Roch. Indeed, Riemann-Roch states that $\ell(D) = d - g + 1 + \ell(K-D)$, and likewise

$$\ell(D-p-q) = d - 2 - g + 1 + \ell(K - D + p + q).$$

However, $\ell(K-D) = \ell(K-D+p+q) = 0$ because these two divisors have degree < 0 (as $\deg K = 2g-2$ and $\deg D$ is large). As a result, the claim follows: $\ell(D-p-q) = \ell(D) - 2$. ▲

One can see more from this.

3.3 Example. If $\mathcal{L}$ is a line bundle of degree $2g$. When would it fail to be an imbedding? The only way it would fail would be if $\ell(K-D+p+q)$ admitted a nonzero section, and since $K - D + p + q$ has degree zero, this would happen precisely when $K - D + p + q$ was principal for some p, q.

Now let's look at the *canonical class.* Let's say $g \geq 2$. This is much more important to us. If we take $\mathcal{L}$ the canonical class, we get a map $\phi_K : X \to \mathbb{P}^{g-1}$ given by the canonical divisor. This will be an imbedding precisely when

$$\ell(K-p-q) = \ell(K) - 2,$$

and the first term, by Riemann-Roch, is

$$2g - 4 - g + 1 + \ell(p+q) = g - 3 + \ell(p+q).$$

Of course, $\ell(p+q) \geq 1$ (because of the constant function), but the problem would be if this happened to be two-dimensional. If there is a nonconstant meromorphic function with just two poles p, q on X, then the canonical divisor fails to be an imbedding. We conclude:

3.4 Corollary. *The map $\phi_K : X \to \mathbb{P}^{g-1}$ from the canonical divisor is an imbedding unless there exists a divisor $D = p + q$ on X with $\ell(D) = 2$. In other words, unless there exists a nonconstant meromorphic function of degree two on X.*

§2 Hyperelliptic curves

As a result of the last section, we make:

3.5 Definition. We say that a compact Riemann surface X is **hyperelliptic** if (equivalently) X has a global meromorphic function of degree 2, or if X is expressible as a 2-sheeted (branched) cover of $\mathbb{P}^1$.

So for non-hyperelliptic curves of genus ≥ 2, the canonical divisor induces an imbedding if and only if X is *not* hyperelliptic. We have now shown that there are non-hyperelliptic curves, but we will see this soon enough. Once you get to genus 3, *most* curves are non-hyperelliptic. One question we'll raise later is what degree one needs in general to express a curve as a branched cover of the sphere.

For a non-hyperelliptic curve, we call $\phi_K(X) \subset \mathbb{P}^{g-1}$ the **canonical model** of X. The whole point is to understand the connection between the abstract Riemann surface X and concrete subvarieties of projective varieties in $\mathbb{P}^n$. There are *lots* of ways of imbedding a Riemann surface in projective space. This is a canonical one for a non-hyperelliptic curve of genus ≥ 2.

Here is a fun application.

3.6 Proposition (Geometric Riemann-Roch)**.** *Let X be a non-hyperelliptic curve of genus ≥ 2, thus imbedded $X \hookrightarrow \mathbb{P}^{g-1}$ via an imbedding of degree $2g-2$. Consider a divisor $D = p_1 + \cdots + p_d$ consisting of d distinct*[1] *points. Then $r(D)$ is the number of linear relations on the points $p_1, \dots, p_d$.*

Note that this relates the intrinsic properties of X to the extrinsic properties of X considered as a subscheme of $\mathbb{P}^{g-1}$ (via the canonical morphism).

Proof. Then, $r(D) = d - g + \ell(K-D)$ by the Riemann-Roch theorem. What is $\ell(K-D)$? Well, this is the space of linear forms on the canonical space that vanish at the $\{p_i\}$. By construction, since $X \subset \mathbb{P}^{g-1}$, the one-forms on X correspond to linear forms on $\mathbb{P}^{g-1}$ because we have used the canonical imbedding. What is the dimension that vanish at $p_1, \dots, p_d$? There are a lot of linear forms on $\mathbb{P}^{g-1}$, and we have imposed d conditions on them if the forms are to vanish on the $\{p_i\}$. However, the $\{p_i\}$ might not be linearly independent. So the dimension of the space of such forms is

$$g - d + \{\# \text{ of relations satisfied by } p_1, \dots, p_d\}.$$

However, this is a cancellation, and this is just the number of linear relations on $p_1, \dots, p_d$. ▲

[1] This applies to arbitrary divisors of degree d if one interprets the terms appropriately.

We say that a curve is *trigonal* if it is a three-sheeted cover of $\mathbb{P}^1$. This means that there is a divisor of degree 3 moving in a pencil (i.e. fits into a nontrivial linear series, corresponding to the map to $\mathbb{P}^1$), and this will be the case if and only if the canonical model contains three collinear points.

§3 Properties of projective varieties

We will be working with curves in projective space, and will need some tools for that.

We start with a general remark. Everything said last week about divisors on curves applies to divisors on an arbitrary smooth projective variety. If X is such an object (or alternatively a complex manifold), of dimension n, Then a *divisor* on X is a formal linear combination of irreducible varieties of codimension one, or of dimension $n-1$. So $D = \sum n_i Y_i$ is a formal linear combination, as before. We have the same notion of *linear equivalence*: $D \sim E$ if there is a meromorphic function on X whose divisor is $D - E$ (we can, as before, talk about the divisor of a meromorphic function). Once more, we say that $\mathrm{Pic}(X)$, the group of *line bundles* on X, is the set of all divisors on X modulo the relation of linear equivalence, at least if such line bundles always admit meromorphic sections. (This is always true for an algebraic variety, though not necessarily for a complex manifold.)

Note that a compact complex manifold need not have any nontrivial meromorphic functions, and in fact no nontrivial divisors. A general torus $\mathbb{C}^2/\Lambda$ for Λ a random rank four lattice in $\mathbb{C}^2$, then one gets a compact complex manifold, homeomorphic to $(S^1)^4$, which in general will have no nonconstant meromorphic functions and no nontrivial subvarieties (besides points). In this case, the previous discussion is somewhat uninteresting. However, one can also find line bundles that *do not come from divisors.*

3.7 Example. Consider $X = \mathbb{C}^2/\Lambda$ where $\Lambda \sim \mathbb{Z}^4$, where Λ is generated by $(1,0), (0,1), v_1, v_2$, for $v_1, v_2 \in \mathbb{C}^2$. Here $v_1 = \alpha e_1 + \beta e_2, v_2 = \gamma e_1, \delta e_2$. We consider the matrix of coefficients,

$$\begin{bmatrix} \alpha & \beta \\ \gamma & \delta. \end{bmatrix}$$

The condition for the existence of non-constant meromorphic functions on X is that, for suitable choices of generators of a finite-index sublattice, the matrix above is *symmetric* with *positive-definite* imaginary part. The space of such matrices is called the Siegel upper-half space. This is analogous to sublattices of $\mathbb{C}$ generated by $\{1, \tau\}$, which parametrize elliptic curves. A general lattice will not satisfy this condition.

We can also talk about maps to projective space induced by a divisor on a smooth variety. We won't need this so much. More important to us will be the *canonical bundle* $K_X = \bigwedge^n T_X^*$, the top exterior power of the cotangent bundle. The sections are holomorphic differential forms of top degree. So these look like $f(z_1, \dots, z_n) dz_1 \wedge \cdots \wedge dz_n$. This corresponds to a divisor class on X, which is the canonical divisor: take a global meromorphic n-form on X, and its divisor represents K_X.

Warning. There is generally no notion of *degree* for divisors on higher-dimensional varieties.

We start with a basic example:

3.8 Example. Take $X = \mathbb{P}^n$. Any irreducible $n-1$-dimensional subvariety $Y \subset \mathbb{P}^n$ is the zero locus of a single irreducible homogeneous polynomial F. The degree of that polynomial is the degree of Y. The basic observation is that Y is linearly equivalent to Y' if and only if $\deg Y = \deg Y'$. If Y is the zero locus of F and Y' the zero locus of G, then $[Y] - [Y']$ is the divisor of the global meromorphic function F/G if F, G have the same degree.

So there are many, many divisors on $\mathbb{P}^n$, but $\mathrm{Pic}(\mathbb{P}^n) = \mathbb{Z}$.

We're going to need a name for the generator of $\mathrm{Pic}(\mathbb{P}^n)$. We could just take the divisor class of a hyperplane, by the previous example.

3.9 Definition. The divisor class associated to a hyperplane (or the associated line bundle) in $\mathbb{P}^n$ is denoted $\mathcal{O}(1)$. We will also just call it h, for "hyperplane." So $\mathcal{O}(1)$ is a line bundle on $\mathbb{P}^n$, which generates the Picard group.

We need to compute the canonical divisor class of $\mathbb{P}^n$. To do this, it's not that complicated: you write down a meromorphic differential, and see where it has zeros and poles. To write down a meromorphic differential, just write something down in an affine open subset.

3.10 Example. In $\mathbb{P}^n$, we could take $dz_1 \wedge \cdots \wedge dz_n$ on the affine subspace $\mathbb{A}^n \subset \mathbb{P}^n$: this is holomorphic and nonzero on the affine subspace. What happens when you go to the hyperplane at ∞? One can make the standard change of coordinates to go to any hyperplane at ∞; one finds that this differential has a *pole* of order $n+1$ at that hyperplane. One has proved:

3.11 Proposition. *The canonical bundle of* $\mathbb{P}^n$ *is* $\mathcal{O}(-n-1)$.

§4 The adjunction formula

Now let's do something that will make the previous section very useful. Let X be a smooth variety, or alternatively a complex manifold, and let $Y \subset X$ be a smooth subvariety (or manifold) of codimension one. Then Y is a divisor, among other things. We want to relate the canonical bundles of X, Y.

3.12 Proposition. *The canonical bundle* K_Y *is* $(K_X \otimes \mathcal{O}(Y))|_Y$*, where* $\mathcal{O}(Y)$ *is the line bundle on* X *associated to* Y.

Proof. This will be based on two facts. Say $\dim X = n$. First, there is an exact sequence

$$0 \to T_Y \to T_X|_Y \to N_{Y/X} \to 0,$$

where $N_{Y/X}$ is the normal bundle. As a result, we find

$$\bigwedge^n T_X^*|_Y = \bigwedge^{n-1} T_Y^* \otimes N_{Y/X}^*$$

or the canonical bundle of Y is the canonical bundle of X (restricted to Y) tensored with the normal bundle. Next, we claim that $N_{Y/X} = \mathcal{O}(Y)|_Y$. This is essentially a tautology in the algebraic geometry world because the conormal bundle is defined by the ideal sheaf mod the ideal sheaf squared. In the analytic sense, one doesn't define the normal bundle in this way, so one has to work a tiny bit to see that $N_{Y/X} = \mathcal{O}(Y)|_Y$. ▲

3.13 Example. Consider a curve $X \subset \mathbb{P}^2$, which is a smooth plane curve of degree d. The line bundle $\mathcal{O}(d)$ on $\mathbb{P}^2$ is also the line bundle $\mathcal{O}(X)$ associated to X (i.e. to the divisor). Since $K_{\mathbb{P}^2} = \mathcal{O}(-3)$, we find

$$K_X = \mathcal{O}(d-3)|_X = \mathcal{O}_X(d-3).$$

(Here $\mathcal{O}_X(1)$ is by definition $\mathcal{O}(1)$ restricted to X.) We can thus obtain the canonical divisor on X simply.

For a general hyperplane (i.e. line) in $\mathbb{P}^2$, it will intersect $\mathbb{P}^2$ in d points. So $\mathcal{O}(1)|_X$ has degree d as a line bundle on X, which implies that $\deg K_X = d(d-3)$. Since $\deg K_X = 2g-2$, we get the familiar formula

$$g = \binom{d-1}{2}.$$

3.14 Example. Consider a smooth quadric $Q \simeq \mathbb{P}^1 \times \mathbb{P}^1 \subset \mathbb{P}^3$. It's not hard to see that $\mathrm{Pic}(Q) = \mathbb{Z} \times \mathbb{Z}$. This is generated by e, f for e, f copies of $\mathbb{P}^1$ in opposite directions. One can show by similar reasoning that

$$K_Q \sim -2e - 2f.$$

Namely, the canonical bundle K_Q is $(\mathcal{O}(2) \otimes \mathcal{O}(-4))|_Q$ because $\mathcal{O}(2)$ is the line bundle on $\mathbb{P}^3$ associated to a quadric, and $\mathcal{O}(-4)$ is the canonical bundle on $\mathbb{P}^3$. One can check that $\mathcal{O}(1)|_Q \sim e + f$, directly. This proves the claim.

3.15 Example. Let $X \subset \mathbb{P}^1 \times \mathbb{P}^1$ be a smooth curve. Suppose we have that $X \sim ae+bf$. This means that X meets a line of the first ruling b times and a line of the second ruling a times. So X is the zero locus of a bihomogeneous polynomial of bidegree (a, b). What is the genus?

The canonical bundle on X is $((-2e-2f) \otimes \mathcal{O}_{\mathbb{P}^1 \times \mathbb{P}^1}(X))_X$. This is $(a-2)e+(b-2)f$, restricted to X. However, e meets X b times and f meets X a times. Consequently, the degree of the canonical bundle is

$$(a-2)b + (b-2)a,$$

so the genus is

$$g = (a-1)(b-1).$$

§5 Starting the course proper

The theory of curves got started in earnest about 200 years ago, with the introduction of the complex numbers and projective space. However, that's a long time, and people have gotten familiar with curves of low degree or genus in projective space. As a result, people know what to expect because there are so many known examples of curves. That's what we'll be doing for the next few weeks in this course. We'll intersperse this with the introduction of new techniques.

Let us first consider the case of genus zero.

3.16 Example. Let X be a compact Riemann surface of genus zero. Topologically, X looks like a sphere. There is only one such Riemann surface: $X \simeq \mathbb{P}^1$. To see this, let $p \in X$, and consider the associated divisor D. Then $\deg D = 0$, and by Riemann-Roch (since $\deg D > 2g - 2$), there are two sections. Ignoring the constant section, we get a degree one map $X \to \mathbb{P}^1$.

Alternatively, we showed in general that a line bundle of degree $2g + 1$ gives an imbedding in projective space. So, for $g = 0$, we just take the line bundle associated to a point. This is the only case where there is only one curve of a given genus.

How can we imbed X in projective space? Given a divisor $D = dp$, we can use it to imbed X in projective space. For instance, take $p = \infty$; then $\mathcal{L}(D)$ consists of the vector space of functions spanned by $1, z, \dots, z^d$. We get a map $X \to \mathbb{P}^d$, $z \mapsto [1, z, \dots, z^d]$. The image is called a *rational normal curve* of degree d. For instance, when $d = 2$, we get that X is isomorphic to a *smooth plane conic.*

For $d = 3$, we get the twisted cubic $X \hookrightarrow \mathbb{P}^3$ via $z \mapsto [1, z, z^2, z^3]$. The equations of the twisted cubic can be written down explicitly. They can be written as minors of a matrix. For $[A, B, C, D] \in \mathbb{P}^3$, the equations defining the twisted cubic are $AC - B^2, AD - BC, BD - C^2$. Note that any two of these are insufficient, though they will give a variety of dimension one.

Remark. For any two quadrics containing a twisted cubic containing the twisted cubic in $\mathbb{P}^3$, the intersection is the twisted cubic plus a chord. This is an interesting question, though I didn't really understand the discussion.

Remark. There is an open problem of whether curves in $\mathbb{P}^3$ are set-theoretic complete intersections. For instance, the twisted cubic is a *set-theoretic* complete intersection.

We now want to at least introduce some of the ideas for the homework problems. When considering rational curves in projective space $\mathbb{P}^2$, we have only so far covered rational normal curves. These are those defined by the complete linear system of a divisor dp, $p \in \mathbb{P}^1$, to imbed $\mathbb{P}^1 \hookrightarrow \mathbb{P}^d$. However, we could use a non-complete linear system. So we could consider a *subspace* of the polynomials of polynomials of degree $\leq d$, and use that to imbed $\mathbb{P}^1$ in some projective space. In this way, a *general* rational curve in projective space come from the previous imbedding in $\mathbb{P}^d$ (the Veronese imbedding) followed by a projection to a smaller projective space.

3.17 Example. Let's consider an example of a rational curve $\phi : \mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ not given in the previous example (rational normal curves), so we use a proper subspace of the polynomials of a given degree, via a quartic curve. So we consider $\mathcal{O}(4)$ on $\mathbb{P}^1$, and take a four-dimensional space of sections (the space of sections is five-dimensional). Now we'd like to ask various questions: what surfaces does this rational curve lie on? Does it lie on a quadric?

By definition, $\phi^* \mathcal{O}_{\mathbb{P}^3}(1) = \mathcal{O}_{\mathbb{P}^1}(4)$. Correspondingly, $\phi^*(\mathcal{O}_{\mathbb{P}^3}(2)) = \mathcal{O}_{\mathbb{P}^1}(8)$, which corresponds to homogeneous polynomials of degree eight. There is a map from homogeneous polynomials of degree two (quadratic) on $\mathbb{P}^3$ to octics on $\mathbb{P}^1$; this is the map

$$H^0(\mathcal{O}_{\mathbb{P}^3}(2)) \to H^0(\mathcal{O}_{\mathbb{P}^1}(8)).$$

However, the latter space has dimension nine. The first space has dimension ten. That implies that this map has a kernel. As a result, the image curve *must lie in a quadric.*

In fact, such a quadric turns out to be necessarily smooth. (**Exercise.**) The image of a map $\mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ given by a non-complete linear system of degree four must have type $(1, 3)$ on a quadric surface.

Lecture 4
9/12

These are notes from the "Basic Notions seminar," which was combined with the lecture for this course.

§1 Motivation

It seems that sometime between the nineteenth and twentieth centuries, there was a fundamental shift in viewpoint. People were still studying the same objects, but the way in which they were studied changed. For instance, in the nineteenth century, what was a *group*? A group was more or less defined to be a *subset* of GL_n, closed under multiplication and inversion. There was a notion of *isomorphism* of groups, but one didn't have groups that didn't come with this inclusion in GL_n. In the twentieth century, what we would now call a group was defined. A group then became just a set with a binary operation satisfying the basic axioms. This was this shift from defining objects as subobjects of some given standard objects to intrinsically defined objects.

As another example, in the nineteenth century, a *manifold* was a subset of $\mathbb{R}^n$ defined by smooth functions with independent differentials. In the twentieth century, it became a *set* with additional structure (a locally ringed space looking locally like $\mathbb{R}^k$). In the twentieth century, manifolds stopped being assumed imbedded in the ambient space; one builds the objects from the ground up.

In algebraic geometry, the situation was very similar. In the nineteenth century, an algebraic variety was simply a subset of projective space $\mathbb{P}^n$ defined over polynomial equations, while now one works with *abstract varieties.*

This shift in viewpoint led to a wholesale restructuring of the various fields. For instance, there is now the classification of *abstract* groups (as opposed to imbedded groups), as well as the second half of *representation theory:* given an abstract group, consider the set of all ways to consider it as imbedded in an ambient GL_n. In the same way, the theory of algebraic curves broke up into two parts after the shift: first was the classification to *abstract curves* (here all curves will be assumed smooth and projective). Here one knows that there is a discrete invariant g, the genus, and in each invariant locus a space M_g parametrizing curves in that space. But there is the other half of the problem: given a curve X, classify maps $X \to \mathbb{P}^r$ to projective space. It's this second half we'd like to discuss today.

Note that with this shift in viewpoint, it's hard to think of classes of abstract objects as parametrized by a given object. There is no moduli space of groups, for instance. But in algebraic geometry, a defining characteristic of the field is that the class of algebro-geometric objects of a given form becomes itself an algebro-geometric object. For instance, there is a space M_g parametrizing curves of genus g. It turns out that we can *stratify* this moduli space M_g according to the answer to the question we're now discussing, that is how the curve can map to projective space.

In other words, we're interested in: does a curve X admit a map of degree d to $\mathbb{P}^r$? The curves that do form a locally closed subset of the moduli space. Given this fact, there should be a *generic* answer to the question. The "generic" curve of genus g should behave similarly. This is the Brill-Noether problem.

1. *For general genus g X, describe the space of non-degenerate*[2] *maps $X \to \mathbb{P}^r$ of degree d.*

In particular, we might ask whether this space is nonempty. We ask further questions:

2 If X is generic of genus g, what is the smallest degree d of a non-constant map $X \to \mathbb{P}^1$?

3 What is the smallest degree of a plane curve (in $\mathbb{P}^2$) birational to X?

4 What is the smallest degree of the smallest imbedding in $\mathbb{P}^3$?

Since the set of maps $X \to \mathbb{P}^r$ is not just a set of maps, but a *space* of maps, we can ask how it behaves. We will just focus today on existence questions.

§2 A really horrible answer

We're going to take a naive approach (which is how people dealt with these things in the nineteenth century). Note that M_g wasn't known about until 1969, but people informally worked with it much earlier. They thought of Riemann surfaces as continuously depending on the coefficients and polynomials.

More formally, let's not consider the space M_g parametrizing abstract curves. Let's instead look at space $H_{d,g}$ of pairs (X, f) where X is an abstract Riemann surface and $f : X \to \mathbb{P}^1$ is a map of degree d, and simply branched. We want a map from $H_{d,g}$ to $\mathbb{P}^b \setminus \Delta$ by sending a branched cover to the branch divisor (we have identified $\mathbb{P}^b$ with the symmetric power of $\mathbb{P}^1$). This is well-behaved. Once you describe the branch divisor, all you have to do to describe X is to describe the monodromy (which is a finite set of data). As a result, $H_{d,g}$ is a finite "cover" of $\mathbb{P}^b - \Delta$, so $H_{d,g}$ has dimension $b = 2g + 2d - 2$ (by Riemann-Hurwitz).

Let us now try to describe the fibers of $H_{d,g} \to M_g$. Once you have a compact Riemann surface, to give a map to $\mathbb{P}^1$ is just a rational function of a given degree. Take $d \gg 0$ (say $d \geq 2g + 1$), which implies that to specify a rational function, we can start by specifying a polar divisor of degree d. When d is large, Riemann-Roch tells us the answer to how many such functions there are. The space of meromorphic functions with poles along D is a vector space of dimension $d - g + 1$; an open subset of this space gives simply branched coverings.

As a result, we can conclude that the dimension of M_g is the difference of $\dim H_{d,g}$ minus the dimension of the fiber,

$$2d + 2g - 2 - (d + d - g + 1) = 3g - 3,$$

[2] The image is not contained in a hyperplane.

which gives the famous dimension of the moduli space M_g (informally), i.e. that it is $3g-3$.

Well, we might also ask: when does $H_{d,g} \to M_g$ become a dominant morphism? We know the dimension of $H_{d,g}$, and we know the fibers have dimension at least three (because any map $X \to \mathbb{P}^1$ can be composed with an automorphism of $\mathbb{P}^1$). In fact, $H_{d,g} \to M_g$ is dominant only when $b \geq 3g$, or when $d \geq \frac{g}{2}+1$. This is one of the first cases of Brill-Noether theory that was proved.

Anyway, the lesson is: the space M_g is not the easier one to deal with. Rather, the space $H_{d,g}$ is better. Still, it is a naive approach, and it takes work to make it precise.

§3 Plane curves birational to a given curve

Let us now consider the question of when a curve is birational to a plane curve of a given degree. We consider the space $V_{d,g}$ of pairs (X, f) where X is a Riemann surface and $f : X \to \mathbb{P}^2$ is of degree d such that f is birational onto a plane curve with only nodes. Just as in the case of branched covers, we had a formula for the number of branch points. We can say that the number of nodes is

$$\delta = \binom{d-1}{2} - g,$$

the difference of the "expected genus" (if the image were smooth) minus the actual genus.

We are going to associate to each element of $V_{d,g}$ an element of $\mathrm{Sym}^\delta(\mathbb{P}^2) - \Delta$ by sending each pairs to the set of nodes. We'll now do something really illegitimate: we'll try to estimate $\dim V_{d,g}$ via this map. The claim is that the fibers of this map have a given dimension. Well, the space of all plane curves has dimension $d(d+3)/2$. To have double points at δ given points, then we impose 3δ conditions. We get that

$$\dim V_{d,g} = \frac{d(d+3)}{2} - \delta = 3d + g - 1.$$

This argument is completely bogus if δ is too big. However, this claim is true, and it can be established via deformation theory.

Anyway, so we want to obtain the fibers of $V_{d,g} \to M_g$. By counting fibers (which have dimension at least 8), we find that $V_{d,g}$ has dimension at least eight plus that of M_g. So the requirement is that

$$3d + g - 9 \geq 3g - 3, \quad d \geq \frac{2}{3}g + 2.$$

In general, it is conjectured—and now proved—that the existence problem is that there exists a nondegenerate map from a general Riemann surface of genus g to $\mathbb{P}^r$ of degree d if and only if $d \geq \frac{r}{r+1}g + r$. This was conjectured in the mid nineteenth-century, and was considered more or less self-evident based on known examples. For instance, Enriques works through all examples up to genus seven, and then just states the general formula and moves on. It wasn't proved until at least the 1970's (though the earlier cases for $r = 1, r = 2$ were proved using proper arguments earlier).

The problem is that people still don't know the dimension of the space of curves of a given genus in, say, $\mathbb{P}^3$. It's not the case that this space is irreducible: there are many components of many dimensions, and nobody knows how many components or what the dimensions are. People don't even know if there are non-trivial examples of non-"rigid" curves. We still lack the ability to give a bound on the dimension of the space of curves in a higher-dimensional projective space.

§4 Statement of the result

Here is a part of the Brill-Noether theorem that we need.

4.1 Theorem. *Let $\rho = g-(r+1)(g-d+r)$. Then a general compact Riemann surface of genus g admits a non-degenerate map $X \to \mathbb{P}^r$ of degree d if and only if $\rho \geq 0$ and, in this case, the dimension of the space of such maps is ρ.*

Moreover, for a general such map, f is an imbedding when $r \geq 3$; f is a birational imbedding when $r = 2$.

This also answers the further questions of maps to $\mathbb{P}^r$ and imbeddings. The general map is an imbedding once $r \geq 3$.

We won't prove this. But we'll say a little more. This is really the beginning of the story. If you start with an abstract algebraic curve, there isn't much structure. There is a little: for instance, things like Weierstrass points. But once it is mapped to projective space, it acquires a rich structure. Once X is realized as a closed subscheme of $\mathbb{P}^r$, it obtains various types of structure. We can talk about geometric things like secant planes and trisecant lines, tangent lines, inflectionary points, the Gauss map to the Grassmannian, and so on. But there's also a great deal of algebraic structure that's not just on the abstract curve. For instance, we have an associated *homogeneous ideal* in the coordinate ring, with a rich structure: we can ask what the generators of this ideal, what degrees they are, etc. We can ask what the defining equations for this curve are.

There, we have a conjecture, still unknown. That's what will be described in the remaining nine minutes. The problem now is: for a general X (a Riemann surface of genus g), and a general map $X \to \mathbb{P}^r$ of degree d (with $r \geq 3$, so it's an imbedding), describe a minimal set of generators for the associated homogeneous ideal of the image.

4.2 Example. A genus two curve X can be described as the intersection of a quadric and a cubic in $\mathbb{P}^3$. How can we describe this? We have to pick an effective divisor D of degree five on this curve; by Riemann-Roch, $\mathcal{L}(D)$ has dimension $5-2+1=4$. The vector space can be written as $[1, f_1, f_2, f_3]$ and we get a map

$$X \to C \subset \mathbb{P}^3, \quad x \mapsto [1, f_1(x), f_2(x), f_3(x)].$$

What are the equations of the ideal? What is the lowest degree polynomial in four variables that vanishes on the curve. It's not one, since the curve is non-degenerate. Does this curve lie on a quadric? If you consider the space of homogeneous quadratic polynomials on $\mathbb{P}^3$, and pull it back to X, we get a map to $\mathcal{L}(2D)$. The first space has dimension ten, while $\dim \mathcal{L}(2D) = 9$. So there is a nontrivial quadric that the curve must lie on. (It is also a unique quadric.) If we do the same thing for cubic polynomials, one sees that C lies on six linearly independent cubics. Four of those are the products of the quadrics with linearly independent lines, but we can choose two cubics for C to lie on. Thus, the ideal of C is generated by a choice of one quadric and two cubics. In fact, we have a resolution of the ideal

$$0 \to \mathcal{O}(-4)^2 \to \mathcal{O}(-3)^2 \oplus \mathcal{O}(-2) \to \mathcal{I} \to 0.$$

This example would be a lot harder in $\mathbb{P}^4$. We don't have such a complete description.

We'd hope that the following conjecture is true. If X is a general curve of genus g, $f : X \to \mathbb{P}^r$ of degree d is a general imbedding, and $r \geq 3$, then we can look for each n at the space of polynomials (in $r+1$-variables) of degree n: we get a map from this space to sections of a very ample line bundle on the curve. The conjecture is that this restriction map has maximal rank.

Lecture 5
9/16

§1 Homework

The problems in the homework to this class tend to be open-ended. For the last week, one problem that confused a lot of people was the following. The statement was, if $C \subset \mathbb{P}^d$ is a rational normal curve of degree d, then the normal bundle $N_{C/\mathbb{P}^d}$ should be isomorphic to the direct sum $\mathcal{O}_{\mathbb{P}^1}(d+2)^{d-1}$. Leaving aside the question of how you do this problem—it was covered in recitation—there are lots of ways of seeing this, none of this are direct applications of things we did in class. But just the question of determining the normal bundle to the rational normal curve gets you thinking. *Suppose C is a copy of $\mathbb{P}^1$ imbedded in $\mathbb{P}^r$, of degree d: what normal bundles occur?*

The answer to this question is known when $r = 3$, but *not* in general. Even in the known case, it is far from obvious. Hopefully the course will help give a sense of these sorts of questions of interest in algebraic geometry, through the exercises. Note in particular, you don't have to do all the homework problems.

Another question explored in the homework that is of interest in current research was determining the Hilbert functions of certain varieties, e.g. curves in $\mathbb{P}^3$, are. This leads to the more general case of minimal resolutions of the ideals.

§2 Abel's theorem

There is one main new idea that we will introduce now; after that, we will focus mostly on examples.

One of the motivations in the early nineteenth century for studying algebraic curves was to study the indefinite integrals of algebraic functions. People had already knocked off integrals such as $\int \frac{dx}{\sqrt{x^2+1}}$; quadratic irrationalities could be evaluated. But when they ran into square roots of *cubics*, such as $\int \frac{dx}{\sqrt{x^3+1}}$, they hit a brick wall.

Ultimately, you want to think of the cubic integral not as just a plain integral, but as a *line* integral $\int \frac{dx}{y}$ on a Riemann surface associated to $y^2 = x^3 + 1$. This helped people to understand why the cubic integral was insoluble. When you integrate a differential on a Riemann surface of genus one (e.g. one associated to a cubic), the choice of path between points matters. In other words, the inverse function of the antiderivative is a doubly periodic function in the complex plane; such cannot be elementary.

Let's now try to set this up properly, from the beginning. We want to consider a compact Riemann surface C, of genus g. We want to consider the expression $\int_{p_0}^{p}$. What does this mean? We can consider this as a linear function on one-forms. Let $H^0(K_C)$ be the space of holomorphic 1-forms; we want to get a map $H^0(K_C) \to \mathbb{C}$ given by integration on the path from p_0 to p. *However,* the linear functional thus defined depends on the choice of path $p_0 \to p$. In order to get a well-defined object, we have to quotient out $H^0(K_C)^*$ by the class of such functionals obtained by integration over closed loops.

In other words, for each closed loop γ in C, then we get an element of $H^0(K_C)^*$ by integration over γ, which gives an inclusion

$$H_1(C, \mathbb{Z}) \to H^0(K_C)^*$$

of a group isomorphic to $\mathbb{Z}^{2g}$ of a complex vector space of dimension $2g$. The image turns out to be a (discrete) lattice. We can thus think of the element $\int_{p_0}^{p}$ as an element of this quotient as $H^0(K_C)^*/H_1(C, \mathbb{Z})$.

5.1 Definition. The **Jacobian** of C, denoted $J(C)$, is the quotient $H^0(K_C)^*/H_1(C, \mathbb{Z})$. This is a complex torus of complex dimension g.

We get a map $C \to J(C)$ by sending a point p to $\int_{p_0}^{p}$. This depends on a choice of p_0. We're going to assume that we have chosen a basepoint, but bear that dependence in mind: it becomes important.

We can extend this additively from the curve C to all divisors on the curve. For instance, we send $D = \sum n_i p_i$ to the linear functional

$$\sum n_i \int_{p_0}^{p_i} \in J(C).$$

In particular, if we restrict to restrict divisors of degree d, then we get a map from the dth symmetric product C_d (the space of degree-d divisors) to the Jacobian $J(C)$. We write

$$u_d : C_d \to J(C).$$

Describing this map is in some sense the problem of abelian integrals.

Abel made a crucial observation in the first half of the nineteenth century. He observed the following. We have a map from divisors to $J(C)$.

5.2 Proposition. *Given two linearly equivalent effective divisors D, D' (i.e. two such that the difference is the divisor of a rational function), then D, D' map to the same point in the Jacobian.*

Proof. Indeed, if D, D' are linearly equivalent, then there is a $\mathbb{P}^1$-family of divisors interpolating between D, D'. If we can think of $D - D'$ as the divisor of a meromorphic function f, then we can just think of the level sets of f as a family of divisors $D_t, t \in \mathbb{P}^1$ interpolating between D, D'. As a result, we get a map

$$\mathbb{P}^1 \to J(C), \quad t \mapsto u_d(D_t).$$

However, $J(C)$ is a complex *torus*, as the quotient of a $\mathbb{C}^g$ by a lattice. Thus there are lots of global 1-forms on the Jacobian: we can in fact take translation-invariant 1-forms to span the cotangent space at each point. On the other hand, there are *no* holomorphic 1-forms on $\mathbb{P}^1$. So given a map $\mathbb{P}^1 \to J(C)$, all the 1-forms on $J(C)$ are pulled back to zero on $\mathbb{P}^1$, and since they span the cotangent space at each point of $J(C)$, it follows that $\mathbb{P}^1 \to J(C)$ must have zero differential. So $\mathbb{P}^1 \to J(C)$ must be constant.

Another way of seeing this is that $\mathbb{P}^1$ is simply connected, so any map $\mathbb{P}^1 \to J(C)$ lifts to a map $\mathbb{P}^1 \to \mathbb{C}^g$ (because $\mathbb{C}^g$ is the universal cover), and must be constant by Liouville's theorem. ▲

So two divisors that are linearly equivalent map to the same point in the Jacobian. Abel still couldn't tell you what these integrals were, but he found that if you take the *sums* of these integrals, it depended only on the linear equivalence class of the endpoints. This in some sense was the beginning of modern mathematics, where you don't actually solve the problem, but just make qualitative statements about it.

Sometime later, we have "Abel's theorem" (due to Clebsch, though he gets no credit):

5.3 Theorem (Abel's theorem)**.** *Let D, E be effective divisors of degree d. Then $u(D) = u(E)$ if and only if $D \sim E$.*

This is a much harder theorem. Just knowing that the integrals turn out to be the same, one has to construct a *rational function* with the required divisors and poles. The point is the conclusion: the fibers of the map

$$u : C_d \to J(C),$$

are *exactly the complete linear systems.* In particular, these fibers are all projective spaces.

The typical notation is that if D is a divisor, $|D|$ denotes the set of effective divisors linearly equivalent to D. (This is the complete linear system corresponding to D.) According to Abel's theorem, these are the fibers.

§3 Consequences of Abel's theorem

Keep the same notation.

Let us consider the map

$$u : C_d \to J(C)$$

from the symmetric power to the Jacobian. Recall now the geometric Riemann-Roch theorem: if C is a canonical curve, sitting inside $\mathbb{P}^{g-1}$, then the dimension of the complete linear system associated to a divisor is the number of linear relations on those divisors. It follows that if one takes $\leq g$ points, then there are generally no other divisors linearly equivalent. It follows that $u : C_d \to J(C)$ is generically injective for $d \leq g$. For $d = g$, we have a generically injective map between g-dimensional varieties, so $u : C_g \to J(C)$ is *a birational isomorphism.* (**N.B.** We have not actually shown that we are working with a variety here as $J(C)$.)

5.4 Corollary (Jacobi inversion)**.** *Given a sum* $\sum_g \int_{p_0}^{p_i} + \sum \int_{p_0}^{q_i}$*, then one could always write this as* $\sum \int_{p_0}^{r_i}$ *for algebraic functions in* p_i, q_i*.*

This is what the surjectivity of $u : C_d \to J(C)$ means.

When $d \geq 2g - 1$, then $C_d \to J(C)$ is a *projective bundle*, in fact a $\mathbb{P}^{d-g}$-bundle. This follows from the fact that the r's of large divisors are all of the same dimension. This lets you realize symmetric powers of C as projective bundles over a complex torus.

5.5 Example. In genus one, the Jacobian $J(C)$ is isomorphic to the curve itself. A course that covers elliptic functions would probably already know this. *This is false in characteristic* p*.*

5.6 Example. In genus two, then there is a birational map $C_2 \to J(C)$: it is generically one-to-one, except on the locus of nontrivial linear systems of degree two. But there is only one nontrivial linear system of degree two, the canonical series. We have a $\mathbb{P}^1$ of divisors in the canonical series in C_2, and those get blown down to a point when projecting to $J(C)$. In other words, C_2 is a blow-up of $J(C)$ at a point.

In some sense, the thing to bring out now is not Abel's theorem itself, but the fact that this group of linear equivalence classes of divisors (of a given degree), has geometric structure. It is a *compact complex manifold.* We can talk about Pic^0 as an *algebraic variety.*

Remark. This is a very old-fashioned construction of the Jacobian, which has many drawbacks.

1. As we've defined it, the Jacobian is only a complex torus so far. Note that a general complex torus is not embeddable in projective space. However, it turns out that the Jacobian has enough meromorphic functions to embed in projective space, so it is a projective variety.

2. The construction just outlined works only over the complex numbers. It doesn't even *begin* to describe what to do in characteristic p. Even if you're over a field of characteristic zero, and you can construct the Jacobian in this way, you'll get something defined over the complex numbers. If you have a curve defined over $\mathbb{Q}$, then you probably want the Jacobian to be defined over $\mathbb{Q}$ as well.

3. This really bothered André Weil, who felt that there has to be a purely algebraic construction of the Jacobian. What he did was to go back to the birational isomorphic $C_g \to J(C)$, and argued that an open subset of $J(C)$ was isomorphic to C_g; after composing with the translates, one sees that $J(C)$ is covered by open sets isomorphic to open subsets of $J(C)$. Weil decided to reverse the process: he

took open subsets of C_g and glued them to form the Jacobian. However, the idea that one can construct a variety by gluing affine varieties had not been constructed yet: Weil *invented abstract varieties* precisely to construct the Jacobians.

4. Weil's construction is also antiquated. Grothendieck then came along, and considered $C_d \to J$ for large d, and developed a general theory of quotienting varieties by equivalence relations. This is how Grothendieck constructed the Jacobian.

We only mentioned abelian integrals because they are historically interesting. Again, the real take-away is that the space of line bundles of degree zero has an algebraic structure, and that there is this birational isomorphism $C_g \to J(C)$.

§4 Curves of genus one

Now we want to do examples of realizing curves in projective space. Last time we talked about rational curves, so now let's do genus one. Let C be a Riemann surface of degree one.

If D is a divisor of degree $d \geq 1$, then Riemann-Roch says that $h^0(\mathcal{L}(D)) = d$, or $r(D) = d - 1$. For instance, a divisor of degree two has a nontrivial element of $\mathcal{L}(D)$. It follows that we get C as a double cover of $\mathbb{P}^1$, branched (by Riemann-Hurwitz) at four points. We get the expression

$$y^2 = x(x-1)(x-\lambda)$$

to describe C. This is the familiar expression for an elliptic curve.

5.7 Example. Now a divisor of degree 3 gives an *embedding* of C as a smooth plane cubic, because as we saw earlier such divisors are *very ample.*

5.8 Example. A divisor D of degree 4 gives an embedding $C \hookrightarrow \mathbb{P}^3$ has a quartic curve. If we look at the map

$$H^0(\mathcal{O}_{\mathbb{P}^3}(2)) \to H^0(C, \mathcal{O}_C(2D)),$$

we find that the last thing has dimension 8 because D has degree four. In particular, the kernel has dimension at least two. Thus C lies on two irreducible quadrics (the quadrics can't be reducible as they would be then union of planes). We conclude:

5.9 Proposition. *An elliptic curve $C \subset \mathbb{P}^3$ is a complete intersection of two quadrics.*

The converse is easy. The complete intersection of two smooth quadrics in $\mathbb{P}^3$, by the adjunction formula, has genus one.

5.10 Example. Recall that a quadric in $\mathbb{P}^3$ is isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$, so it has two rulings of lines on it. If $C \subset \mathbb{P}^3$ is an elliptic curve, then we have two quadrics Q, Q' that C lies on. But in fact we have a pencil of quadrics given by $t_0 Q + t_1 Q'$, and we thus get a *pencil* of quadrics, each of which contains the given curve C.

Exactly four of these quadrics in the pencil will turn out to be singular, by a determinant argument. So occasionally, as this quadric varies, it will degenerate to a quadric cone.

§5 Genus two, beginnings

We want to do one more example. Up to now, when we talked about curves of genus zero and one, we considered maps to projective space coming from divisors on these curves. The geometry of the imbedding turned out to be uniform over all curves of genus zero or one and over all line bundles. In genus two, the geometry of these maps will depend on the choice of line bundle.

Let C be a curve of genus two. Let D be a divisor or divisor class of degree d. We want to look at what sorts of maps we get.

5.11 Example. When $d = 2$, then Riemann-Roch states that

$$\dim \mathcal{L}(D) = 2 \quad \text{if } D = K_C \text{ or } 1 \text{ otherwise.}$$

We saw this earlier when we said that the map $C_2 \to J(C)$ consists of a blow-up at the canonical divisor. But we do get one nontrivial linear series, which gives a two-to-one map

$$C \to \mathbb{P}^1.$$

This expresses C as a double cover of $\mathbb{P}^1$, now branched at six points. This is thus in the form $y^2 = \prod^6 (x - \lambda_i)$.

5.12 Corollary. *A curve of genus two is hyperelliptic.*

Note that this gives an *involution* of the curve that exchanges the two sheets of this cover. Any curve of genus two has a canonical involution from its canonical double cover of $\mathbb{P}^1$.

5.13 Example. Now let's consider the case $d = 3$. Riemann-Roch states that $h^0(D)$ is the same for all such divisors D, and we get a pencil of dimension two in each case. It seems that we get a family of maps $C \to \mathbb{P}^1$ of degree three. Well, not quite.

1. If $D \sim K + p$ for some $p \in C$, then p is a base-point of the linear series D. In other words, $\mathcal{L}(D) = \mathcal{L}(D - p) = \mathcal{L}(K)$. Every divisor of the linear system D is just a canonical divisor plus this fixed point. We *don't* get a map to $\mathbb{P}^1$.
2. If $D \not\sim K + p$, we do get a degree-three map $C \to \mathbb{P}^1$.

Now do we know that there is a divisor class of degree three not linearly equivalent to $K+p$? If we know there is, then we get a degree-three map to $\mathbb{P}^1$. Note that the space of divisor classes of degree three has the structure of a compact complex manifold of dimension *two*, by the theory of the Jacobian. The divisors of the form $K + p, p \in C$ fill out a subspace of dimension at most one. We find that the generic divisor of degree two is not of the form $K + p$ for some $p \in C$.

In particular:

5.14 Proposition. *A curve of genus two is a triple cover of* $\mathbb{P}^1$*.*

However, so far we haven't gotten an imbedding in projective space, even a birational one.

5.15 Example. Let's take $d = 4$, so D is a divisor of degree four. Then $\mathcal{L}(D)$ is a space of dimension three. There are no base-points (again by Riemann-Roch). Thus, for any such D, there is a map of degree four

$$\phi_D : C \to \mathbb{P}^2$$

given by the complete linear system associated to D.

There are three possibilities. Note that any such D can be written as $K+p+q$ (by using facts about the Jacobian, as before). *Proofs are not included here.*

1. $D = K + p + q$ such that p, q lie over distinct points of $\mathbb{P}^1$ via the double cover $C \to \mathbb{P}^1$ constructed earlier. Then $\phi_D : C \to \mathbb{P}^2$ is not an imbedding; it is in fact singular. In fact, $\phi_D(p) = \phi_D(q)$ because $\dim \mathcal{L}(D-p-q) = \dim \mathcal{L}(D) - 1$. We get a map ϕ_D to $\mathbb{P}^2$ which is one-to-one and an immersion *except* that p, q are collapsed. We're not going to do this out, and it may appear on the homework, but in fact $\phi_D(C) \subset \mathbb{P}^2$ is a quartic plane curve with a node. (**Prove this.**)

2. The second possibility is that $D = K + 2p$ for p is not a branch point in the two-sheeted cover. Here the map fails to be an immersion at p, and the image curve $\phi_D(C)$ has a cusp at $\phi_D(p)$.

3. If $D = K + p + q$ where p, q are conjugate under the involution, In this case, every pair of points conjugate under the hyperelliptic involution are mapped to the same point. So if $\iota : C \to C$ is the hyperelliptic involution of C, then $\phi_D(r) = \phi_D(\iota(r))$. It follows that the map $\phi_D : C \to \mathbb{P}^2$ factors as the canonical cover $C \to \mathbb{P}^1$ followed by a conic imbedding $\mathbb{P}^1 \to \mathbb{P}^2$. Here the map is not even birational.

All these examples occur. Indeed, the divisors of the form $K + 2p$ form a curve in the Jacobian, so the generic divisor of degree four gives a node.

Lecture 6
9/21

Today we are going to fix a loose end. We gave a bogus proof of Riemann-Roch earlier, and asserted without proof that the genus g (defined topologically) was the dimension of the space of holomorphic differentials.

Let us consider the following problem. Given a smooth projective curve C and a divisor D, we want to find the complete linear system $|D|$, namely the set of all effective divisors linearly equivalent to D. Equivalently, we want to write down all the rational functions with poles at these points. We'll discuss a simple way of doing this using only linear algebra.

In particular, we need to do this for the canonical bundle. In other words, we should write down the space of all holomorphic differentials.

§1 Differentials on smooth plane curves

Let's start with a relatively simple case. Let $C \subset \mathbb{P}^2$ be a *smooth plane curve,* of degree d (and genus[3] given by the usual formula $\binom{d}{2}$)). Choose an affine open in $\mathbb{P}^2$, so say a complement of a line. Let's take the line at ∞, which we'll call L_∞. We have $\mathbb{A}^2 = \mathbb{P}^2 - L_\infty$ (the usual plane before they projectivized it). We can choose affine coordinates (x, y) on this affine plane $\mathbb{A}^2 \subset \mathbb{P}^2$, and let the corresponding homogeneous coordinates on $\mathbb{P}^2$ be $\{X, Y, Z\}$. (The line L_∞ is thus given by $Z = 0$.) We're going to single out the vertical point at ∞, which is $[0, 1, 0]$.

We now want to make a couple of assumptions about the coordinate plane:

1. L_∞ intersects the curve C transversely, in d distinct points.

2. C does not contain the vertical point at ∞, $[0, 1, 0]$.

We can achieve this by rotating the axes.

In $\mathbb{A}^2$, let us say that the curve is the zero locus of an inhomogeneous polynomial f in x, y (so the projective curve is the zero locus of the homogeneization). Let's now write *one* holomorphic differential on the curve. The standard way to do this is to pick a *meromorphic* differential, and then to kill the poles by multiplying by a meromorphic function. So, consider the meromorphic differential dx: this is a perfectly good meromorphic differential, which is even *holomorphic* in the finite plane (but not at ∞). To say this in another way, if we *project* from the point $[0, 1, 0] \notin C$ (in the affine picture: just the vertical projection, or projection to the x-coordinate), we get a map

$$\pi : C \to \mathbb{P}^1_x$$

which expresses C as a d-sheeted cover of $\mathbb{P}^1$, unramified at ∞ (because we assumed there were d points in L_∞). The differential dx on C is the pull-back of the differential dx on $\mathbb{P}^1$ (which is meromorphic).

So, there's our differential: dx. What's its divisor? In the finite plane, there are no poles; where are the zeros? The zeros are the points where π is ramified. In fact, this is equivalently where $\frac{\partial f}{\partial y} = 0$. To see this, note that

$$\left(\frac{\partial f}{\partial x} dx + \frac{\partial f}{\partial y} dy \right) |_C = 0,$$

because it is df and $f = 0$ on C. Consequently, if $\frac{\partial f}{\partial y}$ vanishes at a point, *smoothness* of the curve (that is, the partials of f have no common zeros), dx must vanish on C to the *same order* as $\frac{\partial f}{\partial y}$.

Consequently:

> In the finite plane, dx is holomorphic, and its divisor is the divisor of the function $\frac{\partial f}{\partial y}|_C$.

[3] Recall that we did this by computing the degree of the canonical divisor in terms of the topological genus, by the Poincaré-Hopf theorem (for instance). In this case, we showed that the canonical bundle K_C was $(d-3)L$ for L a line in $\mathbb{P}^2$: the degree was thus $d(d-3)$, which was $2g-2$, and let us compute g.

All right. What about at ∞? Again, we have $C \to \mathbb{P}^1$ as a d-sheeted cover. Well, the claim is that dx has a pole of order 2 at ∞: this is because dx on $\mathbb{P}^1$ has a pole of order 2 at ∞ (easy to check), and π is unramified above the points at ∞. So, if D_∞ is the intersection $L_\infty \cap C$, the points at ∞, then:

The differential dx has double poles at each point in D_∞.

In particular, we find:

$$\operatorname{div}(dx) = \operatorname{div}(\frac{\partial f}{\partial y})|_{\mathbb{A}^2} - 2D_\infty.$$

Now, we have to divide by a function with poles at ∞ to get rid of the poles of dx. Consider a polynomial of degree 2 in x, y: there will be a pole at ∞. If we divide by this, we can get rid of the poles at ∞ of dx. Of course, we have to be careful: this division will introduce *new poles* where this polynomial by which we divide vanishes. I.e., we're looking for a holomorphic differential of the form dx/h where $h \in \mathbb{C}[x, y]$ has degree at least two—so has no poles at ∞—but we have to worry about the zeros of P. If we can arrange the zeros of P so they occur at pre-existing zeros of dx, then we won't have a problem: the zeros will just cancel.

Let's consider $P(x, y) = \frac{\partial f}{\partial y}$; let's assume this is degree ≥ 2. Given the expression for $\operatorname{div}(dx)$, we see that the zeros of P are precisely at the zeros of dx. Consequently, we have obtained a holomorphic differential on the curve:

6.1 Proposition. *Suppose the degree is at least* 3. *Then* $\omega = \frac{dx}{\frac{\partial f}{\partial y}}$ *is a holomorphic differential on the curve* C. *Moreover,* ω *is nonzero on the affine part of* C*, i.e.* $\mathbb{A}^2 \cap C$. *There will be zeros of order* $d - 3$ *along the divisor at* ∞.

Note that dividing by P precisely eliminated the zeros in the affine part. The divisor at ∞ was obtained by considering the pole of $P = \frac{\partial f}{\partial y}$ at ∞. Note also that since the differential $\operatorname{div}(\omega)$ is $(d-3)D_\infty$, we have obtained again that the canonical divisor is $d(d-3)$, as expected by the Riemann-Roch theorem and the genus formula.

How can we write down other holomorphic differentials? We don't have to worry about poles except at ∞; we just have to multiply ω as above by *any* rational function with poles at most $d-3$ at ∞. These things are known as polynomials of degree at most $d-3$. In other words, if $g(x, y)$ is a polynomial of degree at most $d-3$, then $g\omega$ is a global holomorphic differential. This space of $\{g\omega\}$ is a *vector space* of differentials, and the dimension is $\binom{d-1}{2}$. As a result, we have expected a g-dimensional vector space of holomorphic differentials, and that's all.

6.2 Corollary. *The holomorphic differentials on* C *are of the form* $g\omega$ *for* ω *as above and* $g \in \mathbb{C}[x, y]$ *a polynomial of degree* $\leq d - 3$.

We've written down all the differentials on a plane curve. This is a special case: *most* curves are not smooth plane curves. The second case that we will later consider is the normalization of a nodal plane curve. Since any curve can be realized in this way, we'll be done.

§2 The more general problem

Let C be as before, in the previous subsection, a plane curve of degree d. Let D be now an effective divisor on C: we want to write down $\mathcal{L}(D)$, or equivalently $|D|$. We did this earlier for D the canonical divisor. Let's now assume D doesn't intersect $D_\infty = L_\infty \cap C$.

If we want to find functions allowed to have poles at D, then we should think about ratios of polynomials—rational functions in the plane—where the denominator is allowed to vanish on D. That's cool, but having written down a denominator, it might vanish at other points too. The starting point, anyway, is to choose a polynomial $g(x, y)$ (not the same g as we had a moment ago!) in the plane, of some degree m, vanishing at the points of D. Let $G = V(g) \subset \mathbb{P}^2$ be the degree m curve that we require to contain the subscheme D. We can think of g as the *denominator* of the rational function. We've got the poles at the points of D, but we also have illegitimate poles coming from the points where g vanishes outside of D. So we have to multiply by a function that vanishes at the points of $V(g)$ not in C.

So, let's write $G.C = D + E$. Here E consists of the points where $\frac{1}{g}$ has poles, in addition to its already chosen poles on D (which are legitimate). We need to multiply by a function $h = h(x, y)$ of the same degree m (to cancel the poles at ∞) that has *zeros* at the points of E. In other words, we want to look for plane curves $H = V(h)$ in $\mathbb{P}^2$ of degree m containing E. If we do this, then we can take h/g as the desired rational function.

Suppose $H.C = E + F$ for another divisor F. Then $H.C$ is the intersection $H \cap C$, linearly equivalent to mD_∞ and thus to $G.C = D + E$. We find that $F \sim D$. In other words, $\operatorname{div}(h/g)$ is linearly equivalent to zero: it has no poles or divisors at ∞, and the remaining divisor is $F - E$. So $F \sim E$.

6.3 Proposition. *Any element of $\mathcal{L}(D)$ can be obtained in this form. In other words, if we fix $g(x, y)$, and consider the set of all divisors of the form* $\operatorname{div}(h/g) + D$ *where $h \in \mathbb{C}[x, y]$ has degree m and vanishes at E, then this is exactly the complete linear system $\mathcal{L}(D)$.*

It seems too easy to be true, but it is true. You saw this in some of the homework problems (batch A in ACGH).

6.4 Example. Those of you interested in elliptic curves (plane cubics) should also note that the addition law on an elliptic curve can be described efficiently using divisors. Recall that if C is a plane cubic, then to add points $p, q \in C$, one draws the line through p, q and considers the third point of intersection $r' \in C$. Then, one draws the line through ∞ and r', and finds the third point $r \in C$ of the intersection. Then one has $p + q = r$. One has, from this construction, that as divisors

$$p + q = r + \infty.$$

This is because each line corresponds to a homogeneous polynomial of degree one, and the quotient of any such will be a rational function.

Anyway, this construction is thus not completely unfamiliar to you: however, the claim is that it works in *complete generality.* It's not too hard to prove, and we'll defer the proof until we've done something more general (not just for smooth plane curves).

6.5 Example. It's instructive to consider what this description tells you about the dimension. Say D is a divisor, and $\deg D = n$. Let's run through the construction and try to conclude something, which will be crucial. It follows that if E is as above, then $\deg E = \deg(G.C) - D = md - n$, where $\deg G = m$. To find the dimension $\dim \mathcal{L}(D)$, we just have to look at polynomials h of degree m that vanish at E. So there are $\binom{m+2}{2}$ linearly independent polynomials, and $\binom{m-d+2}{2}$ polynomials that vanish on C. We have

$$\binom{m+2}{2} - \binom{m-d+2}{2}$$

for the total number of polynomials, and when we restrict to vanish on E, we impose $md - n$ conditions. As a result,

$$\dim \mathcal{L}(D) \geq \binom{m+2}{2} - \binom{m-d+2}{2} - (md - n).$$

What happens when you multiply this out? There'll be quadratic terms that cancel. Asymptotically, you can find

$$\dim \mathcal{L}(D) \geq n - \binom{d-1}{2} + 1.$$

That's *exactly* what Riemann-Roch predicts. This is *really important.*

Remark. The original notion of genus was not the one we think of. When Riemann introduced the notion of Riemann surface, and topologists classified compact surfaces, they got the modern notion. But the $\binom{d-1}{2}$ is what people initially thought of. The classical geometers observed that if you have any set of $2m$ points on a conic, then any m will be linearly equivalent to the other m, but that on a plane cubic this isn't a case: given $3m$ points in a plane cubic, there is *one condition* that they be the intersection of the cubic with a degree m curve. In general, on a genus g surface, there should be g conditions: this was considered a *deficiency* of the curve (and this was the original terminology).

Remark. We assumed that D was effective earlier. That was for simplicity of notation and language. If $D = D' - D''$ where D', D'' are effective, then we can do the same thing: just take a curve $G = V(g) \supset D'$, and write $G.C = D' + E$, and choose an equal-degree curve $H = V(h)$ such that $H \supset E + D''$. If we write $H.C = E + D'' + F$, what we see is that $F \sim D' - D''$. In other words, we should allow poles on D', but we want zeros on D'' as well as removing the extra poles from g. Again, we can get all elements of $\mathcal{L}(D)$ as h/g.

§3 Differentials on general curves

We now want to consider a generalization of the previous case. Consider a plane curve C with *nodes.* In this case, we're mostly talking about smooth curves, so *we're implicitly talking about the normalization.* Alternatively, the set-up is this: consider a smooth projective curve C of genus g, and consider a *birational map* $\nu : C \to C_0 \subset \mathbb{P}^2$ where C_0 has degree d and has only nodes. Essentially, we're going to try to solve the problem posed earlier—i.e., writing down holomorphic differentials and more generally rational functions—using this plane model C_0.

The main point is that, given any smooth abstract projective curve, there is *always* such a birational map to $\mathbb{P}^2$. One can always embed in $\mathbb{P}^3$, and then a general projection $\mathbb{P}^3 \to \mathbb{P}^2$ will realize the curve as birational to something in $\mathbb{P}^2$ with only nodes. (This will be on the homework.)

All right. Let's start by trying to find all holomorphic differentials on C. Assume first that at each node, there are no vertical tangents (we can use a general automorphism of $\mathbb{P}^2$ to arrange this). We do *exactly* the same thing: consider a meromorphic differential, say dx—or really, to be technical, $\nu^* dx$ on C. Last time, we said that we had the basic relation

$$\frac{\partial f}{\partial x} dx + \frac{\partial f}{\partial y} dy = 0$$

on C, and the two partials have no common zeros in the finite plane. So we concluded that the order of dx at a point in the plane was the order of $\frac{\partial f}{\partial y}$. This is *no longer true.* Outside of a node, this is still valid (the partials have no common zeros), but at the nodes, this fails.

At a node, dx won't vanish, because of what we've assumed about the lack of vertical tangents at nodes, but the argument doesn't work. Let $r_1, \dots, r_\delta$ be the nodes on C_0. On the curve C, there are two points p_i, q_i that map to r_i. If we consider the differential $\nu^*(dx/\frac{\partial f}{\partial y})$, the general hypothesis we made shows that this has no poles in the finite plane *except at the nodes.* In fact, if $\Delta = \sum (p_i + q_i)$, then

$$\operatorname{div}(\nu^*(dx/\frac{\partial f}{\partial y}) = (d-3)D_\infty - \Delta.$$

So, given the nodes, we still have to cancel the poles corresponding to the nodes. To cancel the poles, we're going to multiply this by a polynomial $g(x, y)$ in the plane of degree $\leq d - 3$, and such that $g(r_i) = 0$ for all i. Then $\nu^*(g dx/\frac{\partial f}{\partial y}$ will be a global holomorphic differential. How many have we got? We have $\binom{d-1}{2}$ total polynomials, and δ conditions, so then at least $\binom{d-1}{2} - \delta$ linearly independent differentials. But we know (Plücker formula!) that this is the genus of C, so these are all the differentials.

So, we have solved the problem of finding differentials on a curve.

Remark. The Plücker formula for the genus of g, as $\binom{d-1}{2} - \delta$, can be obtained by applying Riemann-Hurwitz to the projection $\pi : C \to \mathbb{P}^1$ given by the x-coordinate (the nodes reduce the number of branch points because they increase the fibers). Alternatively, the Plücker formula can be obtained by computing the degree of the *meromorphic differential* obtained above.

Remark. We saw that the $\{r_i\}$ imposed *precisely* δ conditions on the polynomials of degree $\leq d - 3$: in particular, these conditions on polynomials are independent. This follows because if there were fewer conditions, we would get more differentials on the curve than we should (by the Plücker formula for the genus). This is important in the deformation theory of plane curves.

§4 Finding $\mathcal{L}(D)$ on a general curve

We're now going to try to find $\mathcal{L}(D)$ on a curve C projecting birationally via $\nu : C \to \mathbb{P}^2$ to a nodal curve $C_0 \subset \mathbb{P}^2$ with δ nodes $r_1, \dots, r_\delta$, as before in the previous section. Consider an effective divisor D on C, and for the moment let's assume its *support* is disjoint from Δ, the nodal divisor.

As before, we want to consider polynomials $g(x, y)$ of degree m, which *vanishes* on the (effective) divisor D. In other words, the plane curve $G = V(g)$ should be such that $G.C = D + \ldots$. Here's the wrinkle: we also want to require that g vanish at the nodes r_i, or that G pass through them. Then the divisor of $\nu^* g$ is going to be $D + \Delta + E + mD_\infty$, where E is the residual divisor.

Note, as a generalization of the adjunction formula, we have

$$K_C = (d-3)D_\infty - \Delta.$$

This follows from the divisor dx.

OK, let h be a polynomial of degree m, vanishing on E and also vanishing on the $\{r_i\}$. We have that $\operatorname{div} \nu^*(h) = E + F + \Delta - mD_\infty$. Consequently, h/g will have divisor $\geq -D$ (in fact, it will be $F - D$). The claim is that *all* rational functions in $\mathcal{L}(D)$ can be obtained in this way. Alternatively:

6.6 Proposition. *$\mathcal{L}(D)$ is the space of rational functions h/g where g is fixed and h is any polynomial such that $h(E) = h(r_i) = 0$.*

The trick is to make both h, g vanish at the r_i. In fact, doing this is going to give us *more functions*, which is what we should expect: nodes make the genus go down. Note that $\deg E = md - n - 2\delta$ (where n is $\deg D$), because the nodes are removed. Consequently, if you consider $\dim \mathcal{L}(D)$, you get

$$\binom{m+2}{2} - \binom{m-d+2}{2} - (md - n - 2\delta) - \delta$$

where the last condition is for the δ conditions that h vanish at those points. So, there's an extra δ here, which makes it different from what there was on the board twenty minutes before. By requiring g to vanish at the r_i, the residual divisor *drops* by two. However, only one condition is imposed on h for each r_i. Consequently, we get *more* rational functions in this way. One gets for $D \gg 0$,

$$\ell(D) \geq n - (\binom{d-1}{2} - \delta) + 1.$$

The presence of nodes *decreases* the deficiency.

Remark. We can do the same splitting $D = D' - D''$ when D is not assumed effective a priori.

Now, you can have some confidence that when you write down a curve in the plane, or even a nodal curve in the plane, you can find *explicitly* the rational functions with a specified polar divisor. We'll use this, probably next time, not today, for proving Riemann-Roch.

The last thing to do before calling it quits is to consider *arbitrary* singularities. Every curve admits a birational map to a nodal plane curve, so we've really handled arbitrary projective curves, but in practice, one wants to deal with curves *as they come to you.* So, practically, it is useful to know what to do—as above—even when there are wilder singularities. That is, there is a birational map $\nu : C \to C_0 \subset \mathbb{P}^2$ where C_0 might have arbitrary singularities. You can figure out the contribution of each such singularity in C_0 to the genus, by projecting to a line and using Riemann-Hurwitz. It's out there, and you can read about it in the exercises in ACGH. But anyway, one can find the complete linear series one wants—in particular, the space of differentials—by replacing the condition above $g(r_i) = 0$ to what was classically called "$g(x, y)$ satisfies the adjoint conditions." The modern way of saying this is that $g(x, y)$ is in the "conductor ideal," i.e. is in the annihilator of $(\nu_* \mathcal{O}_C)/\mathcal{O}_{C_0}$.

Lecture 7
9/23

§1 More on $\mathcal{L}(D)$

Last time, the claim was that there was an algorithm for finding the space $\mathcal{L}(D)$ associated to a divisor, and in particular obtaining the space of holomorphic differentials.

The idea was to map this birationally to the plane. Let C be a smooth projective curve of genus g. Let $\nu : C \to C_0 \subset \mathbb{P}^2$ be a birational morphism, where $C_0 = V(f)$ is the zero locus of a homogeneous polynomial f in three variables. Suppose there are δ nodes $r_1, \dots, r_\delta$ of C_0, and *no other singularities.* We can always do this. We saw that δ here is the difference between the arithmetic genus $\binom{d-1}{2}$ and the actual genus g; this follows from the Plücker formulas, or follows from writing down meromorphic differentials.

The strategy, for our problem, is to consider the differential $\omega = \frac{g dx}{f_y}$ where g is a polynomial of degree $m \leq d-3$ vanishing at the nodes, $g(r_i) = 0$ for each i. These will be *holomorphic* differentials on the curve.

As we saw:

7.1 Proposition. *All the holomorphic differentials on C are of this form.*

Proof. We observed that the dimension of this vector space of ω constructed in this way is of dimension at least $\binom{d-1}{2} - \delta$, because the space of such polynomials g is of dimension $\binom{d-1}{2}$ and we have imposed δ linear conditions. Consequently, we have at least g holomorphic differentials, and this must give all of them. ▲

Remark. The fact that the dimension of $H^0(K)$ is at most g can be seen in an elementary way as follows. Note that any holomorphic differential is closed, but it *can't* be exact: if it were, it would be the differential of a holomorphic function, and those are constant. So $H^0(K)$ *injects* into the first de Rham cohomology group $H^1(C, \mathbb{C})$, which is a dimension $2g$ on a Riemann surface of genus g. By the same token, we have the *antiholomorphic* differentials, which include in $H^1(C, \mathbb{C})$ as a subspace of dimension g, and the claim is that

$$H^0(K) \cap \overline{H^0(K)} = 0.$$

Here $\overline{H^0(K)}$ consists of the conjugates of the subspace $H^0(K)$, or equivalently the antiholomorphic differentials. Once we've proved it, we'll be done: $\dim H^0(K) = \dim \overline{H^0(K)} \leq g$.

But we have a *positive definite* hermitian form on $H^0(K)$, which is given by $\omega \mapsto i\int_C \omega \wedge \overline{\omega}$. Call this $H(\omega)$. The positive definiteness follows because locally $\omega = f(z)dz$ and $\overline{\omega} = \overline{f(z)}d\overline{z}$, so

$$\omega \wedge \overline{\omega} = |f(z)|^2 (dx + idy) \wedge (dx - idy) = -i|f(z)|^2 dx \wedge dy.$$

Anyway, if we had a form ω which was both holmorphic and *homologous* to an antiholomorphic form, then the integral of $\omega \wedge \overline{\omega}$ would be zero (this is because the integral $\int \eta \wedge \eta'$ depends only on the de Rham cohomology classes of η, η').

It follows that if one starts from first principles, then we can determine that the dimension of $H^0(K)$ is g.

§2 Riemann-Roch

Let now D be an effective divisor, say of degree n. We can choose a polynomial $g(x, y)$ of degree m and vanishes on D and also vanishes at the nodes, $g(r_i) = 0$. We can write $(g) = D + \Delta + E$ for Δ the sum of the points lying over the nodes, and E some residual divisor. We then look at functions of the form h/g, where h is a polynomial of the same degree or less, such that h vanishes on E and all the nodes. Once we've fixed g and let h vary, these ratios h/g are elements of $\mathcal{L}(D)$ because the poles along E, Δ cancel.

7.2 Proposition. *All of $\mathcal{L}(D)$ can be obtained in this way, as h/g (where g is fixed).*

Note also the dimension count. We have $\deg E = md - 2\delta - n$, by counting. The dimension of the space of $\{h/g\}$ thus constructed is

$$\geq \binom{m+2}{2} - \binom{m-d+2}{2} - \delta - (md - 2\delta - n) = n - g + 1,$$

because we have counted the number of polynomials, and counted the number of relations imposed. (We have required that $h(r_i) = 0$ and $h(E) = 0$.)

Let's prove that this construction gives all the functions, and also prove Riemann-Roch:

7.3 Theorem (Riemann-Roch)**.** *If $D = p_1 + \cdots + p_d$ is an effective divisor, then*

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

Let's recall how we "proved" this earlier. Let ω be a holomorphic differential on C, and let $f \in \mathcal{L}(D)$. We observed that $\sum \operatorname{res}_{p_i}(f\omega) = 0$ where $f\omega$ is a meromorphic differential on C.

Note that an element of $\mathcal{L}(D)$ is determined up to a constant by the polar parts at each p_i, and consequently the space of polar parts plus constants has dimension $\deg D + 1$. However, we get a condition on the polar parts for each $\omega \in H^0(K)$, which gives us g conditions. Nonetheless, these conditions are not linearly independent, because whenever a differential vanishes identically on the p_i the relation is trivial. We get that there are at least $g - \ell(K - D)$ relations imposed on the polar parts for it to be a meromorphic function, so

$$\ell(D) \leq d + 1 - (g - \ell(K - D)).$$

Now the argument went downhill; we applied the above inequality to $K - D$, and added the two to deduce Riemann-Roch. But to apply this argument properly, we would have to assume that $K - D$ was effective. If $K - D$ is effective, then we get a reasonable proof of Riemann-Roch. If, however, $\ell(K - D) = 0$, then this argument doesn't work, so we need to prove for a "non-special divisor"

$$\ell(D) = d + 1 - g.$$

We only need to prove

$$\ell(D) \geq d + 1 - g.$$

However, we did this on Wednesday. We *wrote down* a vector space of rational functions of dimension $d + 1 - g$. So, we've done it, and it follows that the construction made on Wednesday *gives us all of* $\mathcal{L}(D)$. In particular, we've proved Riemann-Roch at least if one of $D, K - D$ is effective.

If $D, K - D$ are both non-effective, then we need to claim that $\deg D, \deg(K - D) < g$. Every divisor class of degree g is represented by an effective divisor: this is Jacobi inversion, so we agree. Since $\deg K = 2g - 2$, the only alternative is that $\deg D = \deg K - D = g - 1$, and the Riemann-Roch formula is clear in this case.

We didn't completely prove that this construction of the elements of $\mathcal{L}(D)$ works in the non-effective case.

Remark. A long time ago, people didn't talk about differential forms. A given Riemann surface can be realized as a projective curve in many different ways. There are various representations of a Riemann surface with different degrees, different numbers of nodes. People noticed that the divisor cut by a hyperplane could vary in different models, but the divisor cut by an *adjoint* of degree $d - 3$ was the *same* for any projective model. No matter what you did to this curve, there was one constant divisor class. That's what led to the whole notion of a canonical series, and to the genus.

The *adjoint series* on a plane curve consists of the plane curves of a given degree m that pass through the nodes. It's a hard exercise to show that the adjoint series is complete.

We now give some examples.

7.4 Example. Consider a smooth conic C in $\mathbb{P}^2$, which is rational: it's isomorphic to $\mathbb{P}^1$. To do this, pick any point $p \in C$ and project from it. Another way of saying this is to consider the linear system of degree one and dimension one that are cut out by the lines through that point (not including the point p), and this moving family of divisors gives a map of degree one to $\mathbb{P}^1$.

7.5 Example. A nodal cubic is birational to $\mathbb{P}^1$. Namely, pick the node, and draw lines through it intersecting the curve at other points. Each of those lines intersects the curve in one other point (because of the node), and that defines a map from $\mathbb{P}^1$ to the nodal cubic.

7.6 Example. A plane quartic with three nodes is birational to $\mathbb{P}^1$. This follows from the genus formula written down earlier. How can we see this? Well, if it's rational, then the divisor consisting of a single point has to move in a pencil, and that should give the meromorphic function of degree one. How do we do that? We have to choose a curve of some degree m that vanishes at one point and at the three nodes. So, let G be a conic plane curve containing the starting point p and the three nodes r_1, r_2, r_3. By Bezout, the intersection has eight points, but we've accounted for seven; so we get a new point q which is the remaining point of intersection. This family of conics vanishing on p, r_1, r_2, r_3 is a $\mathbb{P}^1$ of conics, and each of them meets the curve exactly one more time. So we get a birational map from $\mathbb{P}^1$ to this plane quartic.

The key point is that each node r_i imposes one condition on the conic, but absorbs two intersection points: this is the "two-for-one deal." Classically, this is why the nodes drop the genus.

In practice, whenever you see a plane curve and the genus formula leads you to conclude that it's rational, you can *write down* the rational parametrization very explicitly.

§3 Sheaf cohomology

Now we'll review a certain language we have already been using.

Consider a set $\Gamma = \{p_1, \dots, p_\delta\} \subset \mathbb{P}^2$. We've already been working with family of plane curves passing through a family of points, or equivalently homogeneous polynomials that vanish at this family of points. In other words, one is interested in the space of homogeneous polynomials of degree m in three variables vanishing at $p_1, \dots, p_\delta$, which sits inside $H^0(\mathbb{P}^2, \mathcal{O}(m))$. That is,

$$H^0(\mathcal{I}_\Gamma(m)) \hookrightarrow H^0(\mathcal{O}_{\mathbb{P}^2}(m)).$$

The condition that a polynomial vanish on a point is one linear condition, but when are those conditions *independent*? This is hard.

7.7 Definition. We say that Γ **imposes independent conditions** on curves of degree m if the codimension of $H^0(\mathcal{I}_\Gamma(m)) \subset H^0(\mathcal{O}_{\mathbb{P}^2}(m))$ is equal to δ (it is at most δ).

In other words, given any collection of complex numbers, you can interpolate at those points. In modern language, you want the surjective sheaf map $\mathcal{O}_{\mathbb{P}^2}(m) \to \mathcal{O}_\Gamma(m)$ to be surjective *on global sections.* This is one of the ways in which the language of sheaves and sheaf cohomology emerges from simply an alternative way of writing down the same objects. When you express it in this way, it amounts to the same thing. Sheaf cohomology was *designed* to answer the question of when a surjective sheaf map is surjective on global sections. Well, there is an exact sequence

$$0 \to \mathcal{I}_\Gamma(m) \to \mathcal{O}_{\mathbb{P}^2}(m) \to \mathcal{O}_\Gamma(m) \to 0,$$

and consequently there is a long exact sequence in cohomology

$$0 \to H^0(\mathcal{I}_\Gamma(m)) \to H^0(\mathcal{O}_{\mathbb{P}^2}(m)) \to H^0(\mathcal{O}_\Gamma(m)) \to H^1(\mathcal{I}_\Gamma(m)) \to 0,$$

because any intermediate cohomology group of a line bundle on projective space is zero. Consequently:

7.8 Proposition. Γ *imposes independent conditions on curves of degree* m *if and only if* $h^1(\mathcal{I}_\Gamma(m)) = 0$.

What does this do for us? Absolutely nothing. It just helps to observe that this same group $H^1(\mathcal{I}_\Gamma(m))$ will occur in other contexts as well.

7.9 Example. Here's a problem. You've seen that you can realize a genus one curve as a quartic space curve (the intersections of two quadrics). Also, a genus zero curve can be embedded in $\mathbb{P}^3$ as a quadric, as a rational normal curve.

Let C be a quartic space curve: it turns out it will be either genus zero or one. How do we detect which one? Note that $C \cap H$ for a hyperplane $H = \mathbb{P}^2 \subset \mathbb{P}^3$ will look the same either way: it will be four points. The question is: there are two conics in H vanishing on Γ. Are these restrictions of quadrics in $\mathbb{P}^3$ containing C? In other words, when can a conic curve in the plane be lifted to a quadric in $\mathbb{P}^3$? Whenever you ask whether a restriction of something comes from something larger, that's what sheaf cohomology is there to answer, and it tells us what the genus must be.

Recall that the space of quadratics vanishing on Γ is just $\mathcal{I}_C(2)$, and we have a surjection

$$\mathcal{I}_C(2) \to \mathcal{I}_\Gamma(2) \to 0,$$

and the kernel is going to consist of polynomials of degree two vanishing on the hyperplane, so we get an exact sequence

$$0 \to \mathcal{I}_C(1) \to \mathcal{I}_C(2) \to \mathcal{I}_\Gamma(2) \to 0,$$

and we wonder if the last map is surjective on global sections. Now, when we look at H^1 of the kernel, we find what we saw earlier in a different context.

Anyway, we will have problem sets on sheaf cohomology, not as a way of calculating things, but as a way of *transferring* information between contexts.

Here is a basic theorem.

7.10 Theorem. *If* $\delta \leq m+1$*, then* any *configuration* Γ *of degree* δ *imposes independent conditions of polynomials of degree* m*. If* $\delta = m+2$*, then* Γ *fails to impose independent conditions if and only if* Γ *is contained in a line.*

Proof. Given $m+1$ points, they impose independent conditions: that is, there is a polynomial of degree m that vanishes at all but one point, and not at that one point. How can we exhibit a polynomial that vanishes at m points but not at a different one? Choose m lines that pass through the m points that don't pass through the remaining one; we get a degree m curve passing through the m points and not through the $m+1$st one.

When we take $\delta = m+2$, we want to cover $m+1$ points with m lines, and well, we'll leave the details to the reader. ▲

Let's apply this language to plane curves.

7.11 Example. Let $C \subset \mathbb{P}^2$ be a smooth plane curve of degree d. We'd like to ask the question: what is the smallest degree of a nonconstant meromorphic function on C?

Equivalently, what is the smallest n such that there exists a map $C \to \mathbb{P}^1$ expressing C as a branched cover of the sphere of degree n?

The coordinate functions in $\mathbb{P}^2$ (or rather, ratios of them), give rational functions $C \to \mathbb{P}^1$ of degree d. We can also get a rational function $C \to \mathbb{P}^1$ of degree $d-1$: draw two lines in $\mathbb{P}^2$ through a point $p \in C$, cut out by equations ℓ, ℓ'; the ratio of the defining equations gives a rational function of degree $d-1$ to $\mathbb{P}^1$ (because the p cancels).

Let $D = p_1 + \cdots + p_m$ be any divisor of degree m. How can we tell when this moves in a pencil? We can use the geometric form of Riemann-Roch: if we look at the canonical model of a (non-hyperelliptic curve), then a bunch of points move in a pencil if and only if they are linearly dependent. In other words, they should fail to impose independent conditions on the canonical series. That is, $r(D) \geq 1$ if and only if $p_1, \dots, p_m$ fail to impose independent conditions on the canonical series on the curve is. By what we said, though, this is just the restriction of the linear system of degree $d-3$ plane curves. That is, $H^0(K) = \frac{g dx}{f_y}$ with g a polynomial of degree $\leq d-3$. In other words, $r(D) \geq 1$ if and only if $p_1, \dots, p_m$ fail to impose independent conditions on polynomials of degree $d-3$ in $\mathbb{P}^3$.

But we've just seen that any $d-2$ points impose independent conditions on the linear system of curves of degree $d-3$. Consequently, if D moves in a pencil, then $m \geq d-1$, and that can only happen $p_1, \dots, p_m$ are collinear. As a result, the best we can do is a degree $d-1$-map:

7.12 Proposition. *There are no nonconstant meromorphic functions on a smooth plane curve of degree d which have degree $\leq d-2$.*

Let's do one more example today.

7.13 Example. Let's try to do this for a plane curve $C_0 \subset \mathbb{P}^2$ with one node, again of degree d. We can ask the same question, again not really for the plane curve but for its normalization C. What is the smallest degree of a nonconstant meromorphic function to $\mathbb{P}^1$? We're looking for divisors $D = p_1 + \cdots + p_m$ that fail to impose independent conditions on the canonical series. The canonical series of the curve is not simply cut by polynomials of degree $d-3$, but rather is cut by polynomials of degree $d-3$ that vanish at the node r. In other words, we want that $\{p_1, \dots, p_m, r\}$ fail to impose independent conditions on degree $d-3$ curves. The answer is pretty much what you would expect: this can happen only when $m+1 \geq d-2$, or $m \geq d-3$.

Typically, when you take a random curve and jam it into the plane, you'll see more than one node, so this isn't necessarily helpful.

Lecture 8
9/28

Today, we want to get back to talking about examples of curves. We'd like to go through curves of low genus, and see how they behave and what we can say about them. In other words, what makes some curves *different* from others, which is a fascinating topic.

Next week, we'll start on one of the main topics of the course, which is Castelnuovo theory. Those of you who want to read ahead can look at chapter 3 of ACGH.

Before last week, we talked about curves of genus zero, one, and two and their various imbeddings in projective space. Here's the thing: if we start with a curve C of genus g, and let D be a divisor class on C (or, a line bundle) of degree d, then we get a map of degree d (if there are no base-points), $\phi : C \to \mathbb{P}^r$. We can talk about the geometry of the image curve of C. Note however that in the case of genus zero or one, the behavior of this map ϕ depends *only* on the degree d. In genus two, however, the behavior of this map *depends* on the choice of line bundle. If we consider $g = 2, d = 4$ we get a map $\phi : C \to \mathbb{P}^2$ which is either birational onto a nodal curve, or birational to a curve with a cusp, or two-to-one onto a conic. All three are possible, and correspond to different choices of a line bundle.

1. If $D = 2K$, then the map is two-to-one onto a conic.
2. If $D = K + p + q$, then we get a birational map.

It's the same picture for *every* curve of genus two. That is, for curves with $g = 2$, the answer depends on the *line bundle,* but not on the curve. For higher genera, we'll see that the answer depends on both, in general.

§1 Divisors for $g = 3$; hyperelliptic curves

Consider $d = 4$, i.e. divisors of degree four, and $g = 3$. We have, staring us right in the face, a canonical map, no choices required. This is a map

$$\phi_K : C \to \mathbb{P}^2.$$

If this is an embedding, then ϕ_K realizes C as a smooth projective quartic curve. However, this map may not be an embedding; it is an imbedding *precisely* if C is not hyperelliptic, or that it is not a two-to-one cover of $\mathbb{P}^1$. We don't have too much to say here as a result, but we might as well talk about hyperelliptic curves. This will work in any genus.

Let C be a curve of *any genus*, say g. We say that C is **hyperelliptic** if it is a branched cover of degree two of $\mathbb{P}^1$, and consequently it is given by the equation $y^2 = \prod_{2g+2}(x - \lambda_i)$. This expresses C as a cover of $\mathbb{P}^1$ branched at the points λ_i, and with two points at ∞ (and no ramification there). The first thing we want to do for such a curve is to write down the holomorphic differentials.

8.1 Example (Differentials on a hyperelliptic curve)**.** For the curve C given by $y^2 = \prod^{2g+2}(x - \lambda_i)$, we can find a meromorphic differential easy: dx. This is nice; if we have a projective curve, as usual, we have an automatic way to generate differentials. However, if you consider a Riemann surface by the usual way via coordinate charts, it's much less obvious that these differentials exist. Let $r_1, \dots, r_{2g+2}$ be the points on C lying over the branch points of $C \to \mathbb{P}^1$ (so the points lying over the λ_i). The differential dx has zeros at the $\{r_i\}$ and nowhere else. At ∞, we find that dx has a double pole at the two points $p, q \in C$ lying over C (since x has a simple pole there).

I.e., $\mathrm{div}(dx) = \sum r_i - p - q$. To make this a holomorphic differential, we multiply by a rational function to kill the poles at ∞, without introducing new poles.

We can consider dx/y: there are no poles and no zeros in the finite plane, because those of dx cancel those of y. At ∞, the point is, there has to be zeros because the *degree* of the divisor of dx/y is $2g-2$. So there has to be a zero of order $g-1$ at each of p, q (by symmetry). So we get a nontrivial holomorphic differential, and we can get more via

$$x^i \frac{dx}{y}, \quad 0 \le i \le g-1.$$

These form a *basis* for the g-dimensional space of holomorphic differentials.

This also tells us what the canonical map $C \to \mathbb{P}^{g-1}$ looks like; it is

$$(x, y) \mapsto [1, x, x^2, \dots, x^{g-1}] \in \mathbb{P}^{g-1}.$$

So the canonical map of a hyperelliptic curve (with, say $g \ge 2$) is 2-to-1 onto a rational normal curve in $\mathbb{P}^{g-1}$.

Note that a hyperelliptic curve comes with a divisor class that moves in a pencil, coming from the meromorphic function to $\mathbb{P}^1$ which is 2-to-one (i.e. the pull-back of a point). As a result, a hyperellliptic curve has a g^1_2, given by $\pi^{-1}(r)$ for $r \in \mathbb{P}^1$ and $\pi : C \to \mathbb{P}^1$ the meromorphic function. The canonical class, from the above analysis of dx, is given by $(g-1)g^1_2$ for this (unique) g^1_2. Now, if $p, q \in C$ are *any* points on C, we observe—by Riemann-Roch—the pair (p, q) moves in a pencil if and only if the linear series $K(-p) = K(-q)$, or only if $\phi_K(p) = \phi_K(q)$. As a result, any two points on the curve that move in a pencil map to the same point on the canonical curve, and conversely.

There's one more important thing to say. Let's recall the *geometric* form of the Riemann-Roch theorem. If C is a *non-hyperelliptic curve*, and $C \to \mathbb{P}^{g-1}$ is the canonical imbedding, and if $D = p_1 + \dots + p_d$ is a divisor on C consisting of distinct points, then the dimension of $\mathcal{L}(D)$ is the number of relations on the points p_i in $\mathbb{P}^{g-1}$. So, for a non-hyperelliptic curve, finding nontrivial linear systems corresponds to finding points on a hyperplane.

Let C be possibly hyperelliptic; we want to generalize this. We get a canonical map $\phi_K : C \to \mathbb{P}^{g-1}$ which may be two-to-one. How can we interpret the geometric Riemann-Roch? If D is any effective divisor on C, not necessarily with distinct points, we are going to define the *span* of the divisor D on the canonical model. That is, we let $\overline{D}$ be the intersection of the hyperplanes $H \subset \mathbb{P}^{g-1}$ whose pre-image $\phi_K^{-1}(H) \supset D$. This is the scheme-theoretic pre-image, and there could be multiplicities. So the span of $2p$ for $p \in C$ is the tangent line at p. Note that the span of a divisor does not require the curve to be hyperelliptic, though on a non-hyperelliptic curve with distinct points, the span of two distinct points is the line through them. Given a hyperelliptic curve with ϕ_K two-to-one, the span of twice a point $p \in C$ is just the point in $\mathbb{P}^{g-1}$ (because of the two-to-oneness).

8.2 Theorem (Riemann-Roch)**.** *The dimension of $\mathcal{L}(D)$ is the difference between $d-1$ (the expected dimension) and the dimension of the projective space $\overline{D}$.*

Something is a little funky here. If the curve C is hyperelliptic, and the map ϕ_K is two-to-one onto a rational normal curve. How are we going to find *divisors* on the curve C whose span is dimension less than $d-1$ (so they will move in a pencil)? Here's the thing. Obviously we can find such a collection if we use a pair of two points mapping to the same thing via ϕ_K. But, since the image is a *rational normal curve*, we note

8.3 Proposition. *Any $d \leq g$ points on $\phi_K(C)$ is linearly independent (spans a $\mathbb{P}^{d-1}$).*

As a result, the *only* way a divisor can move in a pencil, for $d \leq g$, is if it contains two points in the g^1_2. The complete linear system D is always a multiple $mg^1_2 + D_0$ for some residual D_0.

8.4 Corollary. *The only special linear series on a hyperelliptic curve are those that contain the g^1_2.*

By "special" we mean that a divisor D is such that $h^0(K-D) > 0$. This amounts to saying that $\overline{D} \subsetneq \mathbb{P}^{g-1}$. We are using the fact that any collection of points on a rational normal curve either spans the projective space or is linearly independent. The only way to get a *birational* embedding of a hyperelliptic curve, as a result, is to take a special divisor. We have to take $d \geq g+2$. If we take $d = g+2$, we can get a birational map, and for a real embedding, we can take $d = g+3$. (This is on the homework.)

Hyperelliptic curves have a rational function of lowest possible degree, but they are the *most* resistant to actually embedding in projective space. Any other curve, as we'll see, can be embedding with smaller degree.

8.5 Example. Consider a hyperelliptic curve C of genus 3. Then, we need to take a divisor of degree at least six to imbed C in projective space. If D is a divisor of degree six *not* of the form $K + p + q, p, q \in C$, one gets an imbedding. (Check this.) The image is then a curve of type $(2,4)$ on a quadric. To see that it lies on a quadric, it's the standard calculation which we'll do in a moment. To see that it's a type $(2,4)$, it's because the degree is six and the genus three.

In general, to embed a hyperelliptic curve, if you really must, this is how you would do it: you would take a random divisor class of degree $g+3$, and the image in $\mathbb{P}^3$ will be a curve of type $(2, g+1)$ on a quadric.

8.6 Example. A non-hyperelliptic curve C of genus 3 imbeds in $\mathbb{P}^2$ as a smooth quartic, via the canonical imbedding. From this, we can see all the linear series on C. For instance, what's the smallest degree of a nonconstant rational function? Three, as we saw last time—project from a point. So, at least, it's a three-sheeted cover of $\mathbb{P}^1$, and you can do this in a 1-parameter family of ways, corresponding to the points of C.

§2 $g = 4$

We're going to *exclude* the hyperelliptic case, because in some sense we know what those look like, from the previous section. To focus our ideas, let's consider a genus four curve C, and wonder whether C is a *three-sheeted* cover of $\mathbb{P}^1$. The key to understanding the geometry of the curve—in every respect—is to look at the canonical model ϕ_K :

$C \hookrightarrow \mathbb{P}^3$. Let's take C_0 as the image $\phi_K(C)$; this is a degree six curve in $\mathbb{P}^3$. As we've discussed, we first want to know what surfaces it lies on.

Consider the pull-back

$$\phi^* : H^0(\mathcal{O}_{\mathbb{P}^3}(2)) \to H^0(C, 2K).$$

The first has dimension ten; the second, by Riemann-Roch, has dimension nine. Consequently, it follows that:

8.7 Proposition. *The canonical imbedding $C_0 \subset \mathbb{P}^3$ lies on a quadric.*

By the same logic, we find:

8.8 Proposition. *C_0 lies on five cubics.*

Proof. This follows from the same analysis, since there are twenty linearly independent cubics in $\mathbb{P}^3$. ▲

Of these five cubics, however, four of them are multiples of the quadric with a hyperplane section. But, there is a fifth linearly independent one. Since the curve C_0 is irreducible, we find that C_0 lies on an irreducible cubic as well, a unique one. By Bezout, and since C_0 is a sextic, we find:

8.9 Proposition. *C_0 is the complete intersection of a quadric and a cubic surface.*

Consequently, the ideal of C_0 is generated by this quadric equation and this cubic equation, by Noether's $AF + BG$ theorem.

Remark. There are different notions of what it means for a variety X in projective space to be cut out by hypersurfaces. One is that X is the set-theoretic intersection of these; one is that X is the scheme-theoretic intersection; one is that the *homogeneous ideal* of X is cut out by the ideals of these hypersurfaces. Noether's theorem states that the second implies the third for a transverse intersection. (The second implies that the homogeneous ideal of X is generated in *sufficiently large degrees* by the equations of the ideals of these hypersurfaces.)

In the plane case, Noether's theorem states that if F, G are three-variable homogeneous polynomials defining transversely intersecting curves in $\mathbb{P}^2$, then any homogeneous polynomial H vanishing on the finite set $V(F) \cap V(G) \subset \mathbb{P}^2$ can be written as $AF + BG$ for polynomials A, G. Noether's theorem is a fun application of sheaf cohomology.

Back to the main goals. Let Q be the quadric containing C_0, the canonically imbedded genus four curve in $\mathbb{P}^3$. We don't know whether Q is a cone or if it's smooth. (Note, also, that the complete intersection of any quadric and a cubic in $\mathbb{P}^3$ has genus four by adjunction.) In other words, the *non-hyperelliptic* curves of genus four are precisely the smooth complete intersections of quadrics and cubics in $\mathbb{P}^3$.

Once again, let's return to the main question: if C is non-hyperelliptic, is there a non-constant meromorphic, degree three function on the curve? Geometric Riemann-Roch gives the answer. A divisor of degree three on the curve C will move in a pencil if and only if it is contained in a line on the canonical model: that is, if their span is not a $\mathbb{P}^2$ but a $\mathbb{P}^1$. Are there such? Well, if such a line L intersects C_0 in three points, it intersects the quadric Q in three points, so it sits inside Q. And conversely: if a line is contained in Q, then it will meet the cubic surface in three points by Bezout. We find:

8.10 Proposition. *The divisors of degree three that move in pencils are precisely the intersections of lines on the quadric Q with C_0.*

This means that there are two distinct such pencils if Q is nonsingular, and one if Q is singular. As a result, C can be expressed as a three-sheeted cover of $\mathbb{P}^1$, but in either one or two ways—not in a one-parameter family of ways.

Remark. C_0 is a $(3,3)$ curve in the quadric Q, so if Q is smooth, we find automatically that there are degree-three maps to $\mathbb{P}^1$ by projecting.

Here we find a situation where different curves behave differently. As you increase the genus, there'll be more and more phenomena. We are going to get up to genus 6 (I assure you that something vaguely theoretical is coming up; we won't be here in December analyzing curves of genus 17), there'll be dozens of different types. There is not a botany course.

§3 $g = 5$

We'll consider non-hyperelliptic curves of genus 5, so any such C is imbedded in $\mathbb{P}^4$ via the canonical map. This is an *octic curve.* We'll look at the quadrics that contain the curve. As before, we consider the map

$$H^0(\mathbb{P}^4, \mathcal{O}_{\mathbb{P}^4}(2)) \to H^0(C, 2K),$$

and counting dimensions (15 and 12), we find that C lies on at least three linearly independent quadrics Q_1, Q_2, Q_3. Then C could be, and generically is, the complete intersection of these three quadrics. (If we start with a smooth curve which is the complete intersection of these three quadrics, then an exercise in adjunction shows that the genus is five.)

But, if C is given, and Q_1, Q_2, Q_3 are found as above, it *might* be that C is *not* the complete intersection of these three. In fact, the three might intersect in a surface. There is only one nondegenerate surface in $\mathbb{P}^4$ that lies on three quadrics, the cubic scroll. For this to happen, $Q_1 \cap Q_2$ must be reducible, and Q_3 would contain one of the components of $Q_1 \cap Q_2$. The only way this can happen if $Q_1 \cap Q_2$ is the union of a plane and a cubic surface, and Q_3 is a third quadric vanishing on that cubic surface. This is the cubic scroll. So, there are two possibilities:

1. $C = Q_1 \cap Q_2 \cap Q_3$.
2. $C \subsetneq Q_1 \cap Q_1 \cap Q_3$, and the latter is a cubic scroll.

As before, we now ask: Can we express C as a triple cover of $\mathbb{P}^1$? Once more, we consider divisors of degree 3, say D, on C. Then D moves in a pencil if and only if D is contained in a line on the canonical model, by geometric Riemann-Roch. But if you had three collinear points on C, the line would be contained in all three quadrics Q_1, Q_2, Q_3 since a line can only intersect a quadric twice. In the first case above, in particular, there are *no* such divisors. So, *not every curve of genus five can be a three-sheeted cover of* $\mathbb{P}^1$.

If, however, C is contained in three quadrics intersecting in a surface, we'll just state the answer. Here $S = Q_1 \cap Q_2 \cap Q_3$ is the rational normal scroll, and is constructed as follows. Choose a line in $\mathbb{P}^4$ and a conic in $\mathbb{P}^2 \subset \mathbb{P}^4$, choose an isomorphism between the line and the conic, and draw lines in $\mathbb{P}^4$ between corresponding points. Intrinsically, it is the blow-up of $\mathbb{P}^2$ at a point, imbedded in $\mathbb{P}^4$ by the divisor of conics in $\mathbb{P}^2$ through that points. (Equivalently, imbed $\mathbb{P}^2$ in $\mathbb{P}^5$ via the Veronese, and project from a point.) When C lies on a cubic scroll, one can check that C meets each line of the ruling of the scroll three times, and in this case we get a *unique* g^1_3.

It's possible for a genus five curve to be a three-sheeted cover of $\mathbb{P}^1$, but generically it does not happen.

Lecture 9
9/30

This class was taught by Anand Deopurkar. Today's topic is not strictly a continuation of what Joe has been doing, but it will be related, and it will involve a preview of the rest of the class.

The main question we are interested in is about *automorphisms* of compact Riemann surfaces. Let X be a smooth projective curve of genus g. The question we are going to be dealing with is, what can we say about the automorphism group $\mathrm{Aut}(X)$?

§1 Low genus examples

9.1 Example. In genus zero, we have $X = \mathbb{P}^1$, and all isomorphic maps $\mathbb{P}^1 \to \mathbb{P}^1$ are just coordinate transformations. They look like $[X, Y] \mapsto [aX + bY, cX + dY]$; the matrix $\begin{bmatrix} a & b \\ c & d. \end{bmatrix}$, for a nonsingular matrix. This depends only up to scaling, so the automorphism group is PGL_2—there is a three-dimensional space of automorphisms. There are *infinitely many.* In fact, given any three ordered distinct points of $\mathbb{P}^1$, there is a unique automorphism sending them to $0, 1, \infty$. The automorphism group is 3-transitive.

9.2 Example. In genus one, X is an *elliptic curve* (once we fix a point $p \in X$) and there is a group law on X with origin p. There are infinitely many automorphisms of X given by translates; i.e. $x \mapsto x + q$ for each $q \in X$. We can ask if these are *all* the automorphisms. In other, we can ask for automorphisms $X \to X$ that fix the origin p (since translations never do). There is obviously one such: the inverse in the group law, $x \mapsto -x$. All elliptic curves have this.

Are there even more? To answer this, we consider the linear series $2p$ on X. This is a g^1_2 on X, and in fact every divisor of degree two is one such, so we get a map of degree two $X \to \mathbb{P}^1$ ramified at four points by Riemann-Hurwitz. p is one of the points of ramification, because $\{p\}$ is a fiber. There are three more. Any automorphism of X that fixes p must fix this linear series $\{2p\}$, and consequently induces an automorphism of $\mathbb{P}^1$. There is a commutative diagram

$$\begin{tikzcd} X \arrow[r] \arrow[d] & X \arrow[d] \\ \mathbb{P}^1 \arrow[r] & \mathbb{P}^1. \end{tikzcd}$$

The ramification points of the linear series are sent to the ramification points, and consequently the four branch points are *permuted*. Let's call the image of p in $\mathbb{P}^1$ ∞; the three remaining branch points of $\mathbb{P}^1$ are permuted by this automorphism, and ∞ is fixed. Conversely, if we have an automorphism of $\mathbb{P}^1$ which sends ∞ to ∞ and permutes the three other points, then we can *lift* it to an automorphism of the cover $X \to \mathbb{P}^1$. This is basically the theory of covering spaces once you forget the branch points. The lift is not unique, but it *is* unique once you fix where a given point (unramified) goes (there are two choices). In other words, the lift is unique up to the hyperelliptic involution (which is also known as inversion in the group).

Consequently, computing the order of automorphism of the elliptic curve amounts to computing the number of automorphisms of $\mathbb{P}^1$ that permute three given distinct points and send $\infty \mapsto \infty$. If $a, b, c \in \mathbb{C} \subset \mathbb{P}^1$ are the three points, then we have a linear map $\mathbb{C} \to \mathbb{C}$ permuting the three a, b, c. Once we count these, we multiply by two to account for the hyperelliptic involution. There are not many tuples (up to isomorphism) for which there is a nontrivial automorphism of $\mathbb{C}$ permuting $\{a, b, c\}$. Simple calculations show that there are only three such: $\{1, 0, -1\}$ (given by inversion), and the second is $\left\{1, \omega, \omega^2\right\}$, given by rotation by ω, a cube root of unity.

The conclusion is that there is *one* elliptic curve (the double cover of $\mathbb{P}^1$ branched over $\{\infty, 1, 0, -1\}$) which has an automorphism group of size four. There is *one* elliptic curve (the double cover of $\mathbb{P}^1$ branched over $\left\{\infty, 1, \omega, \omega^2\right\}$) which has an automorphism group of size six. There are two *special* elliptic curves with more than the normal share of automorphisms.

We can also see this as follows. An elliptic curve is just a torus $\mathbb{C}/L$ for $L \subset \mathbb{C}$ a lattice, and its universal cover is $\mathbb{C}$. Any automorphism of the torus is (by the theory of covering spaces) the same thing as an automorphism of $\mathbb{C}$ (say, fixing zero), that sends L into itself. Such an automorphism will be necessarily linear. In other words, to determine the automorphism group of $\mathbb{C}/L$, we have to find all $\alpha \in \mathbb{C}$ such that $\alpha L \subset L$. For almost all L, the only possible choices are $\alpha = \pm 1$. There are two special lattices for which there are more. The first is the square lattice $\mathbb{Z}[i] \subset \mathbb{C}$ (which has $\alpha = \pm i$ as well) and the hexagonal lattice $\mathbb{Z}[\omega] \subset \mathbb{C}$ (which has $\alpha = \pm\omega, \pm\omega^2$ as well).

In particular, the automorphism group of an elliptic curve can only be $\mathbb{Z}/2\mathbb{Z}, \mathbb{Z}/4\mathbb{Z}, \mathbb{Z}/6\mathbb{Z}$. Again, however, there are *infinitely many* automorphisms of a genus one curve (where the origin is not fixed).

§2 The Hurwitz bound

For the rest of the time, we want to prove:

9.3 Theorem. *Let X be a smooth projective curve over $\mathbb{C}$ of genus $g \geq 2$. Then* $\mathrm{Aut}(X)$ *has cardinality at most* $84(g-1) = -42\chi(X)$.

There is a sudden change from genus one and genus zero to genus two and up—suddenly, there are not so many automorphisms.

Proof. We have two clearly distinct steps. First, we have to show that $|\mathrm{Aut}(X)| < \infty$. We'll do this in one way and sketch a couple of other ways. The second case is, assuming the group is finite, to show that its order is at most what claimed. You really have to do the first step before the second.

It is surprising where the seemingly random number 84 comes up. In fact, 42 will be pretty prominent in the proof.

2.1 Step 1

Let's start with the finiteness. Let $G = \mathrm{Aut}(X)$ be the automorphism group. Associated to X, there is a vector space $V = H^0(K, X)$. If $g \in G$ is an automorphism of X, then it acts on V by pull-back, so V is a G-representation. In other words, we get a map $G \to \mathrm{GL}(V)$.

Let us start by showing that the image of G in $\mathrm{GL}(V)$ is finite. To see this, we recall further structure associated with V; there is a canonically defined *bilinear pairing* $V \otimes_{\mathbb{C}} \overline{V} \to \mathbb{C}$. This was introduced earlier; it sends (ω, η) to $-i\int_X \omega \wedge \overline{\eta}$ (the constant factor is to make it positive definite). Clearly the action of G on V respects this hermitian form. So, in fact, G maps into the unitary group $U(V)$. Now the unitary group is compact.

There's even more structure to V, though: there's a distinct lattice. The integral cohomology $H^1(X, \mathbb{Z}) \subset H^0(K)$ which is a lattice (recall that this is how the Jacobian was constructed). Any automorphism of X preserves this lattice as it must preserve the cohomology. Consequently, G also lands in $\mathrm{GL}_{2g}(\mathbb{Z})$, the space of automorphisms of this lattice. This is *discrete.*

But a group which is discrete and which lies in a compact group must be finite. Consequently, the image of G is finite. (We haven't used $g \geq 2$ yet.)

Now, we claim that $G \to \mathrm{GL}(V)$ is *injective.* Combined with the previous argument, this will prove that G is finite. Consider an element $\sigma \in G$ which acts trivially on $H^0(K)$, and let's consider the de Rham cohomology $H^*(X; \mathbb{C})$. Here we have

$$H^1(X; \mathbb{C}) = H^0(K) \oplus \overline{H^0(K)},$$

as we saw earlier. Any automorphism acts trivially on H^0, H^2. Any automorphism in the kernel of $G \to \mathrm{GL}(V)$, from the above decomposition, acts trivially on H^1, and thus on *all* the topological cohomology.

We'll now use the Lefschetz fixed-point theorem. Recall what it states. If $\phi : K \to K$ is an endomorphism of a finite CW complex, then the number of fixed points of ϕ ("morally") is the *Lefschetz number*,i.e. the alternating sum of the traces $\phi_* : H^i(K) \to H^i(K)$. To be precise, if the number of fixed points is finite, then one can say that the number of fixed points counted with the right multiplicity is the Lefschetz number. To get the multiplicities, one has to take the intersection numbers of the graph of ϕ with the diagonal $\Delta \subset K \times K$, with orientation; in the case of complex manifolds, we don't have to worry about orientation (all intersection numbers are positive).

In our case, we find that the Lefschetz number of $\sigma \in G$ is $\chi(X)$ because any automorphism σ in $\ker G \to \mathrm{GL}(V)$ acts trivially on the cohomology, as we saw above. This Lefschetz number is $2 - 2g < 0$. Consequently, there can't be finitely many fixed points by the above (as $g \geq 2$), so we find that σ must fix an *infinite* number of points. So $\sigma = 1_X$ because the fixed points will have a limit point. (Here is where the genus ≥ 2 thing is used.)

2.2 Step $1'$

Alright, that's one proof of the finiteness, but there are others. Here is a sketch. We don't have all the tools yet. This goes back to the case of elliptic curves; recall that we had, for elliptic curves, marked special points on the curves and automorphisms that preserved the origin had to preserve them.

More generally, given an automorphism $\sigma : X \to X$ of X which fixes more than $2g + 2$ points, the claim is that it must be the identity.

Once again one can prove this using the Lefschetz fixed-point theorem: when one computes the action on cohomology, one gets integer-valued matrices whose eigenvalues must have absolute value one (since they are unitary as we saw above) and consequently the Lefschetz number can be at most $2g + 2$.

Then, one can argue as follows: one associates to a Riemann surface X a finite set $W \subset X$ *canonically*, such that any automorphism has to permute these points. These are called *Weierstrass points.* We will show (later) that there are enough Weierstrass points that there is an injection $\mathrm{Aut}(X) \to \mathrm{Aut}(W)$ (i.e. such that $|W| > 2g + 2$). Consequently, we win.

2.3 Step $1''$

Here is a sketch of something that works in characteristic p which proves the claim about finiteness of the automorphism group (for $g \geq 2$). The theme is really nice. Look at this group $G = \mathrm{Aut}(X)$, as before. The neat thing that happens in algebraic geometry is that this itself turns out to be a scheme. (This happens a lot: for instance, divisors up to linear equivalence are not just a set, but a variety—the Jacobian.) Anyway, so $G = \mathrm{Aut}(X)$ is this quasi-projective scheme. (We already saw this in genus zero and one.)

We want to show that it is finite—to do this, we can show that it is zero-dimensional. To do this, we can show that the *tangent spaces* are zero. Let's compute the tangent space at the identity (enough by translation); this is the space of holomorphic vector fields of X. Indeed, you can think of as a vector field as an infinitesimal deformation of the identity. But this is zero when $g \geq 2$.

2.4 Step 2

This is not too hard. Now, once we've accepted that $\mathrm{Aut}(X)$ for X a compact Riemann surface of genus at least two is finite, we're going to show that its cardinality is at most $84(g - 1)$. Let G be the automorphism group. We form the quotient $Y = X/G$, which is also a compact Riemann surface such that $X \to Y$ is holomorphic. This is not too hard. Details are in Miranda's book "Algebraic curves and Riemann surfaces." (It is generally not a covering map, if G has fixed points.) We use two observations:

1. There are only finitely many points of X such that there exists $g \in G$ fixing the point.

2. The stabilizer of each point of X is cyclic. This is amusing to see using a power series expansion. For instance, let z be a local coordinate near a point $p \in X$ with nontrivial stablizer. For any element σ in the stabilizer, we get an element of $\mathbb{C}^*$ by sending σ to the first coefficient in the local coordinate $\sigma(z)$. This is injective: that is, there can't be a nontrivial automorphism of X which fixes a point p and whose tangent space at that point is the identity. This is an exercise using the finite order.

With these in mind, it's not too hard to make X/G into a Riemann surface.

There is also a purely algebraic way of forming X/G, via the "categorical quotient" process of quotienting a quasi-projective variety by a finite group. Since this preserves normality, X/G constructed in this way will be a smooth scheme.

OK, anyway, we have our map

$$\pi : X \to Y = X/G.$$

This has degree $|G|$. Let g be the genus of X and h that of Y. Let's apply Riemann-Hurwitz. We find

$$2g - 2 = |G|(2h - 2) + R,$$

for R the ramification. What can we say about ramification? Suppose there are b branch points, $p_1, \dots, p_d \in Y$. Say that above p_i, there are f_i distinct points in X above p_i; this will be one orbit of G, and the behavior at each of these distinct points will be uniform. That is, the ramification index at each of these f_i points above f_i will be the same, say e_i. We know that $e_i f_i = |G|$, the degree of the map. The degree of the ramification divisor is consequently

$$\sum_i f_i(e_i - 1).$$

This is because there are f_i points over the ith branch points, each contributing $e_i - 1$. This is also $|G| \sum_i (1 - \frac{1}{e_i})$, and we find

$$2g - 2 = |G|(2h - 2 + \sum_i (1 - \frac{1}{e_i})).$$

Let us call Q the quantity within parentheses. Here Q is positive, and if we find out the *minimum* positive value for such an expression, then we will have an upper bound for the genus. This is where the number comes from.

We are trying to minimize $Q = 2h - 2 + \sum_i \left(1 - \frac{1}{e_i}\right)$ subject to the conditions that it be positive, $h, b \geq 0$, and all the e_i are integers at least two. For $h \geq 2$, one checks that $Q \geq 1$. For $h = 1$, necessarily $Q \geq 1/2$. Then one has to check for $h = 0$. The best one can achieve, as one may check, is $-2 + (1 - 1/2) + (1 - 1/3) + (1 - 1/7) = \frac{1}{42}$. This gives the minimum value of Q, and consequently the bound for $|G|$. ▲

One might ask whether this bound is *sharp.* It is actually sharp, and it turns out:

1. There are infinitely many genera $g \geq 2$ for which the bound is achieved (i.e. for which there exists a curve X of genus g with $84(g-1)$ automorphisms).

2. In genus two, this bound is unattainable.

3. In genus three, the *Klein quartic* attains the curve. This is a quartic in $\mathbb{P}^2$, given by $X^3Y+Y^3X+Z^3X=0$, with 168 automorphisms. All its automorphisms come from automorphisms from $\mathbb{P}^2$, but they're not all easy to see. Cyclic permutation is one. An automorphism of order seven is given by $[X, Y, Z] \mapsto [\zeta^4 X, \zeta^2 Y, \zeta Z]$ for ζ a root of order seven. There is also an element of order two, which is given by a rather complicated matrix.

4. For a long time, the Klein quartic was the only curve known for which the bound was achieved. But appropriate covers of the Klein quartic can be used to achieve the bound in higher genera.

5. There are also infinitely many genera for which the bound is *not* sharp, i.e. for which no curve reaching the bound is achieved.

The function $N(g)$ defined as the size of the largest automorphism group of a curve of genus g is an object of current research. One can also ask about *lower* bounds for $N(g)$; here, one has to exhibit curves with large automorphism groups.

9.4 Example. Using hyperelliptic curves with $2g+2$ branch points whose branch points are vertices of a regular $2g+2$-gon (i.e. $y^2 = x^{2g+2}-1$), one can permute these branch points (and thus find automorphisms of the curve by lifting) to get $N(g) \geq 2(4g+4)$ (where the 2 comes from the hyperelliptic involution). Amusingly, $N(g) = 8g+8$ holds for infinitely many g.

In positive characteristic, the automorphism group of a curve of genus g is bounded by something like $16g^4$; it can grow more rapidly than in the characteristic zero case.

9.5 Example. The curve $y^2 = x^p - x$ in characteristic p is branched over all the $\mathbb{F}_p$-points, and consequently the genus (proportional to p) is a lot less than the number of automorphisms (which is at least $|\mathrm{PGL}_2(\mathbb{F}_q)|$, coming from automorphisms of the line).

Lecture 10
10/7

It's time to get serious. We've been looking at a lot of examples and basic results, and now we should prove a theorem. We've been studying linear series on curves, so now let's ask a very basic question. We'll answer it completely, and then find that it was the wrong question—so ask a better question, and answer that too.

So here's the question. C is, as always, a smooth projective curve of genus g, and D a divisor (or divisor class) of degree d. What's the dimension $r(D)$ of the linear system $|D|$ in which D moves? Alternatively, what is $h^0(D) = r(D)+1$?

We have to choose one notation or the other, so let's use $h^0(D)$, and express everything in terms of that.

§1 Preliminary remarks

First, we can make some obvious remarks about the above question.

1. If $d < 0$, then $h^0(D) = 0$: there are no sections.
2. If $d \geq 2g - 1$, then $h^0(D) = d - g + 1$ (by Riemann-Roch).
3. In general, $h^0(D) \leq d + 1$.
4. For a divisor of degree $2g - 2$, we have $h^0 < g$ except for the canonical class.

In other words, we can draw a graph where the vertical axis represents the number of sections, and the horizontal axis represents the degree, then the graph starts out at zero—and at the end, it becomes a line of slope one. Riemann-Roch implies that there are upper and lower bounds to this graph (which should be drawn). So we get a *range* of values allowed by Riemann-Roch. However, that's not the full story—only about half of the parallelogram actually occurs.

§2 The next theorem

We'll now prove *Clifford's theorem,* which will give a complete answer to the question originally posed. There are a lot of forms, but we'll quote one suited to the picture.

10.1 Theorem. *If D is a special divisor, then $h^0(D) \leq \frac{d}{2} + 1$; alternatively, a special divisor can't move in a projective space of dimension $> d/2$. Equality holds if and only if one of three things is true:*

1. *$D = 0$.*
2. *$D = K$ is the canonical series.*
3. *C is hyperelliptic and D is a multiple of the base g^1_2.*

Recall that a divisor is *special* if $h^0(K - D) > 0$.

Proof. It is easy to check that the boundary cases verify equality in the inequality; this will be omitted. Anyway, if we think in terms of the parallelogram that restricts the graph of the dimension versus the degree, then Clifford's theorem restricts the graph to about half of it.

The proof of the inequality is quite simple. One bit of basic language. If $\mathscr{D}$ is a linear series—not necessarily complete—and $\mathscr{E}$ is a linear series—not necessarily complete—then $\mathscr{D}+\mathscr{E}$ is the subspace of the associated *complete* linear system spanned by divisors of the form $D' + E'$ for $D' \in \mathscr{D}, E' \in \mathscr{E}$. If you want to think of it in terms of vector spaces and sections, then $\mathscr{D}$ corresponds to a line bundle $\mathcal{L}$ and $V \subset H^0(\mathcal{L})$, and $\mathscr{E}$ corresponds to $\mathcal{L}'$ and $W \subset H^0(\mathcal{L}')$; we take the image of $V \otimes W \to H^0(\mathcal{L} \otimes \mathcal{L}')$. So we can *add* linear series.

Here's an amusing fact.

10.2 Lemma. *If we have a linear map $\phi : V \otimes W \to U$ of complex vector spaces, which is one-to-one on something of the form $V \otimes \{w\}$ for some $w \in W$, and on $\{v\} \otimes W$ for some $v \in V$, then* $\dim(V \otimes W) \geq \dim V + \dim W - 1$.

We're not going to use this in the proof, but note that it is *false* over the reals. It's an exercise to find a counterexample.

Anyway, if we apply it here, we find that the dimension of $\mathscr{D} + \mathscr{E}$ is at least $r(\mathscr{D}) + r(\mathscr{E})$ (be careful with projective dimension and vector space dimension). The conclusion is:

10.3 Lemma. $r(\mathscr{D} + \mathscr{E}) \geq r(\mathscr{D}) + r(\mathscr{E})$.

But this lemma is a *completely obvious fact*. What does it mean for a linear series to have dimension r? To say that $r(\mathscr{D}) \geq r$ is to say that for any choice of r points on the curve $p_1, \ldots, p_r$, there exists a divisor $D \in \mathscr{D}$ containing these points. If we use this fact, then the lemma becomes a simple statement. If we have $r(\mathscr{D}) + r(\mathscr{E})$ points on the curve, we can find a divisor in $\mathscr{D}$ containing the first $r(\mathscr{D})$ of them, and a divisor in $\mathscr{E}$ containing the second $r(\mathscr{E})$ of them, and then add them to get something in $r(\mathscr{D} + \mathscr{E})$ containing all of them.

So you don't need this tensor product stuff to prove the lemma.

Now let's prove Clifford's theorem. Let $\mathscr{D} = |D|$ be the complete linear series associated to the special divisor D, and let $\mathscr{E} = |K - D|$; then $\mathscr{D} + \mathscr{E}$ is contained in the canonical series, and the above lemma states that

$$h^0(\mathscr{D}) + h^0(\mathscr{E}) \leq g + 1.$$

But we also have Riemann-Roch, which states that

$$h^0(\mathscr{D}) - h^0(\mathscr{E}) = d - g + 1 \quad (\mathscr{E} = |K - D|).$$

If we add these and divide by two, we get $h^0(D) \leq \frac{d+2}{2}$, which is Clifford's theorem. We haven't yet shown that if equality holds implies the special cases, which we won't. ▲

§3 Remarks

So Clifford's theorem was pretty elementary. Now we want to point out that this was the wrong question to ask. There are two variants of this question which will be of interest to us for the rest of the semester. If we thought of the graph of the dimension-degree thing in a parallelogram, then we cut the thing in half by Clifford's theorem. But we also showed that things in the middle of the graph—where equality holds in the inequality—holds for Clifford's theorem.

The main point is, on most curves, even the half of the parallelogram isn't filled. This is the basic subject of Brill-Noether theory—what linear series exist on a *general* curve? In other words, it's not only the case that points on the edge of this diagonal occur only on special curves, but that even near the diagonal, most curves get near it.

So we might ask the following. **For which r, d does every curve of genus g possess a g^r_d?** We talked about this in the basic notions talk, and the answer is that there is a slow-growing conic that bounds the graph of the dimension-degree thing generally, and gives a better bound than Clifford's theorem for a general curve.

Anyway, there's something else wrong with this question. Remember we said at the beginning of the course that the way to understand algebraic curves is to think of them simultaneously in two ways—as abstract curves, and as curves in projective space. We need to know how to pass between the two viewpoints. If we start with an abstract curve, we'd like to know how to imbed it in projective space. Note something about the extremal examples in Clifford's theorem—the ones that give inequality are those that don't give imbeddings of the curve (they're multiples of the g_2^1).

So here's another question? On a curve of genus g, **what very ample linear series exist?** We are mostly interested in linear series that give imbeddings, or at least birational imbeddings. In other words: on a (smooth, projective) curve C of genus g, D a divisor class of degree d that is very ample (such that the map ϕ_D associated to D imbeds C as a curve in projective space $\mathbb{P}^r$), what relations can we say about possible values of (g, r, d)? This includes, for instance, fixing r, d and considering what possibilities for g can arise.

Here's how we will answer the question today.

> **Question.** If we start with $C \subset \mathbb{P}^r$ is an irreducible, smooth, nondegenerate curve of degree d, how large can its genus be?

Anything that involves with three quantities requires a chart. For example, it's not hard to see what we know and don't know $r = 3$. If $r = 2$, then we have the usual genus formula—there isn't any ambiguity, but in higher dimensions, the degree *doesn't determine the genus.*

10.4 Example. When $r = 3$, and the degree is three, then we get the rational normal curve (of genus zero). For $d = 4$, one can get a projection of a rational normal curve (genus zero) or an elliptic curve which is the intersection of two quadrics (genus one). Note that a curve $C \subset \mathbb{P}^3$ of degree four, nondegenerate, can't have genus ≥ 2. For, by taking $D = \mathcal{O}(1)|_C$, then $h^0(D) \geq 4$ and D must be non-special by Clifford's theorem. So $g = 4 - h^0(D) + 1$ by Riemann-Roch. One can, by the same reasoning, show that if the degree is five, then the genus can be $0, 1, 2$ only. With a little more work, one can show that for degree six, then the genus is at most four. However, Clifford's theorem doesn't help us in degree seven and up.

Note that in the above example, there is always a maximum genus in each degree, and the lower genera actually occur. However, this isn't always true. We will get a sharp upper bound for the genus in terms of the degree, but not every number below the upper bound actually occurs as a genus.

§4 The main result

Here's Castelnuovo's approach to the previous problem. This succeeds much better than it has any right to. We start with a curve $C \subset \mathbb{P}^r$—smooth, irreducible, nondegenerate. Let D be a hyperplane section of C, typically for D general. We will simultaneously think of D as a divisor on C and as a configuration of points in $\mathbb{P}^{r-1}$.

Castelnuovo's approach in this set up is to bound from below $h^0(mD)$. Put it this way—when the degree is low (this is a general pattern), Clifford's theorem tells us that the linear series is non-special. Once we know that D is non-special, we can read the genus off from the degree and its dimension. But once we get to a certain degree, Clifford's theorem no longer tells us anything. Even Riemann-Roch doesn't tell us anything.

Anyway, if you can bound from below $h^0(mD)$, then for $m \gg 0$ we'll apply Riemann-Roch (because mD will eventually become non-special). Then, we will have $g = md - h^0(mD) + 1$, so we can do stuff with that. This is the basic outline of the argument.

Remark. This illustrates a general principle: we can ask abstract questions about linear series, but it's crucial that we use the imbedding in projective space.

OK. How do we estimate $h^0(mD)$, as outlined in the above approach? Consider the vector space $H^0(\mathcal{O}_C(mD))$, which we'll sometimes just write as $H^0(mD)$. We want to relate this to $H^0((m-1)D)$. To do this, note that D is a sum of d points, which will be the intersection $C \cap H$ for a general hyperplane, intersecting C transversely. We'll evaluate the sections of $H^0(mD)$ at these points. So we take the map

$$H^0(\mathcal{O}_C(mD)) \to H^0(\mathcal{O}_D(mD)) \simeq \mathbb{C}^d,$$

which is just evaluating the sections at the points of D, or taking the fibers of these sections. (This isn't the best notation.)

Note that the kernel of $H^0(mD) \to \mathbb{C}^d$ is just $H^0((m-1)D)$. We have an exact sequence

$$0 \to H^0((m-1)D) \to H^0(mD) \overset{\rho_m}{\to} H^0(\mathcal{O}_D(mD)).$$

To bound from below $H^0(mD)$, we will bound from below the *rank* of ρ_m. The rank of ρ_m is just the number of linear conditions on a section of $H^0(mD)$ to vanish at the points of D. Among sections of mD, there are sections of $H^0(\mathbb{P}^r, \mathcal{O}_{\mathbb{P}^r}(m))$. In other words, there is an inclusion $H^0(\mathbb{P}^r, \mathcal{O}_{\mathbb{P}^r}(m)) \hookrightarrow H^0(C, mD)$. (This inclusion might be proper—the linear system of hyperplanes might not be complete.)

So, we'll do something a little weaker. We'll estimate from below the rank of the map $H^0(\mathbb{P}^r, \mathcal{O}_{\mathbb{P}^r}(m))) \to H^0(mD) \to H^0(\mathcal{O}_D(mD))$. In other words, the number that we want (the rank of ρ_m) is at least the number of linear conditions on a *polynomial* of degree m in $\mathbb{P}^r$ to vanish on the d points of D. This has to do with "imposing independent conditions" on degree m hypersurfaces, as we discussed some lectures back.

We want to know: if we have a collection of d points, what is the *smallest* number of conditions they can impose on degree m hypersurfaces? At least $\min(d, m+1)$. When is it acheived? For collinear points. But, well, if you take a D from a hyperplane section, then D won't be contained in a line (e.g. you can get such a D by choosing three points not on a line and taking a plane containing them). So this bound is way too weak for our purposes.

Here is the crucial lemma.

10.5 Lemma (General position lemma)**.** *If C is irreducible and nondegenerate in $\mathbb{P}^r$, and $H \subset \mathbb{P}^r$ a general hyperplane, $D = H \cap C = p_1 + \cdots + p_d$ the intersection, then the points of D (i.e. $p_1, \ldots, p_d$) are linear general position in that $H = \mathbb{P}^{r-1}$.*

"Linear general position" means that no r are linearly dependent. Again, forget r—take $r = 3$. Then what the result states is that no three of the points on a general hyperplane section are collinear. This is a fun, fun theorem to try to prove directly; we'll give a proof in the next class (as a consequence of a more general statement). It's fun to try to prove a priori.

Let's accept the general position lemma and move on ahead. This is where it gets kind of interesting. We've taken a question about algebraic curves, and now we are talking about a simpler question: if we have a bunch of points in projective space, how many conditions do they impose? Let's ask a sub-question.

> **Sub-question.** Say $p_1, \dots, p_d \in \mathbb{P}^n$. Assume they span $\mathbb{P}^n$ and are in linear general position. How many conditions do they impose on degree m polynomials?

Before we answer this, there is a standard way of writing this. If we let D be the set of points $\{p_1, \dots, p_d\}$, then the number of conditions imposed by hypersurfaces of degree m is what we called the *Hilbert function* of D, $h_D(m)$. So we want to bound from below the Hilbert function.

We now answer the sub-question.

Answer. The number of conditions imposed is at least $h_D(m) \geq \min(d, mn+1)$. Here is a simple argument for the claim. Say first $d > mn+1$, so the minimum is $mn+1$. Then we want to say that, given $mn+1$ points of D, we claim that there exists a hypersurface $Z \subset \mathbb{P}^n$ of degree m, containing all but any one of these points. We can do this using the linear general position hypothesis.

Suppose we start with q among the $mn+1$-points; can we produce a hypersurface of degree m containing q but not the others. Well, take a bunch of hyperplanes. For any q, we just group the remaining mn points into m sets of n, and take the union of the hyperplanes spanned by each of these groups. By linear general position, none of those hyperplanes contains q, so we get the claim.

The same argument works also if $d < mn+1$.

So, we get a lower bound on the Hilbert function on a set of points in linear general position in projective space. This is certainly better than $\min(d, m+1)$, but it doesn't seem like it could be sharp—we're using incredibly crude methods (just looking at products of linear forms). But it is, remarkably, sharp.

In fact, for any configuration $D = \{p_1, \dots, p_d\}$ as above contained in a rational normal curve, the bound is achieved. The point is that a hyperplane containing any $mn+1$ points will contain the full curve.

Now we have all the ingredients in place. We can finish the Castelnuovo theorem. Going back to our original curve $C \subset \mathbb{P}^r$ of degree d, we *apply* this lemma to the points of a hyperplane section. Let $n = r-1$ be the hyperplane H in which $D = C \cap H$ lies. We find that, using the previous short exact sequence,

$$h^0(\mathcal{O}_C(mD)) \geq h^0(\mathcal{O}_C((m-1)D)) + \min(d, m(r-1)+1).$$

We know a priori that $h^0(D) \geq r+1$, because the hyperplanes cut a linear series of dimension r. Also we find that $h^0(2D) \geq (r+1) + (2r-1) = 3r$, assuming d is large enough, from the above inequality. Again, $h^0(3D) \geq 6r-2$ from the above inequality. At some point repeating this argument will fail because $\min(d, m(r-1)+1)$ will become

d. Let $d = m_0(r-1) + \epsilon + 1$ where $0 \le \epsilon \le r-2$. This will detect the point where the minimum shifts over.

The point is that this rapid increase occurs up until m_0. So we get

$$h^0(m_0 D) \ge \binom{m_0}{2}(r-1) + m_0 + 1.$$

If we continue to do this, we just add d. After m_0, we just keep adding d. So

$$h^0((m_0+k)D) \ge \binom{m_0}{2}(r-1) + m_0 + 1 + kd.$$

Now we're almost done.

Let's recall that $d = m_0(r-1) + \epsilon + 1$. Now, for large k, we have that $(m_0+k)D$ is non-special, and we can just apply Riemann-Roch. We find that

$$g = \deg(m_0+k)D - h^0((m_0+k)D) + 1.$$

This, however, is *at most* (by the lower bounds)

$$(m_0+k)d - \left(\binom{m_0}{2}(r-1) + m_0 + 1 + kd\right) + 1.$$

Here the kd's cancel, as it should be. We now lose the k's, and we just plug in d in terms of m_0; we find in the end

$$\boxed{g \le \pi(d,r) := \binom{m_0}{2}(r-1) + m_0\epsilon.}$$

This is the main result. Next time, we have to prove the general position lemma, and go through the derivation of this result more carefully.

Lecture 11
10/12

Today's class should be fun. We're going to review the material from last time on bounding the genus of a curve in projective space. We'll then prove some of the necessary lemmas for said estimate, and see some of the consequences.

§1 Recap

From last time, the situation was: we have $C \subset \mathbb{P}^r$, a smooth curve which is nondegenerate, irreducible. (Actually, much of this applies if C is just birational onto its image in $\mathbb{P}^r$, and not necessarily imbedded. The arguments go through.) The *basic goal* is to estimate the genus of C. The key part of the construction is to focus on a particular hyperplane section. That is, we take a general hyperplane $H \subset \mathbb{P}^r$, and let $\Gamma = H \cap C$ be the corresponding hyperplane section. Then Γ is a d-tuple of points $p_1, \dots, p_d$ (where the degree of C is d). They play a dual role: first, as a *divisor* on the curve; and second, as a configuration of points in projective space.

We are now going to estimate $r(m\Gamma)$ for various m, the dimension of the various linear series. The **basic estimate** is a series of inequalities.

1. $r(k\Gamma) - r((k-1)\Gamma) \geq$ the number of conditions imposed by Γ on $H^0(k\Gamma)$. This is pretty straightforward, and it's actually equality. This in turn is at least the number of conditions on the *subseries* $H^0(\mathbb{P}^r, \mathcal{O}_{\mathbb{P}^r}(k))$.

2. Castelnuovo estimates the number of conditions. If H is general, then the claim is that the points of Γ are in *linear general position.* We didn't prove this last time. Consequently, we find (by elementary means) that the number of conditions imposed is at least $\min(d, k(r-1)+1)$.

3. Let's see what that gives us. We know that $h^0(\Gamma) \geq r+1$ (because the map $H^0(\mathbb{P}^r, \mathcal{O}_{\mathbb{P}^r}(1)) \to H^0(\Gamma)$ is injective by nondegeneracy). Now, we apply the basic inequalities:

$$h^0(\Gamma) \geq r+1 \tag{2}$$
$$h^0(2\Gamma) \geq r+1+2r-1 = 3r \tag{3}$$
$$h^0(3\Gamma) \geq 6r-2 \tag{4}$$
$$\cdots \tag{5}$$

and this continues *as long as* $k(r-1)+1 \leq d$. This stops, though, once $d < k(r-1)+1$. That suggests we introduce m, the integer part of $\frac{d-1}{r-1}$, and we let $d-1 = m(r-1)+\epsilon$ where $0 \leq \epsilon \leq r-2$. (This because we need the difference.) So, the above pattern persists up until $h^0(m\Gamma)$, which gives

$$h^0(m\Gamma) \geq \binom{m+1}{2}(r-1)+m+1.$$

At this point, we could argue by Clifford that $m\Gamma$ is a *non-special* linear series, and then plug this into Riemann-Roch to determine the genus. However, it isn't necessary. We can just keep iterating by adding Γ and adding d to the inequality. We have

$$\boxed{h^0((m+n)\Gamma) \geq \binom{m+1}{2}(r-1)+m+1+nd.}$$

To *this*, we apply Riemann-Roch. For $n \gg 0$, this is a non-special linear series, so we find that

$$g(C) = \deg((m+n)\Gamma) - h^0((m+n)\Gamma)+1,$$

and this becomes (after cancellation of md and the fact that $\deg((m+n)\Gamma) = (m+n)d$),

$$g \leq md - \binom{m+1}{2}(r-1)-m-1.$$

Here this becomes

$$g \leq m^2(r-1)+m\epsilon+m.$$

Note that the m's now cancel, and what we're finally left with is

$$\boxed{g \leq \pi(d,r) \equiv \binom{m}{2}(r-1)+m\epsilon.} \tag{6}$$

So, we've reproduced the basic definition, *modulo* the linear general position of the points of Γ. Two remaining things we want to do today: first, how this bound behaves in practice. Then, prove the linear general position.

§2 The bound, again

Let's make a chart.

d	m	ϵ	π
r	1	0	0
$r+1$	1	1	1
$r+2$	1	2	2
$\dots$			
$2r-2$	1	$r-2$	$r-2$
$2r-1$	2	0	$r-1$
$2r$	2	1	$r+1$
$\dots$			
$3r-3$	2	$r-2$	$3r-5$
$3r-2$	3	0	$3r-3$
$3r-1$	3	1	$3r$

The first row states that a degree r curve in $\mathbb{P}^r$ is a rational normal curve. Note that for $r \le d \le 2r-2$, we find $d \le d-r$. In this range for $d < 2r-2$, then Clifford's theorem states that Γ is non-special as a divisor, and the genus is in fact equal $d - h^0(\Gamma) + 1 \le d - r$. Nothing interesting happens here.

The first curves imbedded by special linear series occur for degree $2r$; this is is because we can consider *canonical* curves of genus $r+1$. As we keep going down the table, the ϵ terms keep increasing by one at each stage and then are reset to zero; then the π's start jumping faster and faster.

For fixed r, as $d \to \infty$, we have

$$\pi(d,r) \sim \frac{d^2}{2(r-1)}.$$

We are going to get a *lot* out of this, other than the intrinsic interest. But we still have to check the basic estimate (which bounded below the number of conditions imposed by Γ). Finally, we have to ask: what are the odds that the upper bound is sharp? Given the crudeness of the methods, you mightn't expect the bound to be sharp. But it is. We are going to show that the Castelnuovo bound is sharp. To see this, we have to *produce* lots of curves in projective space with maximal possible genus. So we'll have to *characterize* when curves have maximal genus. Finally, we have to obtain consequences: we're going to prove a number of basic facts about, e.g. canonical curves (such as Noether's theorem that canonical curves are projectively normal, and the Enriques-Babbage theorem that non-trigonal canonical curves are cut out by quadrics).

That's the agenda for the next week.

§3 The general position lemma

Let $C \subset \mathbb{P}^r$ be an irreducible, nondegenerate curve. We need to prove:

11.1 Theorem. *For H a general hyperplane in $\mathbb{P}^r$ and $\Gamma = C \cap H$, then Γ is in linear general position in $H \simeq \mathbb{P}^{r-1}$.*

11.2 Example. If $r = 3$, we have a space curve, and what we're saying is that the general plane section of a space curve has no three collinear points. Let's prove this. We have a two-parameter family of secant lines, parametrized by pairs of points. The *family* of secant lines is irreducible for this reason. A general secant line will meet the curve again; it won't be a trisecant. There is at most a one-dimensional family of trisecants. Every trisecant lies only on a one-parameter family of planes, so the dimension of the space of planes containing a trisecant has dimension at most two. So there are planes that contain no trisecants. This means that there are planes H such that $C \cap H$ contains no three collinear points.

Except showing that not every secant is a trisecant is not obvious at all. Castelnuovo did this with difficulty.

We're going to look at this more generally. Again, let's think of a curve $C \subset \mathbb{P}^3$. We have this curve and a plane H, general, which meets the curve in d distinct points. Let's move the plane H around, taking care to avoid planes tangent to the curve. So let's move the plane in a 1-parameter family (parametrized by a *real* parameter t) H_t of planes *transverse* to the curve. As we move the plane, the set of points of intersection $H_t \cap C$ varies with the parameter t. As H wanders, we can try to follow each of these points individually. Since t is a real parameter, we can do this in such a way to avoid the tangent hyperplanes.

At the end of the day, we want H to wander in such a way to come back to the original position. While the points have to *collectively* come back to the original position, they might get permuted. If $\{p_1, \ldots, p_d\}$ is the initial set of points, following the loop on H will induce a *permutation* in S_d. The claim is going to be that *any* permutation can be induced—the monodromy group is full the symmetric group. The consequence is that there will be no distinguished subsets of these points $H \cap C$. It consequently could not be the case that *some* triples are collinear and some triples and noncollinear. This is going to easily imply the general position lemma.

§4 Monodromy

Let's discuss this in general. This is a subject you can approach topologically, algebraically, or geometrically. Here's the basic setup. We have a map $\phi : X \to Y$ of varieties, smooth and projective. This is going to be a generically finite, dominant morphism. We want to say that there is a Zariski open subset $U \subset Y$ such that $\phi^{-1}(y) \subset X$ consists of *reduced* points, of cardinality equal to the degree. In fact, $\phi^{-1}(U) \to U$ will be a *topological* covering space.

So let's look at this covering space $\phi^{-1}(U) \to U$. You have to go back to the early days when you were taught about covering spaces and fundamental groups. If $p \in U$, then the fiber $\phi^{-1}(p)$ is a finite set. The basic fact is that $\pi_1(U, p)$ *acts* on the fiber $\phi^{-1}(p)$: the way it does so is to take a loop γ at p and a choice of point q in the fiber, take a unique lifting $\widetilde{\gamma}$ to $\phi^{-1}(U)$ starting from q and ending at some other point, which is $\gamma.q$.

11.3 Definition. The induced subgroup of $S_{\phi^{-1}(p)}$ which is the image of $\pi_1(U, p)$ is called the **monodromy group.**

This is not the same thing as the group of deck transformations. In general, the covering space will have no deck transformations; to say that it does is to say that the covering space is normal, or Galois.

OK, back to the varieties. We can associate the monodromy group associated to any dominant, generically finite map of varieties associated to the covering space generically induced.

There are other ways of characterizing monodromy, as we said. Let's do the algebraic version. We won't justify this, as we won't use it. But if you don't like the complex topology, you can do it algebraically. Let $\mathbb{C}(Y) \hookrightarrow \mathbb{C}(X)$ be the finite extension of quotient fields. Take the Galois normalization $L \supset \mathbb{C}(X)$ (inside some algebraic closure). The *monodromy group*, say G, is the Galois group of $L/\mathbb{C}(Y)$.

It's interesting—if you go back to the early history of Galois theory, what Galois himself wrote was unreadable. You have to interpret very carefully, and nobody understood what Galois was talking about until 1870 when Jordan wrote a book on them. But half the examples are very familiar to us—number fields, and then the Galois groups associated to extensions of function fields. This stuff was a part of Galois theory that fell into disuse in the twentieth century.

Here's the geometric description (partial). We will use it. Let's make some observations about the monodromy group G.

11.4 Proposition. *G is transitive on the fiber if and only if the total space X is connected. G is twice transitive (i.e. acts transitively on ordered pairs) if and only if $X \times_U X - \Delta$ is connected. More generally, if $X_U^r = X \times_U X \times_U \cdots \times_U X - \Delta$ (r times) with all diagonals removed, so that*

$$X_U^r = \{(q, p_1, \ldots, p_r) \in X^r : q \in U, \phi(p_i) = q, \quad p_i \neq p_j\}.$$

Then G is r-transitive if and only if X_U^r is connected.

Note that $X \times_U X - \Delta$ parametrizes distinct pairs of X which lie in the same fiber. This is really a direct consequence of the definitions. Note that X_U^r is a covering space of U.

We aren't determining G, but we can at least say something about it.

§5 Proof of the general position lemma

We will prove something more general:

Let $C \subset \mathbb{P}^r$ be a nondegenerate curve of degree d. Let $C^* \subset \mathbb{P}^{r*}$ be the dual hypersurface, which is the family of tangent (or non-transverse) hyperplanes. Take $U = \mathbb{P}^{r*} - C^*$ be the family of transverse hyperplanes (transverse to C). Let X be the set of pairs (H, p) such that $p \in H \cap C$ and $H \in U$. Consequently,

$$X \subset U \times C$$

is a branched cover of U, of degree d (since every transverse intersection has d points).

11.5 Theorem (Uniform position theorem)**.** *The monodromy group of $X \to U$ is the symmetric group S_d.*

This is true *only in characteristic zero.*

Proof. The first claim is that the monodromy group G is twice transitive. To see this, well, we have to consider $X_U^2 = X \times_U X - \Delta$: this is the set of triples (H, p, q) such that H is transverse to the curve, $p \neq q$, and $p, q \in H \cap C$. We have to show that this is irreducible; then we will get the double transitivity that we claimed. The proof with this incidence correspondence, as always, is to look at it not as a covering of U but as a covering of $C \times C - \Delta$, then it's easy. Consider the map

$$X_U^{(2)} \to C \times C - \Delta, \quad (H, p, q) \mapsto (p, q).$$

The fibers of this map are (Zariski open subspaces of) projective spaces of dimension $r-2$, because once you fix two points, then we have a $\mathbb{P}^{r-2}$ of projective spaces containing the two points and we throw out tangent hyperplanes. So we have an essentially $\mathbb{P}^{r-2}$-bundle over an irreducible surface, which implies $X_U^{(2)}$ is irreducible.

We've used: you have a dominant map, with irreducible fibers of the same dimension. Then the total space is usually irreducible—not always. Well, it's a bit subtle to prove this completely rigorously.

Anyway, the point is that one can deduce multiple transitivity by using this geometric argument.

Now, we claim that G contains a simple transposition. If we know that the group $G \subset S_d$ is twice transitive and contains a transposition, it must be the symmetric group (as it contains all transpositions). So, we have to show that it contains a transposition. We'll do this by *deliberately* going to a point in $C^* \subset \mathbb{P}^{r*}$, a tangent hyperplane, and look at the covering in a neighborhood of that point. The claim is that we can let H wander in a neighborhood of that point and induce a simple transposition in monodromy.

So let H_0 be a hyperplane simply tangent to C. So $H_0 \cap C = 2p + d - 2$ other points . There's contact of order two at one point and transversality everywhere else. (**That's where characteristic zero enters: there are curves in characteristic p such that when a hyperplane is tangent, it is multiply tangent.**) Choose a small analytic neighborhood V of H_0, and we look at the family $X_{V \cap U} \to V \cap U$. The claim is that the cover over this $V \cap U$ is going to consist of $d-2$ copies of that same neighborhood, all disjoint and all mapping isomorphically and one component which is still connected but which maps 2-to-one onto the base. For this purpose, it makes sense to define the incidence correspondence for all hyperplanes, not just transverse ones. So the claim is that there are two sheets such that you can get from one to another by wandering the hyperplane.

Let's say $\overline{X} = \{(H, p) : p \in H \cap C\}$ which maps to $\mathbb{P}^{r*}$. Note that $\overline{X}$ is still smooth, and $d-2$ of those connected components map isomorphically near H_0; one maps two-to-one near H_0. This will therefore induce a transposition in the monodromy.

▲

So, we've proved the uniform position theorem, and thus we've proved the basic estimate. By that very elementary lemma using hyperplanes, we get the bound on Hilbert functions and we have completed the proof of Castelnuovo's theorem.

This is stronger than the version in the book.

Lecture 12
10/14

§1 General position lemma

Before we do anything else, let's go over the last step in the general position lemma.

Let $C \subset \mathbb{P}^r$ be a nondegenerate curve of degree d, and let $H \subset \mathbb{P}^r$ be a general hyperplane. The claim is that $H \cap C$ is in linear general position. Last time, we showed that if you vary H, and as you follow the points of intersection around, the monodromy induced—the group of permutations of these points—is the full symmetric group. From this we want to claim:

12.1 Proposition. *No r points of $H \cap C$ are linearly dependent.*

Proof. Let $U = \mathbb{P}^{r^*} - C^*$ be the family of transverse hyperplanes to the curve. Let us form the incidence correspondence $X = \{(H, p) : p \in H \cap C\} \subset U \times C$; this is a d-sheeted cover of U. We saw that the monodromy acts as the full symmetric group on these d sheets. That is, the fiber product of X, any number of times, with itself (over U), with diagonals removed, will be an irreducible space. This is a consequence of the monodromy result, and in fact is true in nonzero characteristic as well.

OK, so let's use this to prove the linear general position fact. Let Z be the set of pairs $\{(H, p_1, \dots, p_r) : p_1, \dots, p_r \text{linearly independent}, p_i \in H\}$ such that all the p_i are distinct and lie over the same point of U. This is a closed subvariety of $X_U^{(r)}$, the r-times fibered product $X \times_U \times \cdots \times_U X$ with the diagonals removed. This Z is nonempty, because we can choose linearly independent points on the curve. Since $X_U^{(r)}$ is irreducible, to say that $Z \subsetneq X_U^{(r)}$ is to say that $\dim Z < \dim X_U^{(r)}$. In particular, it follows that the image of Z in U is a *proper* subvariety of U, because the map $X_U^{(r)} \to U$ is a quasi-finite map.

Things in U that are outside the image of Z are hyperplanes whose intersection with the curve consist of linearly independent points. ▲

This is a common use of monodromy: we used the fact that the monodromy group is the full symmetric group to conclude that certain varieties were irreducible.

Remark. We proved this for curves. It's true for varieties in arbitrary dimension. If you have a k-dimensional variety X in $\mathbb{P}^r$, then the intersection with a general $r-k$-plane will give points in linear general position. You can deduce that from the statement for curves (because a $r-k$-plane is an intersection of k hyperplanes, and the intersection of X with all but one of them is a curve). Or, one can imitate this proof directly.

So now we have *completely* proved the Castelnuovo bound.

§2 Projective normality

Let's recall the notion of *projective normality.*

12.2 Definition. Let $C \subset \mathbb{P}^r$ be a smooth curve. We say that C is **projectively normal** if, for every m, the hypersurfaces of degree m in $\mathbb{P}^r$ cut out a *complete* linear series on C. In other words, if D is a divisor on C and D is linearly equivalent to m times the hyperplane divisor H (i.e. $D \sim mH$), then D is the intersection of C with a hypersurface.

Or, the restriction map

$$H^0(\mathbb{P}^r, \mathcal{O}_{\mathbb{P}^r}(m)) \to H^0(\mathcal{O}_C(m))$$

is surjective. This is equivalent to saying that $H^1(\mathcal{I}_{C,\mathbb{P}^r}(m)) = 0$ for all m, where $\mathcal{I}_{C,\mathbb{P}^r}$ is the ideal sheaf of $C \subset \mathbb{P}^r$.

What does this mean geometrically? I can't tell you that, but if you're an algebraist—one of those people who read Atiyah-Macdonald and enjoyed the experience—here is an equivalent formulation (which is an exercise in Hartshorne).

12.3 Proposition. *C is projectively normal if and only if the cone*[4] *$S \subset \mathbb{P}^{r+1}$ has the property that the local ring $\mathcal{O}_{S,p}$ at the cone vertex $p \in S$ is a Cohen-Macaulay ring. This, if and only if, S is a normal variety.*

Let's go back to Castelnuovo theory. Consider an irreducible, nondegenerate, smooth curve $C \subset \mathbb{P}^r$. If we don't assume C projectively normal, then we have a *subspace* $H^0(\mathcal{O}_{\mathbb{P}^r}(k))|_C$ of $H^0(\mathcal{O}_C(k))$ consisting of the image of $H^0(\mathcal{O}_{\mathbb{P}^r}(k))$. This is a subseries; we call it E_k. If E_k is always the full linear series, then we have projective normality.

Last time, we proved a basic fact about the E_m. We showed

$$\dim E_k - \dim E_{k-1} \geq \min(d, k(r-1)+1).$$

For large values of n, we were able to bound below $\dim E_n$, and were able to use Riemann-Roch to bound above the genus $g \leq \pi(d,r)$.

The observation is that, if we have equality at the *end*, then we must have equality at every stage. We must have not only

$$\dim E_k - \dim E_{k-1} = \min(d, k(r-1)+1),$$

but we must also have that each E_k is the full linear series (or we would have strict inequalities at some stage).

Thus:

12.4 Proposition. *If $g(C) = \pi(d,r)$, then C is projectively normal.*

So extremal curves with respect to the Castelnuovo bound are projectively normal.

[4] I.e., imbed $\mathbb{P}^r \subset \mathbb{P}^{r+1}$, take a point outside it and draw lines from that point to C.

12.5 Corollary. *If $C \subset \mathbb{P}^{g-1}$ is a canonical curve of genus g, then C is projectively normal.*

This is because canonical curves are extremal.
This means, in particular:

12.6 Corollary. *Let C be non-hyeprelliptic. Every quadratic differential (section of K^2) is a quadratic polynomial in the holomorphic differentials. More generally, any pluricanonical differential on the curve is a polynomial in the holomorphic differentials.*

This is interesting because quadratic differentials are useful in the geometry of, say, Teichmuller theory.

§3 Sharpness of the Castelnuovo bound

We're now going to launch into an extended discussion of a topic whose relevance may seem obscure. Remember what we said we had to do—we have now a bound $g \leq \pi(d, r)$ on the genus of a curve imbedded in projective space. That bound was based on a series of inequalities that we just wrote down, on the differences of the dimensions of the series $\mathcal{O}_C(n)$. That seems very crude, but we saw that it is achieved, e.g. by points on a rational normal curve. That suggests that if we want to look for curves that achieve this bound on the genus—our next goal—we need that for a general hyperplane section Γ of the curve, the Hilbert function of Γ is the minimum $\min(d, k(r-1))$ we gave. (An example of such a configuration Γ is any set of points on a rational normal curve.)

This means, once more, that if C is our curve, and $\Gamma = C \cap H$ is a general hyperplane section, then we want Γ to *lie on a rational normal curve.* So what if $C \subset \mathbb{P}^3$, for instance; then we're saying that we need to look for configurations of points in the *plane* which impose the smallest number of conditions on hypersurfaces. Here we would be looking for points lying on *conics*. So we might ask: can we take a curve in $\mathbb{P}^3$ such that the general hyperplane section consists of points lying on a conic? Sure; just take C lying on a quadric. If C lies on a quadric, then the hyperplane section Γ will lie on a conic, and consequently the Hilbert function of Γ will be exactly $\min(d, k(r-1)+1)$. This does *not*, however, mean that curves on a quadric in $\mathbb{P}^3$ have maximal genus.

12.7 Example. If C lies on a smooth quadric $Q \subset \mathbb{P}^3$, and is of type (a, b), then we have that $a + b = d$ (for d the degree) and $g = (a-1)(b-1)$. The genus is maximized when a, b as close to each other as possible, e.g. to $\frac{d}{2}$. So a general curve on a smooth quadric won't maximize the genus, though one can see that the *best* curve on a smooth quadric will achieve the bound.

If $d = 2k$, then one can check that $g(C)$ can go up to $(k-1)^2$, and for $d = 2k+1$, the genus can go up to $k(k-1)$. These are in fact extremal curves. *So the Castelnuovo bound in $\mathbb{P}^3$ is sharp.* When, however, one works with curves on a smooth which are of the form (a, b) for $|a - b| > 1$, then even though the hyperplane sections have the right Hilbert polynomials, projective normality fails (cf. the homework).

12.8 Example. Curves that lie on an irreducible cubic surface satisfy stronger bounds on the genus (cf. the homework). The largest genus comes when you're on a quadric.

In general, we're looking for curves $C \subset \mathbb{P}^r$ such that $H \cap C$ lies on a *rational normal curve.* We were able to achieve this in $\mathbb{P}^3$ by putting C on a quadric surface, which guarantees that the hyperplane section lies on a conic. We might ask ourselves: is there an analogous surface in $\mathbb{P}^r$? Can we find a surface $S \subset \mathbb{P}^r$ such that $H \cap S$ is always a rational normal curve? (One good suggestion is to take a cone on a rational normal curve; but we need something slightly different.) If we can find one such, and if we can understand what curves on this surface look like, then we might be able to hunt down the curves we want of maximal genus.

We're going to take the next twenty and thirty minutes and talk about a more general topic. At the end of the day, we'll have a complete description of these surfaces. We won't prove most of these assertions. If you want to see proofs, they're in various sources.

§4 Scrolls

We start with:

Question. What is the smallest possible degree of a variety $X \subset \mathbb{P}^r$ which is irreducible, nondegenerate, of dimension k?

For $k = 0$, there isn't much to say.

12.9 Example. Let's consider $k = 1$; we're looking for nondegenerate curves in $\mathbb{P}^r$ of smallest degree. The answer is that the degree is at least r,

$$\deg C \geq r.$$

and achieved for the rational normal curve. You can't have anything with smaller degree—if you have a curve in $\mathbb{P}^r$, then r general points will span a hyperplane H; if the degree of the curve were $< r$, then $C \subset H$ by Bezout, contradicting nondegeneracy.

12.10 Example. Let $k = 2$. We're looking at surfaces in $\mathbb{P}^r$. The claim is that an irreducible, nondegenerate surface S in $\mathbb{P}^r$ has the property that a general hyperplane section will *again be irreducible and nondegenerate.* I.e., if H is a general hyperplane, then $S \cap H$ will be a $\deg S$ curve in $\mathbb{P}^{r-1}$ which is irreducible and nondegenerate, so we find that

$$\deg S = \deg S \cap H \geq r - 1.$$

12.11 Example. In general, given a variety $X \subset \mathbb{P}^r$ of dimension k, we can use induction and the above type of argument to see that

$$\deg X \geq r - k + 1.$$

It turns out that this bound is achieved. If you don't care about smoothness, then you can see that this bound is achieved by taking cones on a rational normal curve repeatedly.

Let's now give an explicit construction of extremal surfaces in $\mathbb{P}^r$, i.e. those of degree $r - 1$. Here's the construction. In $\mathbb{P}^r$, take a direct sum decomposition of the corresponding vector space; i.e., choose completementary linear spaces $\mathbb{P}^a, \mathbb{P}^b \subset \mathbb{P}^r$. These *span* $\mathbb{P}^r$ and are disjoint. This implies that $a + b + 1 = r$. For instance, we could take two skew lines in $\mathbb{P}^3$; a less simple example is a line and a 2-plane in $\mathbb{P}^4$. Next, *choose* a rational normal curve in each of these two subspaces; call these $C_a \subset \mathbb{P}^a, C_b \subset \mathbb{P}^b$. Choose an isomorphism of C_a with C_b (which we can do; both are isomorphic to $\mathbb{P}^1$). Call this isomorphism ϕ. Finally, take the surface S to be the *union* of the lines joining p to $\phi(p)$ for $p \in C_a$, i.e.

$$S = \bigcup_{p \in C_a} \overline{p, \phi(p)}.$$

For instance, we could draw a conic in $\mathbb{P}^2 \subset \mathbb{P}^4$, and connect corresponding points on this conic to a line $\mathbb{P}^1$.

12.12 Definition. These are called **rational normal (surface) scrolls.** They are sometimes denoted $S_{a,b}$.

12.13 Proposition. *These surfaces S just constructed have the minimal degree.*

Proof. There are many ways to calculate the degree of this surface; we'll do it in a concrete way. Choose a general hyperplane H containing *one* of those two linear subspaces $\mathbb{P}^a, \mathbb{P}^b$. What is $H \cap S$? Well, obviously $H \cap S$ will *contain* the curve C_a, but there's more. The hyperplane intersects $\mathbb{P}^b$ in a hyperplane in $\mathbb{P}^b$, and that intersection will intersect the rational normal curve in d points. For each point $p \in C_b$ that lies in H, the line through p and its image in C_a is also in H (because $C_a \subset H$). So, H contains the *lines* through the lines in $C_b \cap H$. It's easy to see that there are no other points in the intersection.

We find that $H \cap S$ is the union of C_a and lines through b points. The degree of this is $a + b$; consequently by Bezout,

$$\deg S = a + b.$$

▲

To be careful, we should check that this calculation is legitimate, i.e. that the intersections occur with multiplicity one.

Remark. Suppose instead of choosing an isomorphism of the two curves, we chose a parametrization $\phi : \mathbb{P}^1 \simeq C, \phi' : \mathbb{P}^1 \simeq C'$. Then we could write $S = \bigcup_{t \in \mathbb{P}^1} \overline{\phi(t)\phi'(t)}$. That's just another way of expressing it, but now we have an option not expressible before; we could take one of $a, b = 0$ (in which case one of the ϕ, ϕ' would be zero). We'd just get a *cone* over the nonconstant one of C_a, C_b. This, depending on your preferences, is also a scroll.

12.14 Theorem (Classical)**.** *The scrolls (including cones) are the surfaces of minimal degree in $\mathbb{P}^r$ with one exception: that is, an irreducible, nondegenerate surface $S \subset \mathbb{P}^r$ of degree $r - 1$ is either a scroll or the Veronese surface (isomorphic to $\mathbb{P}^2$) in $\mathbb{P}^5$.*

We state this without proof. Note also that the Veronese surface contains no lines, and is very far from being a scroll. There is a generalization of this to more general varieties of minimal degree.

Remark. The projective equivalence class of $S = S_{a,b}$ only depends on (a, b), because the choices we made don't really amount to much (a different choice, with a given (a, b), would be conjugate by automorphisms of projective space). Conversely, in a given $\mathbb{P}^r$, no two of these $S_{a,b}$ are projectively isomorphic; each is distinguished by the presence of a curve of degree $\min(a, b)$ on it.

Here are two other things about the scrolls. These are both going to require some thought, but they'll say something about them.

12.15 Proposition. *Each surface scroll S is the projectivization*[5] $\mathbb{P}(\mathcal{O}(a) \oplus \mathcal{O}(b))$ *of the vector bundle $\mathcal{O}(a) \oplus \mathcal{O}(b)$ on $\mathbb{P}^1$. The embedding is given by the tautological bundle on the projectivization.*

One thing this result tells us which is worth remembering is that:

12.16 Corollary. $S_{a,b} \simeq S_{a',b'}$ *if and only if* $|a - b| = |a' - b'|$.

This is because tensoring a vector bundle by a line bundle doesn't affect the projectivization.

Next, let's recall that a rational normal curve is a *determinantal variety.* In $\mathbb{P}^r$, if you have a 2-by-r matrix of linear forms that looks like

$$A = \begin{bmatrix} L_1 & \dots & L_r \\ M_1 & \dots & M_r \end{bmatrix}$$

which are *general*, then the locus Φ of points $p \in \mathbb{P}^r$ such that this matrix $A(p)$ has rank one, then this locus is a rational normal curve. The claim is that we can do something similar and construct the $S_{a,b}$. OK, let's still take A a matrix,[6] but one column short:

$$A = \begin{bmatrix} L_1 & \dots & L_{r-1} \\ M_1 & \dots & M_{r-1} \end{bmatrix}$$

and consider the locus of points $p \in \mathbb{P}^r$ where the rank is one. This is a scroll—we won't justify it, but it's useful. For instance, what does it mean for a two-by-whatever matrix to have rank one? It means that the rows are independent. If we specified a specific *relation* of linear independence, i.e. considered a pair $[\alpha, \beta] \in \mathbb{P}^1$ and considered the equations $\alpha L_i + \beta M_i = 0, 1 \leq i \leq r - 1$, then this is a *line* $\ell_{\alpha,\beta}$. As we vary $[\alpha, \beta]$, we get a *family* of lines parametrized by $\mathbb{P}^1$. This is in fact the scroll. (Determining the a, b in the scroll appears to be hard. The generic behavior is balanced.)

12.17 Example. Let $X_0, \dots, X_r$ be the linear coordinates on $\mathbb{P}^r$. Consider the matrix

$$\begin{bmatrix} X_0 & X_1 & \dots & X_{a-1} & X_{a+1} & \dots & X_{r-1} \\ X_1 & X_2 & \dots & X_a & X_{a+2} & \dots & X_r \end{bmatrix};$$

this is of type $(a, r - a + 1)$.

[5] This is the post-Grothendieckian construction. That is, the projectivization of a vector space is the set of hyperplanes in it.

[6] Technically, this should be 1-generic: this means no entry is zero, and no single row or column operation will make an entry zero.

Ultimately, we are going to be talking about *curves* lying on these surfaces S. So we'll need a similar description of curves on a scroll, just as we had a nice description (in terms of a pair of integers) of curves on a smooth quadric, that will enable us to determine the genus. So, in other words, we need the intersection theory on a $S_{a,b}$. They're not hard to derive, but now is not the time to do that. We'll do that on *Monday.*

Let's just say one more thing about these scrolls. We started out this whole discussion with a simple question on varieties of minimal degree. Here's another equally simple question:

Question. Let $X \subset \mathbb{P}^r$ be a k-dimensional, irreducible, nondegenerate, variety of dimension k. We have asked in the past how many hypersurfaces a curve lies on. So, what is the maximal number of (linearly independent) quadrics can X lie on?

Once we know an upper bound on the number of quadrics containing a variety, we might also try to characterize the extremal ones.

In both cases, we are talking in some sense of the *size* of a variety. In some sense, the degree of a variety is a measure of its size; another is the number of hypersurfaces of some degree it lies on. We won't prove this, but in fact:

12.18 Proposition. *The maximal number of quadrics is* $\binom{r+1-k}{2}$.

Proof. OK, let $Y = H \cap X$ be the intersection of X with a general hyperplane H. We have a short exact sequence of sheaves (where $\mathcal{I}$ denotes the ideal sheaf)

$$0 \to \mathcal{I}_{X,\mathbb{P}^r}(1) \to \mathcal{I}_{X,\mathbb{P}^r}(2) \to \mathcal{I}_{Y,\mathbb{P}^{r-1}}(2) \to 0$$

where the first thing is elements of the ideal sheaf vanishing on H, and we find that

$$H^0(\mathcal{I}_{X,\mathbb{P}^r(2)}) \hookrightarrow H^0(\mathcal{I}_{Y,\mathbb{P}^{r-1}(2)}).$$

By induction on the degree (and an analysis for points in $\mathbb{P}^2$), we can obtain the bound claimed. ▲

Among surfaces, the extremal ones are the Veronese surface and the scrolls. One can see that scrolls lie on a lot of quadrics by representing them as a determinantal variety, and looking at 2-by-2 minors.

Lecture 13
10/17

§1 Motivation

From last time, let's recall: $C \subset \mathbb{P}^r$ is a curve, smooth and irreducible and nondegenerate of degree d. We proved a bound

$$g(C) \leq \pi(d,r) \equiv \binom{m}{2}(r-1) + m\epsilon$$

where $d-1 = m(r-1)+\epsilon$ such that $0 \le \epsilon \le r-2$. We want to find curves of *maximal genus.* To do this, (cf. the proof of the Castelnuovo bound) we want to maximize the Hilbert function of the intersection $C \cap H$ with a generic hyperplane section H. We know that a configuration of points lying on a rational normal curve has a large Hilbert function, so we want to find C such that *every hyperplane section lies on a rational normal curve.*

Let $d > 2r-2$. Note also that any extremal C is projectively normal. This means that any quadric in H containing a hyperplane section $H \cap C$ *lifts* to a quadric containing the curve C. The quadrics $Q \subset H = \mathbb{P}^{r-1}$ containing the hyperplane section $\Gamma = H \cap C$ cut out a rational normal curve in $\mathbb{P}^{r-1}$ (because the degree is $> 2r-2$), so the quadrics in $\mathbb{P}^r$ containing C cut out a *surface* whose hyperplane section is that rational normal curve.

So we're looking for *curves* on a surface S of minimal degree. This will be, with one intersection, a scroll.

§2 Scrolls again

Now we need to describe the geometry of scrolls and curves on scrolls. To construct a scroll, depending on parameters $a \le b \in \mathbb{Z}$, we take $\mathbb{P}^r$, two complementary subspaces $\mathbb{P}^a, \mathbb{P}^b$, and draw rational curves $C_a \subset \mathbb{P}^a, C_b \subset \mathbb{P}^b$, choose an isomorphism between the two $C_a \simeq C_b$, and draw lines between corresponding points. One can show that one in fact gets a smooth surface $S_{a,b}$, whose degree must be $a+b$.

13.1 Definition. The curve $C_a \subset S_{a,b}$ is called the **directrix.**

Let us describe a few facts.

If you take a scroll, and project from a line on the ruling, then you project to $\mathbb{P}^{r-2}$: lines map to lines, C_a maps to a rational normal curve of degree $a-1$, and C_b maps to a rational normal curve of degree $b-1$. So, if you project from a line of the ruling, you *get an isomorphism of $S_{a,b}$ with $S_{a-1,b-1}$*. If you projected from a point of C_a or C_b, then you would only get a *rational* map. Projecting from a Cartier divisor gives a regular map.

Consequently, $S_{a,a} \simeq S_{1,1} \simeq \mathbb{P}^1 \times \mathbb{P}^1$, because $S_{1,1}$ is a quadric surface. Moreover, $S_{a,a+1}$ is isomorphic to a cubic scroll in $\mathbb{P}^4$. This is, moreover, isomorphic to the blow-up of $\mathbb{P}^2$ at a point p. Why is this? If you blow up a point in $\mathbb{P}^2$, then you replace a point in $\mathbb{P}^2$ with an exceptional divisor, and the lines passing through that point become disjoint. If you choose a second line in the plane, not passing through the blown up point, it's exactly the same picture: it follows that the blown up surface is swept out by lines joining the two, disjoint. More formally, consider the space of conics in $\mathbb{P}^2$ through the point p; this gives a rational map from $\mathbb{P}^2 \to \mathbb{P}^4$, which becomes regular on the blow-up, and imbeds the blow-up in $\mathbb{P}^4$.

§3 Intersection numbers

OK, let's go back to where we were. We have a scroll $S_{a,b}$ obtained by joining corresponding points on curves in complementary projective spaces. We can think of this as swept out by lines, and it's a bundle over $\mathbb{P}^1$. The fibers are these lines. Let H be a hyperplane divisor in $S_{a,b}$.

We now want to understand the Picard group of $S_{a,b}$ and the intersection pairing. This is a straightforward exercise. Here is an outline:

13.2 Proposition. $\mathrm{Pic}(S_{a,b})$ *is freely generated by the classes* f *(the class of a fiber, i.e. the lines in the ruling) and* h *(the class of a hyperplane* H*).*

Proof. In fact, we just have to make the following observation. If we take $S_{a,b}$ and remove a fiber and remove a hyperplane section of S, what we have left is a bundle of affine lines over an affine space—so it's $\mathbb{A}^2$, which has trivial Picard group. Thus the Picard group must be (by the standard exact sequence, or because the triviality of $\mathrm{Pic}(\mathbb{A}^2)$ means that divisors can be "pushed" onto the union of h and f) generated by h and f. Consequently, the rank of the Picard group is at most 2.

More generally, if you're familiar with projective bundles, you can appeal to a general fact about the Picard groups of them.

OK, let's compute intersection products. **We will use these numbers below.** The intersection $h.h$ is the intersection of two hyperplanes, which is the *degree* of the surface $r-1$. We also easily get $h.f = 1$ and $f.f = 0$ (two fibers are disjoint, and a hyperplane class hits f, which is a line, once). This also shows that h, f are linearly independent in the Picard group, which must thus be $\mathbb{Z}^2$.

▲

We could also compute the classes of C_a, C_b in $S_{a,b}$. A general hyperplane H containing C_a intersects C_b in a total of b points and consequently $C_a \cap H$ will be the union C_a and b lines of the ruling. So the divisor of C_a is $h - bf$ and that of C_b is $h - af$.

Note that the self-intersection of C_a is $a - b \leq 0$, while $C_b.C_b = b - a \leq 0$. The two curves really behave differently from each other. In fact, C_a is the *unique* curve of degree $\leq a$ (other than the lines) on this surface if $a < b$.

Finally, we need to know the canonical class of $S_{a,b}$. This is a natural thing to ask when you work out the Picard group. There are probably better abstract ways of doing this. For instance, you could write down an explicit 2-form by writing down equations and then count zeros and poles. But that's too much like work, and "cheat" (it's not cheating; it's legitimate)

$$K_S = \alpha h + \beta f$$

for some $\alpha, \beta \in \mathbb{Z}$. How can we find α, β?

Let's apply adjunction to the curves we *know about* on the surface. Let's apply it to the fiber f itself, which is a fiber: we know that f is a line, which has genus zero. So the canonical class of f has degree -2. But the canonical class of f is

$$(K_S + f)|_f,$$

so the degree is $(\alpha h + \beta f).f = \alpha$. It follows that $\alpha = -2$. Now, we can do it one more time to figure out β; let's take a hyperplane class h.

The hyperplane section is a rational normal curve. So we do this again, and have

$$-2 = (K_S + h).h = (\alpha h + \beta f).h + r - 1 = \beta - (r-1).$$

Thus $\beta = r - 3$.

We find:

13.3 Proposition. *The canonical series* $K_S = -2h + (r-3)f$.

That's why this basis for the Picard group is nice; note that the answer doesn't depend on a, b.

Remark. You should check that this is consistent with applying adjunction to C_a, C_b and the self-intersection of these curves. You should *immediately* check this in the case where $r = 3$, where the scroll is a quadric is $\mathbb{P}^3$, and where adjunction in $\mathbb{P}^3$ gives the canonical class of the scroll, which we found earlier to be $-2h$. This is just what the above analysis gave.

Now that we know the canonical class and the intersection products, we can say exactly what the genus of a curve on the scroll is. Let $C \subset S_{a,b} \subset \mathbb{P}^r$ is a scroll. As a divisor in S, we can write $C = \alpha h + \beta f$; this means that α is the number of times C meets a line of the ruling. What we'd like to is to compute the *degree* and *genus* of C via α, β.

1. The degree of C in $\mathbb{P}^r$ is the intersection of $C = \alpha h + \beta f$ with the hyperplane class h, which is $\alpha + (r-1)\beta$ by the earlier basis and the intersection numbers computed.

2. The genus of C can be found via adjunction. Namely, we can get it from

   $$C.(K_S + C) = (\alpha h + \beta f).((\alpha-2)h + (\beta + r - 3)f) = \alpha(\alpha-2)(r-1) + \beta(\alpha-2) + \alpha(\beta + r - 3).$$

   This can be simplified to

   $$\alpha(\alpha-1)(r-1) + 2\alpha\beta - 2\alpha - 2\beta.$$

   The conclusion is that

   $$g(C) = \binom{\alpha}{2}(r-1) + (\alpha-1)(\beta-1).$$

Note that we're using a different basis for the Picard group that we would use for the quadric in $\mathbb{P}^3$.

OK, now we want to find extremal curves, i.e. choose α, β appropriately to maximize the genus.[7] For a *fixed value of the degree, we want to find the maximal value of the genus by the above answer.* You can assign it to your calculus class—you'll see what you get. We'll just state the answer.

Answer: To maximize g, take

$$\alpha = m + 1, \quad \beta = \epsilon - r + 2.$$

The degree in this case is $m(r-1) + \epsilon + 1$, which is what you want. The genus works out to be $\binom{m}{2}(r-1) + m\epsilon$. This in fact gives the right answer; we saw this was the upper bound. (If $\epsilon = 1$, we can also take $\alpha = m, \beta = 1$.)

Is it clear that any such divisor class is represented by a smooth curve on a scroll? Let's answer this question. Here's a little lemma, worth pointing out.

[7] We have to show that a given divisor class contains a smooth curve, though.

13.4 Lemma. *The class $\alpha h + \beta f$ on $S_{a,b}$ is represented by a smooth (not necessarily irreducible) curve on the $S_{a,b}$ if $\alpha = 0, \beta \geq 0$ (this is a collection of lines on the curve) or if $\alpha > 0$, and $\beta \geq -\alpha a$.*

The only if case is more complicated and has more cases.

Proof. Use Bertini's theorem. This states if you have a linear series of divisors without base points, then the general member of that linear series is smooth. That applies in this case, because $\alpha h + \beta f = \alpha(C_b) + (\beta + a\alpha)f$ (this is because $C_b = h - af$). If the second term is nonnegative, one can check that the linear series has no basepoints; for this one needs to check that C_b moves in a linear series without basepoints. Thus among the curves in this class are smooth curves. ▲

To show that you can get these curves to be irreducible, you need a further argument with intersection numbers: if you had two curves whose classes added up to $\alpha h + \beta f$ in this way, they would have to intersect, which can't happen.

Bertini's theorem is frequently used to show that a family has smooth divisor classes by exhibiting non-smooth elements of a divisor class.

Remark. In case $r = 5$, and d is even, then we can take C to lie on a Veronese surface. The Picard group of a Veronese surface (which is isomorphic to $\mathbb{P}^2$ is rank one) and C is the image of a plane curve under the Veronese imbedding. So you can check that taking images under the Veronese of plane curves can also be used to satisfy the Castelnuovo bound. We won't do this.

Now what? We've established the maximum possible genus of a curve in $\mathbb{P}^r$, which is cool, but we want to know more. What comes next is a really remarkable theorem. It gives a sort of converse. What we're going to prove is, you get in one of these ways, *all* curves of maximal genus.

We proved, first of all, an inequality on the Hilbert function of a number of points in linear general position in $\mathbb{P}^r$. By applying that to hyperplane sections, we estimated $H^0(C, \mathcal{O}_C(m))$ from below for a curve C in $\mathbb{P}^r$. We saw that the inequality on the Hilbert function could be satisfied by points on a rational normal curve. Now we are going to state something much stronger. We'll give an indication of the proof, and do it properly on Wednesday.

13.5 Theorem (Castelnuovo's lemma). *If $\Gamma \subset \mathbb{P}^n$ is a configuration of $d \geq 2n+3$ points in linear general position, which impose only $2n+1$ conditions on quadrics,*[8] *then Γ lies on a rational normal curve.*

We'll go through the proof on Wednesday, and describe why it's so important. Basically, it's going to tell us that a curve of maximal genus (and degree at least $2n+3$) lies on a rational normal scroll or Veronese surface, because the only way to get the hyperplane sections to lie on a rational normal curve is (as we saw earlier—in Jefferson) is to have the curve lie on a surface.

[8] We saw that the number of conditions imposed on quadrics is at least the minimum of the degree and $2n+1$, and here that minimum is $2n+1$.

The idea of the proof will be a beautiful, old-fashioned idea called the **Steiner construction.** Let's say we have $\Lambda_1, \ldots, \Lambda_n \simeq \mathbb{P}^{n-2}$, codimension two subspaces in $\mathbb{P}^n$. Let $\left\{H_t^i\right\}_{t\in\mathbb{P}^1}$ be the pencil of hyperplanes containing the plane Λ_i. For instance, in case $n = 2$, we'd be looking at pencils of lines containing points. Consider the intersection of the hyperplanes in each of these pencils, i.e.

$$H_t^1 \cap \cdots \cap H_t^n$$

of the n hyperplanes parametrized by the same parameter. In general, this intersection should be a point. Let's suppose this is true; it is true generally. Then, for each value of t, we can take the *union* of these points $\bigcup_t H_t^1 \cap \cdots \cap H_t^n$, which will be a *rational normal curve* in $\mathbb{P}^n$. (*Exercise.*)

In fact, we can give a slightly stronger statement. Instead of taking n pencils of hyperplanes, let's say we take $l < n$ pencils of hyperplanes H_t^i for $1 \leq i \leq l$ and $t \in \mathbb{P}^1$. Again, we consider the variety swept out by the intersections $\bigcap_i H_t^i$ over all t; we assume that these intersections are dimension $n - l$. So we get a variety swept out by $\mathbb{P}^{n-l}$'s, which in fact will be a *rational normal scroll* of dimension $n - l + 1$. This is not obvious.

Proof. Observe that the map $\phi : \mathbb{P}^1 \to \mathrm{Gr}(n-l, n)$ sending $t \mapsto \bigcap_i H_t^i$ is given by the wedge product of the linear forms $\mathbb{P}^1 \to \mathrm{Gr}(n-1, n)$ that describe each H_t^i. So this map ϕ has degree l (when the Grassmannian is imbedded in some projective space). In general, when you have a map to the Grassmannian, the variety swept out is equal to the degree of this map. So it follows that the degree of the variety X swept out by these intersections is degree l, and consequently one sees that it has minimal degree and must be a scroll (being swept out by linear spaces, it's not a Veronese surface). ▲

This is another old-fashioned construction of the scroll, which we'll use to prove Castelnuovo's lemma.

Lecture 14
10/19

Last time, we were about to embark upon a project to completely characterize curves of maximal genus. We had derived Castelnuovo's inequality, which gives an upper bound for the genus of a curve in projective space. By doing intersection theory on scrolls, we showed that the bound was sharp. Now we want to do something even stronger—we want to characterize these extremal curves. The key element was this lemma of Castelnuovo, which we're about to prove.

§1 A basic lemma

We first need:

If $\Lambda_1, \ldots, \Lambda_k$ is a collection of $\mathbb{P}^{n-2}$'s in $\mathbb{P}^n$ (a bunch of codimension two planes, e.g. points in $\mathbb{P}^2$ or lines in $\mathbb{P}^3$) where $k \leq n$, consider the *pencil of hyperplanes* containing each of these hyperplanes containing them.[9] Let $\{H_t^i\}_{t\in\mathbb{P}^1}$ be the pencil of hyperplanes containing Λ_i. For each value of t, consider the intersection $H_t^1 \cap \cdots \cap H_t^k$; assume that this is a transverse intersection always, so the intersection is a $\mathbb{P}^{n-k}$ at all times.

Consider the union $X = \bigcup_{t\in\mathbb{P}^1} H_t^1 \cap \cdots \cap H_t^k$. This should sweep out some variety X.

14.1 Example. If you have two lines in $\mathbb{P}^3$, say ℓ_1, ℓ_2, then we're taking the set of planes containing each of the two lines, pairwise intersecting them, and taking the surface swept out by these intersections. But this corresponds to joining corresponding points on the two lines, so gives a scroll.

14.2 Lemma. *In general, X is always a rational normal scroll.*

Proof. Consider the map $\phi : \mathbb{P}^1 \to \mathbb{G}(n-k, k)$ (the Grassmannian in projective space) sending $t \mapsto H_t^1 \cap \cdots \cap H_t^k$. What does this map look like? If we dualize and think of it as $\mathbb{G}(k-1, n^*)$ of $k-1$-planes in the dual space, then each pencil of hyperplanes is given by a *linear form* varying linearly with t. So $H_t^i = V(\ell_t^i)$ for ℓ_t^i a linear form varying linearly with t. Then the map ϕ sends $t \mapsto V(\ell_1^t \wedge \cdots \wedge \ell_k^t)$; this is thus a curve of degree k under the Plücker embedding. That means (we'll just quote this) that the corresponding X has degree k in $\mathbb{P}^n$, and since this is minimality, we get that X is a scroll. ▲

We're invoking way too much here. Let's say it alternatively. If we write H_t^i as the zero locus of a linear combination $H_t^i = V(t_0 L^i + t_1 M^i)$ for L^i, M^i linear forms, then $\bigcap_i H_t^i$ is thus the locus where the matrix

$$\begin{bmatrix} L^1 & \dots & L_k \\ M^1 & \dots & M_k \end{bmatrix}$$

is such that t_0 times the first row plus t_1 times the first row vanishes, and as t varies, this sweeps out a scroll (where this matrix has rank ≤ 1).

We need another little lemma.

14.3 Lemma. *In $\mathbb{P}^n$, if $p_1, \dots, p_{n+3}$ are points in linear general position, then $p_1, \dots, p_{n+3}$ lie on a rational normal curve.*

For instance, five points in $\mathbb{P}^2$, no three of which are collinear, lie on a smooth conic. In $\mathbb{P}^3$, we find that six points, no four of which lie in a plane, lie on a twisted cubic. We want to use the above construction to give such examples.

Proof. Consider the first n points $p_1, \dots, p_n$. Take the planes in the above construction to be the $\mathbb{P}^{n-2}$'s spanned by $n-1$ of the above points. So,

$$\Lambda_i = \overline{p_1, \dots, \hat{p_i}, \dots, p_n}$$

so this is the span of $p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_n$. Then, choose the parametrization of each pencil H_t^i containing Λ_i such that the element of this pencil that passes through p_{n+1} has $t = 0$, and similarly for the others. That is, and more precisely, we require

$$p_{n+1} \in H_0^i, \quad p_{n+2} \in H_\infty^i, \quad p_{n+3} \in H_1^i.$$

[9] Because they're codimension two, each lies on a pencil of hyperplanes.

We can do this by reparametrizing.

Now, if we take $X = \bigcup_{t\in\mathbb{P}^1}\bigcap_i H_t^i$, then this contains p_1—that's because p_1 will lie on $\Lambda_2, \dots, \Lambda_n$ and thus on *all* the $H_t^2, \dots, H_t^n$, and it will lie on *some* H_t^1. So $p_1 \in X$. Similarly all the $p_i \in X$ for $1 \leq i \leq n$, and the last three do as well by the above construction (taking $t = 0, 1, \infty$). But by the above lemma, this X is a rational normal curve.

This is called *Steiner's construction.* ▲

§2 Castelnuovo's lemma

Now let's restate and prove Castelnuovo's lemma.

14.4 Theorem (Castelnuovo)**.** *Consider $\Gamma = \{p_1, \dots, p_d\} \subset \mathbb{P}^n$. Let's make the following assumptions:*

1. *Γ is in linear general position.*
2. *$d \geq 2n+3$.*

If $h_\Gamma(2) = 2n+1$,[10] *then Γ is contained in a rational normal curve.*

By way of an example, let's see why this is necessary.

14.5 Example. Take $n = 3$. So we have $\mathbb{P}^3$, and let's say you have eight points in linear general position. Can they impose seven conditions on quadrics? Since the space of quadrics is ten-dimensional, this is to say that the eight points lie on a *net* of quadrics. The answer is yes—just take three quadrics, and pick the intersection. In other words, the transverse intersection of three quadrics will be eight points, generically in linear general position, which will impose only seven conditions. They don't lie on a rational normal curve because then the rational normal curve would lie on each of the quadrics.

14.6 Example. If we have *nine* points in $\mathbb{P}^3$, and they lie on three quadrics, then by Bezout the intersection of three quadrics must be positive-dimensional. Ultimately one sees that the intersection is a twisted cubic (after some work).

Proof. Again, let Λ_i be $\overline{p_1, \dots, \hat{p_i}, \dots, p_n}$ be the span of the first n points with the ith omitted. Choose the parametrizations of the pencils H_t^i (for $t \in \mathbb{P}^1$) of hyperplanes containing Λ^i such that

$$p_{n+1} \in H_0^i, \quad p_{n+2} \in H_\infty^i, \quad p_{n+3} \in H_1^i, \quad \forall i.$$

We know that the rational normal curve X swept out by $\bigcap_i H_t^i$ containing these contains $p_1, \dots, p_{n+3}$, as in the earlier argument today. We want to see that this contains them all.

OK, take $\Lambda = \overline{p_{n+4}, \dots, p_{2n+2}}$. Take again $\{H_t\}$ as the pencil of hyperplanes containing Γ. Again, we normalize so that

$$p_{n+1} \in H_0, \quad p_{n+2} \in H_\infty, \quad p_{n+3} \in H_1.$$

[10] This is the Hilbert function.

Let $Q_i = \bigcup_{t\in\mathbb{P}^1} H_t^i \cap H_t$. This is a union of codimension two planes, and it's varying with one degree of freedom. It has codimension one. The claim is that Q_i is a quadric; in fact, it's a scroll of codimension one, so it's a quadric hypersurface by the earlier lemma today. It contains a bunch of points. It contains the points in Λ^i, so it contains $p_1, \dots, \hat{p_i}, \dots, p_n$, and it also contains $p_{n+1}, p_{n+2}, p_{n+3}$, and also $p_{n+4}, \dots, p_{2n+2}$. If we count these up, we get $2n+1$ points of the configuration Γ. If any $2n+1$ points of Γ impose independent conditions on quadrics, then any quadric containing $2n+1$ points of Γ must contain Γ.

Now we're done, because all the remaining points have to match up. Well, that says in turn that for all points from $2n+3$ onto d, that is for p among $p_{2n+3}, \dots, p_d$ we find that the value of t such that H_t^i contains p is equal to the value of t such that H_t contains p. (That's because all the remaining points lie on the quadric Q_i.) This means in particular that the value of t such that H_t^i contains p is also the same as the value of t such that H_t^j contains p. *That* implies in turn that $p_{2n+3}, \dots, p_d$ all lie on X.

We still have to check that $p_{n+4}, \dots, p_{2n+2}$ lie on X; by construction they lie on each quadric, but they don't obviously match up. However, since there's at least more point, you can swap it and use the same argument: we know by know that X contains $p_1, \dots, p_{n+3}$ and there is a *unique* rational normal curve containing those. You can swap points to get the same rational curve. ▲

(Being tired, I didn't really follow this, and this argument may not be faithfully written down.)

§3 Equality in the Castelnuovo bound

Let $C \subset \mathbb{P}^r$ be a nondegenerate curve of degree d. Then $g(C) \leq \pi(d,r)$. If equality holds there, i.e. $g(C) = \pi(d,r)$ and $d \geq 2r+1$, then the generic hyperplane section $\Gamma = C \cap H$ must lie on a rational normal curve in $\mathbb{P}^{r-1}$ by the previous lemma of Castelnuovo. The intersection of quadrics in $\mathbb{P}^{r-1}$ containing Γ must be that rational normal curve, and since those quadrics in $\mathbb{P}^{r-1}$ lift to quadrics in $\mathbb{P}^r$ containing the curve (by projective normality).

So we have used two facts:

1. The intersection of quadrics in $\mathbb{P}^{r-1}$ containing Γ is this rational normal curve B in $\mathbb{P}^{r-1}$ that the points lie on.

2. The intersection of quadrics in $\mathbb{P}^r$ containing C must be a surface of minimal degree, which is thus either a Veronese surface or a scroll.

Anyway, we have seen that curves on a scroll or on a Veronese surface can achieve this bound.

We now have an exact list of curves of maximal genus.

14.7 Theorem. *A curve in $\mathbb{P}^r$ of degree d meeting the Castelnuovo bound and that satisfies $d \geq 2r+1$ lies on a Veronese surface or a scroll.*

Remark. A canonical curve does not satisfy $d \geq 2r+1$.

Let's draw one interesting consequence. If you take the quadrics containing such a curve, you'll get a scroll. Let's ask the same question with a canonical curve.

Take a canonical curve $C \subset \mathbb{P}^r$ with $g = r - 1$ and $d = 2g - 2 - 2r$. The above doesn't apply, but we might consider the intersection over quadrics $X = \bigcap_{C \subset Q \subset \mathbb{P}^r} Q$; this is an intersection of quadrics which contains C. Suppose X contains one more point $p \notin C$.

Take a general hyperplane H containing that point p. Consider the corresponding plane section $H \cap C$: there are $2r$ points of intersection $p_1, \dots, p_{2r}$. You'll see that in this case, if you review the proof of the uniform position lemma, the points of a hyperplane section of a curve *together with the point* p, are in linear general position. That is, $p_1, \dots, p_{2r}, p$ are in linear general position. That's not obvious. The first $2r$ are in linear general position by the general monodromy argument. Then you have to argue that p together with any $2r-1$ points among $p_1, \dots, p_{2r}$ will still give you a large monodromy group. That's a crucial observation.

Now, every quadric in projective space containing the curve C contains p. It follows that any quadric in H containing $p_1, \dots, p_{2r}$ contains p. It follows that $\{p_1, \dots, p_{2r}, p\}$ impose only $2r - 1$ conditions on quadrics. We have enough points that Castelnuovo tells us that $p_1, \dots, p_{2r}, p$ all lie on a rational normal curve. Now the same logic goes through as before. It follows that the intersection of all quadrics in projective space is such that the general hyperplane section is a rational normal curve, by the same logic. So this intersection of quadrics, if it's not C, is a scroll or Veronese surface.

We have proved:

14.8 Proposition. *A canonical curve lies on a scroll or a Veronese surface, or it's the intersection of quadrics.*

One can see that the class of the curve is $3h - (r-3)f$.

Remark. This means that the canonical curve will meet the lines of the ruling three times. In other words, the lines of the ruling will cut out a g^1_3 on the curve. This implies that *a canonical curve lying on a scroll is trigonal.* Conversely, if we start with a trigonal curve and look at the canonical model, the geometric Riemann-Roch theorem shows that there are triples of collinear points lying on the points. Every quadric containing the canonical model must contain the line through these points. Then in that case the canonical model lies on the scroll and the lines of the ruling cut out the g^1_3.

If C lies on a Veronese surface and is a canonical curve, then C must be the image of a smooth plane quintic.

We find:

14.9 Theorem (Enriques)**.** *If $C \subset \mathbb{P}^{g-1}$ is a canonical curve, then C is cut out by quadrics containing it unless C is either trigonal or C is a plane quintic.*

This is saying something genuinely interesting about the geometry of canonical curves. Once you get beyond degree $2r$ in Castelnuovo theory, the curves you're looking at of maximal genus are very special. This argument, as it stands, tells us nothing except about curves about that single maximum genus. It says nothing about curves of genus one less at the maximum. The homework assignment will hint at these questions. It's interesting to ask how much of this analysis applies to curves of *relatively* high genus.

We have a sequence of logical steps: we're looking for a curve of largest possible genus. A curve of higher genus has a smaller Hilbert polynomial, so it lies on more hypersurfaces of large degree. To find curves of largest degree, Castelnuovo found collections of points with minimal Hilbert function. These were points on rational normal curves; this is reasonable, because a rational normal curve is a curve minimizing the Hilbert function.

Now, suppose you want to look for configurations of points that have a small, but not really small, Hilbert function—a little larger than the minimum possible Hilbert function. We saw that if Γ is in linear general position, then $h_\Gamma(m) \geq \min(d, m(r-1)+1)$ for d the degree. What's another set of points, not lying on a rational normal curve, that is close to this? One example: points lying on an *elliptic* curve. Or a projection of a rational normal curve. It turns out the elliptic curve wins.

If Γ is contained in an elliptic normal curve, then

$$h_\Gamma(m) \sim \min(d, mr).$$

If we start with a curve whose Hilbert function is this, we can derive in the same way a bunch of inequalities on the genus as before, the bound on $\pi(d,r)$ is strikingly different. We find in this case that the genus $g(C) \leq \pi_1(d,r) \sim \frac{d^2}{2r}$. So, in the range between $\frac{d^2}{2r-2}$ (the Castelnuovo bound) and this $\frac{d^2}{2r}$, we'll be able to conclude that they have to lie on scrolls.

Conjecturally, curves of large genus should lie on surfaces of small degree.

Lecture 15
10/21

§1 Introduction

We want to talk today about *extensions* of Castelnuovo theory. Today, we'll talk about something which is still being researched. We won't use it in the rest of the course, so if you're not looking for an interesting problem, just treat this as an occasion to get acquainted with an interesting problem, or just ignore it altogether. You have my permission.

You have to have a nineteenth-century mindset for this. If someone asked you today to classify curves, most people would answer the question by saying that "smooth projective curves have one discrete invariant (the genus), and within each genus a single irreducible family of dimension $3g-3$." If you wanted to say more, you would describe the space M_g parametrizing curves. In the nineteenth century, abstract curves weren't defined, so if you asked someone like Castelnuovo how to classify all curves, they would interpret the word "curve" as a subset of $\mathbb{P}^n$ for some n. In this case there are *two* discrete numerical invariants: the genus and the degree. So the first question you would ask is the analog of the statement that all genera occur—which degrees and genera occur? Castelnuovo's theorem, which we've gone through and proved, is the first step towards answering that: it gives a sharp upper bound.

But it doesn't answer the question. We now know what the maximum genus is, but we don't know about genera between zero and the maximal bound. The proposal that I'm going to make is to investigate this by going back to Castelnuovo's argument and seeing how it might be modified to address curves of large genus.

In the course of proving Castelnuovo's theorem, we derived a lower bound on the Hilbert function configuration of points in a hyperplane section. To do this, we used the fact that these hyperplane sections had points in linear general position. But this fact was a consequence of something stronger. We said that if $C \subset \mathbb{P}^r$ was a curve, then when you look the set of hyperplanes transverse to C, the *monodromy* of these points as you vary the hyperplane gives the full symmetric group. But you can say more.

Let H be a *general* hyperplane. Consider the monodromy on $H \cap C = \Gamma$ as H varies. This is the symmetric group on d letters (for d the degree). This implies that for all $\Gamma', \Gamma'' \subset \Gamma$ of the same cardinality and for all m, the Hilbert function $h_{\Gamma'}(m) = h_{\Gamma}(m)$. In other words, *you cannot distinguish between subsets of* Γ *of the same cardinality.* Remember how we set this up—we looked at

$$X = \{(H, p) : p \in H \cap C\} \subset U \times C$$

where U is the open subset in $\mathbb{P}^{r*}$ consisting of transverse hyperplanes, and argued that the fiber product $X_U^{(k)}$ was irreducible. That implies that the monodromy is k times transitive. Inside $X_U^{(k)}$ there is an open subset where the number of conditions is maximal. For a general point in U, *all* the k tuples of points in the corresponding hyperplane impose the same numbers of conditions. This again follows from irreducibility.

15.1 Proposition (Uniform position)**.** *For any general hyperplane H, and two $\Gamma', \Gamma'' \subset \Gamma = C \cap H$ of the same cardinality k and for all m, we have $h_{\Gamma'}(m) = h_{\Gamma''}(m), \forall m$.*

So you can't distinguish between subsets of the same cardinality using intrinsic invariants. We are going to use this strengthening of the linear general position lemma today.

§2 Castelnuovo's argument in the second case

Castelnuovo started this off by arguing that if a configuration of d points $\Gamma \subset \mathbb{P}^n$ is in linear general position, then

$$h_\Gamma(m) \geq h_0(m) \equiv \min(d, mn + 1).$$

This is *achieved* by taking points on a rational normal curve, where in fact: if $\Gamma \subset B$ for B a rational normal curve, then $h_B(m) = mn + 1$ for all m. When d is large, then any hypersurface of degree m which contains Γ will contain B and the number of conditions imposed by Γ and by B will be the same.

So we have this "minimal Hilbert function" which is achieved by points on a rational normal curve. Last time, we observed that $h_0(m)$ was the *smallest* possible Hilbert function of such a configuration of d points, under linear general position. Now let's assume *uniform position*, and ask what the *second* smallest possibility for the Hilbert function.

The point was $h_0(m)$ was the minimum of d and $mn+1$, where $mn+1$ was the Hilbert function of a curve that was as small as possible: the rational normal curve. So we were looking at points that lie on curves of small Hilbert function. What's the second smallest candidate? The claim is that we can do second-best by taking the set Γ *to lie on a linearly normal elliptic curve.*

Recall:

15.2 Definition. A curve $C \subset \mathbb{P}^r$ is **linearly normal** if $H^0(\mathcal{O}_{\mathbb{P}^r}(1)) \to H^0(\mathcal{O}_C(1))$ is surjective. This is equivalent to saying that the imbedding $C \hookrightarrow \mathbb{P}^r$ does not factor as an imbedding $C \hookrightarrow \mathbb{P}^{r+1}$ followed by a projection, or that the linear system cut out by hyperplanes is complete.

When the degree of the curve is small (less than $2r$), you checked on the homework that linear normality implies *projective* normality.

Let's make the following claim:

> The second smallest (irreducible, non-degenerate) curve in $\mathbb{P}^n$ is an elliptic (linearly) normal curve of degree $n+1$. In other words, one takes an elliptic curve B, an ample line bundle $\mathcal{L}$ of degree $n+1$, and imbeds E in $\mathbb{P}^n$ via the complete linear series $\mathcal{L}$.

In this case, the Hilbert function of B is

$$h^0(\mathcal{O}_B(\mathcal{L}^m)) = m(n+1)$$

by Riemann-Roch. Notice that the Hilbert function of a rational normal curve is just $mn+1$; the difference here is a shift in the degree, while the genus has jumped.

Another candidate would be to consider a *rational* (non-normal) curve of degree $n+1$ in $\mathbb{P}^n$, but the Hilbert function would be a little more because the genus is zero.

Let's state the result more formally:

15.3 Proposition. *Among all irreducible non-degenerate curves $B \subset \mathbb{P}^n$, the linearly normal elliptic curve of degree $n+1$ has the second-smallest Hilbert function.*

The proof is left as an exercise.

Now suppose $\Gamma \subset B$ is a finite set of cardinality d, in uniform position. Then

$$h_\Gamma(m) = \begin{cases} h_B(n) = m(n+1) & \text{if } d > m(n+1), \text{ by Bezout} \\ d & \text{if } d < m(n+1) \\ d-1 & \text{if } d = m(n+1) \end{cases}$$

In the latter case, the points impose independent conditions. Note that there is a slight gap in here. It's possible that the points Γ could be exactly the intersection of a hypersurface of degree m with the elliptic normal curve. We won't write out all the details. This is exactly what is described in part III of the Montreal books.

Remark. It is *nearly* true that points with the Hilbert function as above lie on an elliptic curve.

Let's call this function $h_1(m)$. We said that $h_0(m)$ was the smallest possible Hilbert function; $h_1(m)$ is the second-smallest. Now recall how Castelnuovo proceeded. He got a bunch of inequalities

$$h^0(\mathcal{O}_C(k)) - h^0(\mathcal{O}_C(k-1)) \geq h_0(k)$$

which he added up for varying k. This he used to bound the genus of C using Riemann-Roch.

So we now have a *second-smallest* Hilbert function. What we now want to conclude is, referring back to the picture of $C \subset \mathbb{P}^r$ and Γ the points of a general hyperplane section, then if Γ *doesn't* lie on a rational normal curve, then

$$h_\Gamma \geq h_1.$$

(**I'm not sure that we proved this.**) We can, in this same way, inequalities

$$h^0(\mathcal{O}_C(k)) - h^0(\mathcal{O}_C(k-1)) \geq h_1(k),$$

which gives *better* lower bounds on the Hilbert function of C, and thus better *upper* bounds the genus. The significant improvement between h_0 and h_1 gives a much better upper bound for the genus. We won't write this all out, so we'll just state:

15.4 Theorem. *For a curve C whose general hyperplane section doesn't lie on a rational normal curve,*

$$g(C) \leq \pi_1(d,r) \sim \frac{d^2}{2r}$$

Contrapositively, if the genus is *bigger* than $\pi_1(d,r)$, then the hyperplane section lies on rational normal curves, and we can in fact conclude:

15.5 Proposition. *If $g(C) > \pi_1(d,r)$, then C lies on a rational normal scroll.*

The point is that there aren't that many genera on a rational normal scroll, because the Picard group is so simple. For instance, there was a homework problem on this sort of thing. If S is a cubic surface in $\mathbb{P}^3$, then any C of genus $g(C) > \pi_1(d,r) \sim \frac{d^2}{6}$ *must* lie on a quadric with balanced type.

This was first observed by Halphen: he analyzed curves lying on cubic surfaces, then quartic surfaces, etc. He gave a complete answer to the answer of what (d,g) may be in $\mathbb{P}^3$, and got a huge prize for it—though it was later discovered his argument had a serious flaw. The standards of proof have shifted since his work.

§3 Converses

We'd now like to talk briefly about what's next. There is an analog of Castelnuovo's lemma in this setting, and it states:

15.6 Proposition (Fano)**.** *If $\Gamma \subset \mathbb{P}^n$ is a configuration of points in uniform position, consisting of $d \geq 2n+5$ points, if Γ imposes the second smallest number of conditions on quadrics—that is, if $h_\Gamma(2) = 2n+2$—then Γ lies on an elliptic normal curve.*

This is proved in the Montreal notes. Where is this all leading? The general conjecture is going to be that if we have a configuration of points that imposes few independent conditions on quadrics, it lies on a "small" curve. The general conjecture is:

Conjecture. For $\alpha = 0, \ldots, n-2$, and $\Gamma \subset \mathbb{P}^n$ a collection of points in uniform position of degree $d \geq 2n+3+2\alpha$, then $h_\Gamma(2) \geq 2n+1+\alpha$ should imply that Γ is contained in a curve $B \subset \mathbb{P}^n$ of degree $\leq n+\alpha$.

The point is that we're saying that curves impose relatively few independent conditions on quadrics. Imagine now that $d > 2^n$, the degree is huge; you would know then the points lie on the intersection of quadrics containing them, which will be an irreducible curve by a monodromy argument and a counting argument.

If we know this conjecture, we can put this together with Castelnuovo's argument. Let $h_\alpha(m)$ be the minimum Hilbert function of a configuration Γ contained in a curve of degree $n+\alpha$ in $\mathbb{P}^r$, we can derive a corresponding *bound* on the genus of C such that a general hyperplane section of C does not lie on a curve of degree $\leq r-1+\alpha$.

The picture is that we have a Castelnuovo bound for the genus, $g \leq \pi(d,r)$, and then a stronger bound $\pi_1(d,r)$ for curves whose hyperplane sections don't lie on curves of small degree, and then a stronger bound $\pi_2(d,r)$ for curves whose hyperplane sections don't lie on curves of slightly less small degree, and so on. It is conjectured that relatively small genera (well less than the Castelnuovo bound) all occur.

The key to using the Castelnuovo statement, or the variant we discussed today, is knowing that a surface of small degree is of a given type—say a scroll or a Veronese surface—and then to analyze curves on that surface. What do we know about surfaces of degree $\leq r-2+\alpha$ for varying α?

15.7 Proposition. *If $S \subset \mathbb{P}^r$ is a surface of degree $r-1+\alpha < 2r-2$,*[11] *then:*

1. *If $\alpha = 0$, S is a scroll or a Veronese surface.*
2. *Let B be a general hyperplane section, so B is a smooth curve with self-intersection $r-1+\alpha$. Then $K.B = 2g(B)-2-(r-1+\alpha) \leq 2\alpha-2-(r-1+\alpha) \leq \alpha-(r-1) < 0$. Thus, the canonical bundle of the surface is not effective. The surface is thus birationally ruled.*

Proof. Only the last claim requires proof: it is a consequence of general classification theory. This tells us that the surface is birational to a $\mathbb{P}^1$-bundle over a curve. ▲

So the point is that these surfaces of small degree are *easy.* There are no K3 surfaces in this case; we can completely classify these surfaces of small degree, and thus classify curves on such surfaces.

Anyway, if we know the truth of the above conjecture, we can answer the general Halphen question in $\mathbb{P}^r$: what genera occur for a given degree in $\mathbb{P}^r$?

15.8 Example. Let's take $\alpha = 1$. So surfaces $S \subset \mathbb{P}^r$ of degree one greater than the minimum, so degree r. What do those look like? Its hyperplane sections are

[11] So the genus of a general hyperplane section is at most $d-(r-1) = \alpha$ by Clifford.

irreducible nondegenerate curves of degree r in $\mathbb{P}^{r-1}$, thus yielding either rational non-normal curves or elliptic normal curves. In the first case,[12] S turns out to be the projection of a scroll in $\mathbb{P}^{r+1}$. In the second case, there are two possibilities: either it is a cone over an elliptic normal curve, or $r \leq 9$ and S is a del Pezzo surface.

Del Pezzo surfaces are interesting. To us in practice here, these are simply either $\mathbb{P}^2$ blown up at a bunch of δ points $p_1, \ldots, p_\delta$, imbedded in $\mathbb{P}^{9-\delta}$ by the linear system of cubics passing through $p_1, \ldots, p_\delta$. An example is the blow-up of $\mathbb{P}^2$ at six points, which sits as a cubic surface in $\mathbb{P}^3$.

So in high degrees, either your surface S is either not linearly normal, or it's a scroll or it is a cone over an elliptic normal curve. A curve on S which is smooth can be seen to have degree $\equiv 0, 1 \mod r$ by counting how many times it passes through a line of the ruling.

Lecture 16
10/26

There is no class on Friday because of the AGNES conference.

§1 Goals

We want to start a new topic today. In some sense, the basic problem we've been dealing with, in one form or another, is to say something about the dimensions of linear series on curves. That isn't the sole object of algebraic curve theory, but it's a good focal point.

The starting point for this is always Riemann-Roch, which tells you that if the degree is at least $2g-2$ (g the genus), the dimension of the complete series is exactly the degree d minus the genus, $d-g$. The answer to this question is pretty straightforward in small degree. The second thing we saw was when $0 \leq d \leq 2g-2$, we have Clifford's inequality $r(D) \leq \frac{d}{2}$ (equality achieved with the canonical series and also on hyperelliptic curves). So Riemann-Roch said that the graph of $r(D)$ with d (well, it's a multivalued graph) lies in a certain parallelogram; Clifford's inequality cuts it in half. That said, Clifford's theorem doesn't answer the question of *embeddings* in projective space, because equality in that inequality comes from the g^1_2 on a hyperelliptic curve, and thus some of the divisors meeting the bound aren't very ample.

So, we could ask another question: what can we say about the dimension of the linear series when we require that the map associated to the linear system to be a *(birational) embedding*? In that case, we saw a much stronger inequality: we had $g \leq \pi(d, r)$, which corresponds to the dimension of the linear series being bounded by a *parabola* which is substantially below the Clifford line. We saw, moreover, that linear series that achieve equality can be described. This was the Castelnuovo theory.

Now, the final question we want to raise, which we won't answer today, is the following. How about for C a *general* curve? That's not obvious; we have to talk about what a general curve is. We have seen, for instance, that some curves of a given genus have more linear series than others (e.g. among genus three curves, there

[12] Exercise.

are hyperelliptic and non-hyperelliptic curves). We want to know what linear series exist on *general* curves. What we'll see is that there is an even better estimate on the dimensions of these linear systems: below the Clifford line, and even below the Castelnuovo curve. This is the Brill-Noether theory.

But today, we'll do something more geometric.

§2 Inflection points

Consider a smooth curve $C \subset \mathbb{P}^r$. For instance, $C \subset \mathbb{P}^2$; let's take this as the initial case. Let's assume it's nondegenerate (i.e. not a line). The claim is that at a general point $p \in C$, the tangent line ℓ at p meets the curve with multiplicity exactly two at p. I.e., $m_p(C.T_pC) = 2$. This is true in *characteristic zero.*

Remark. This fails in characteristic p.

16.1 Definition. We say that $p \in C$ is a **flex point** if the tangent line at p meets C with multiplicity ≥ 3.

In the real picture, when one has a non-flex point, the plane curve stays locally to one side of its tangent line. If the plane curve goes to both sides of the tangent line, then there is a flex. One can see from the geometry of this that it's *obvious* that not every point can be a flex point, except that it fails in characteristic p.

Now, let's ask the same question with $C \subset \mathbb{P}^r$ where r could be greater than 2. Consider a curve $C \subset \mathbb{P}^r$, smooth and nondegenerate. When one looks at the curve, what does one see? We can talk about how the curve meets planes, with multiplicities. We can also talk about the curve meets various lines, which might not be a hyperplane. Let's introduce an additional notion.

16.2 Definition. For $\Lambda \subset \mathbb{P}^r$ a k-plane, we define the **order** $\mathrm{ord}_p(\Lambda.C)$ via

$$\mathrm{ord}_p(\Lambda.C) = \min_{H \supset \Lambda} m_p(C.H)$$

Here $m_p(H.C)$ is the multiplicity of the intersection; H is a *hyperplane.* So we can talk about **flex points** of a non-plane curve.

We want to say that most curves are such that no hyperplane intersects the curve in multiplicity ≥ 3. I.e., most curves in $\mathbb{P}^r, r \geq 3$, have *no* flex points (though this fails in $\mathbb{P}^2$). On next week's problem set, you'll see that if you take a general surface, and a general curve on that surface, it will have *no* flex lines.

16.3 Example. Consider a non-flex point $p \in C$, and consider the tangent line ℓ at C. Then ℓ has contact of order two with the curve at p. If we look at the hyperplanes containing ℓ, then these will have contact at least two with C at p. The claim is that there is a *unique* plane containing ℓ which has contact ≥ 3 with C.

If you recall studying real arcs in $\mathbb{R}^3$, then you know at every point there is the notion of *curvature* and *torsion.* The curvature has to do with how rapidly the tangent line moves. At every point, you also have a canonical normal vector, and the plane spanned by the tangent line and the normal vector has contact of order three with the arc. If you look at the planes containing the tangent line, there will be a *unique* plane such that the curve lies on both sides of the plane in a neighborhood of that point.

This is called the **osculating plane.** Another way of saying this that if you take the scheme $3p$ (the thickened point), then the claim is that since it is not contained in a line, it spans a plane.

You can pretty much visualize what's going on. The word "flex" in higher dimensional spaces still means the same thing; the tangent line has contact ≥ 3 at that point, as we said. Classically, as people looked at space curves, they described the special behavior at flex points.

16.4 Definition. Let C be a space curve. At some points $p \in C$, we may have that the multiplicity of intersection of C with the osculating hyperplane H is > 3. These are called **stalls.** In differential-geometric terms, this states that the curve is not moving rapidly away from its osculating plane; the *torsion* of the curve is zero there.

This is a completely archaic, nineteenth-century word, and I suggest you forget it immediately.

§3 A modern reformulation

All this is just to suggest that there's a question here we need to address. We will introduce some of the definitions that we *do* use, not the nineteenth-century version.

We're ultimately going to be applying this to curves in projective space, but let's apply this to linear systems.

Let C be a smooth curve. Consider a line bundle $\mathcal{L}$ and a space $V \subset H^0(\mathcal{L})$, of degree d. Let $\dim V = r$, so we have a g^r_d. We will be using this notation throughout today's lecture.

16.5 Proposition. *Let $p \in C$. Consider the set $\{\mathrm{ord}_p(\sigma)\}_{\sigma \in V - \{0\}}$; this is the set of orders of vanishing of sections in V at p. The cardinality of this set is exactly $r+1$.*

So there are exactly $r+1$ different orders of vanishing.

Proof. One direction is clear: a bunch of sections with different orders of vanishing are necessarily linearly independent. So we get $\leq$. The other direction is also simple. Choose a basis $\sigma_0, \ldots, \sigma_r \in V$, and if $\mathrm{ord}_p(\sigma_i) = \mathrm{ord}_p(\sigma_j)$, then replace σ_j with an appropriate linear combination of σ_i, σ_j vanishing to higher order. Keep doing this, and obtain a basis for V consisting of elements which have distinct orders of vanishing at p. ▲

We now give some notation and terminology.

16.6 Definition. We write this collection of orders of vanishing of sections in V at p, i.e. $\{\mathrm{ord}_p\sigma\}_{\sigma \in V \setminus \{0\}}$, will be written as $\{a_0, \ldots, a_r\}$ with $a_i < a_{i+1}$ for each i. This is called the **vanishing sequence** of the linear series V at that point, and sometimes we will write $a_i(V, p)$ to indicate the dependence.

We write $\alpha_i = a_i - i$, which is a *nondecreasing* sequence $0 \leq \alpha_0 \leq \alpha_1 \cdots \leq \alpha_r$; this is called the **ramification sequence.** The **total ramification index** $\alpha = \alpha(V, p)$ as

$\alpha = \sum \alpha_i$. We say that p is an **inflectionary point** of the linear series V if $\alpha > 0$. This is equivalent to saying that $\alpha_r > 0$, or that $a_r > r$. What that means, at the end of the day, is that there is $\sigma \in V$ which vanishes at p to order strictly bigger than r.

The claim is that this notion of an inflectionary point includes the classical ideas of inflection points, stalls, etc. Here we look at the linear series cut out by hyperplanes. If C is a plane curve, V the linear series of lines, then the highest multiplicity of vanishing of a section is the intersection multiplicity at that point of the tangent line. So an inflection point in this sense corresponds to an inflectionary point in the modern sense.

Let us make some observations. Let $p \in C$.

1. To say that $a_0 = \alpha_0 > 0$ is to say that p is a *base point* of the linear series.
2. If $a_0 = \alpha_0$, to say that $\alpha_1 > 0$ (or that $a_1 \geq 2$) is to say that there is a section of the line bundle that doesn't vanish at that point, but every section which vanishes at that point vanishes to degree two. In other words, the *differential* of the map associated to the linear series at that point is zero; the associated map to projective space (which is defined at p) fails to be an immersion at that point.

Now we want to state and prove the basic lemma.

16.7 Theorem (Basic lemma, characteristic zero)**.** *For any linear series $(\mathcal{L}, V)$ on any smooth curve C, for a general point $p \in C$, $\alpha(V,p) = 0$: that is, a general point is not an inflectionary point.*

I.e., at a general point, the sequence $a_0, \dots, a_r$ is just $0, 1, \dots, r$.

Proof. Pick $p \in C$. Say the map $\phi_V : C \to \mathbb{P}^r$ is given locally by a vector-valued function $v(z)$; we've chosen local coordinates around p. We can do this by locally trivializing the line bundle. So $v(z) = (\sigma_0(z), \dots, \sigma_r(z))$, where we have written each σ_i as a function by trivializing the line bundle. Now the observation is that if p is an inflectionary point, there is a linear combination of these σ_i that vanishes to order $r+1$ or more at p. That is, this vector-valued function $v(z)$ is such that $v(p) \wedge v'(p) \wedge \cdots \wedge v^{(r)}(p) = 0$; that's what it means for p to be an inflectionary point. Suppose every point were an inflectionary point, i.e. this function $v \wedge v' \wedge \cdots \wedge v^{(r)} \equiv 0$.

OK, so we ask: let's say k is the smallest integer such that $v \wedge v' \wedge \cdots \wedge v^{(k)} \equiv 0$. We can assume that the wedge product of v with its first $k-1$ derivatives is nonzero. Now, if we look at this expression, we know that $\{v, \dots, v^{(k-1)}\}$ are linearly independent, so the hypothesis is that $v^{(k)}$ is in the span $\{v, v', \dots, v^{(k-1)}\}$. Let's take derivatives of the identity

$$v \wedge v' \cdots \wedge v^{(k)} \equiv 0;$$

we find by differentiation (since almost everything cancels when you apply the product rule)

$$v \wedge v' \cdots \wedge v^{(k-1)} \wedge v^{(k+1)} \equiv 0.$$

So $v^{(k+1)}$ is in the span of $\{v, v', \dots, v^{(k-1)}\}$. Now take the derivative of this thing (which is zero) *again.* A priori, we could get two nonzero terms. We'll get

$$v \wedge v' \cdots \wedge v^{(k-1)} \wedge v^{(k+2)} \pm v \wedge v' \cdots \wedge v^{(k-1)} \wedge v^{(k)} \wedge v^{(k+1)} \equiv 0$$

since we've differentiated zero. But the second term is zero because all the vectors are in the span of the k-dimensional space $\{v, v', \dots, v^{(k-1)}\}$. So the first term is also zero. We find $v^{(k+2)}$ is likewise in the span of $\{v, v', \dots, v^{(k-1)}\}$. Now you see where this is headed—each time you keep taking higher derivatives of this identity, you find that all the higher derivatives lie in the span of $\{v, v', \dots, v^{(k-1)}\}$.

So we're saying that *all derivatives* of v lie in the span of the first $k-1$. That is a linear series properly contained in $\mathbb{C}^{r+1}$. All the derivatives lie in this subspace, which means that the image of this curve under v lies in a proper subspace of projective space.[13] This contradicts the fact that a map associated to a linear system is nondegenerate. ▲

16.8 Example. Remember: there are plane curves in characteristic p all of whose tangent lines pass through a point. Consider for instance $XY - Z^2$ in characteristic two; this is a smooth curve in $\mathbb{P}^2$. The tangent lines *all pass through* $[0, 0, 1]$. So the tangent vectors all lie in a proper linear subspace, even though the curve is nondegenerate.

One more observation and then we'll state and prove the Plücker formulas. We have seen that one way to view an inflectionary point is to think of the curve and the linear series as a vector-valued function, and then to consider the place where the vector-valued function and its derivatives are linearly independent.

Define v as in the previous proof. Consider the function $v \wedge v' \wedge \cdots \wedge v^{(r)}$. This is a function which is not identically zero, as we're in characteristic zero. This is the wedge product of $r+1$ vectors in $r+1$ space, so it's a scalar. This is just the *Wronskian* determinant of the matrix

$$\det \begin{bmatrix} \sigma_0 & \dots & \sigma_r \\ \sigma_0' & \dots & \sigma_r' \\ \sigma_0'' & \dots & \sigma_r'' \\ \vdots & \vdots & \vdots \\ \sigma_r^{(r)} & \dots & \sigma_r^{(r)} \end{bmatrix}.$$

Here we are thinking of the sections $\sigma_i \in V$ as functions by choosing a local trivialization.

16.9 Proposition. *The* order of vanishing *at a point of this Wronskian is* exactly equal *to the total ramification index $\alpha(V, p)$.*

Proof. Choose σ_0 to vanish to minimal order at p, σ_1 to next-to-minimal order, and so on. Then compute what the determinant looks like near p. For instance, σ_0 looks like z^{a_0} in a local coordinate system, σ_1 looks like z^{a_1}, and so on. Thus we can figure out how each row looks like. The first nonzero derivative of the matrix can be seen to be the derivative of order $\alpha_0 + \cdots + \alpha_r$. ▲

From this, we want to draw a conclusion. To write this Wronskian out, we had to first choose a trivialization, and then choose an appropriate basis of V. To what extent were our choices dependent on these choices? Let's say we have an open cover $\{U_\alpha\}$ of C and transition functions $f_{\alpha\beta}$ that define the line bundle $\mathcal{L}$. Then the *ratio* of the

[13] This is the only part of the argument that fails in characteristic p.

two different expressions for the σ_i given by local trivializations will be given by $f_{\alpha\beta}$. One can check easily that the lowest order term in the Wronskian only changes by a nonzero constant. (Well, one should note: the operation of "taking a derivative with respect to the local coordinate" changes when one changes local coordinates.)

In fact, this determinant Wronskian is a section of $\mathcal{L}^{r+1} \otimes K^{\binom{r+1}{2}}$ for K the canonical bundle. This follows by the same idea of computing in local coordinates. There are more modern ways of saying this. It follows that the *total number of zeros* of the Wronskian is just the degree of this line bundle, which is

$$\deg(\mathcal{L}^{r+1} \otimes K^{\binom{r+1}{2}}) = d(r+1) + \binom{r+1}{2}(2g-2) = (r+1)(d + r(g-1)).$$

(Recall that d was the degree of the line bundle $\mathcal{L}$.) The conclusion:

16.10 Proposition (Plücker formula)**.** *The sum of the total ramification indices, $\sum_p \alpha(V,p)$, is $(r+1)(d+r(g-1))$.*

Let's look at some examples of this. We should also check a bunch of things at once.

16.11 Example. What goes on when $g = 0$ and the linear series is complete? There is a unique complete linear series in each degree. Then the map to projective space is the Veronese map, which imbeds $\mathbb{P}^1$ in $\mathbb{P}^d$ as a rational normal curve. This is *homogeneous* in projective space; the automorphism group of projective space fixing the rational normal curve permutes the points on it. So there can't be any ramification. More simply, a hyperplane can't meet the curve with degree $> d$ because the curve has degree d. Note that the formula gives zero.

We can also get something if the linear series isn't complete. We could consider the rational quartic curve, where $g = 0, d = 4, r = 3$. This is a projection of a rational normal quartic. How many flex points does this curve have? We can just plug it into the formula. The curve has 4 flex points. The flex points are as follows: each point in the rational normal curve has an osculating plane at each point. The osculating plane of the rational normal curve has contact of order three at each point. To get a point in $\mathbb{P}^3$ where the osculating plane in the *projection* where the intersection was order > 3, then the claim is that the point of projection lies on the osculating plane. So, a general point in $\mathbb{P}^4$ lies on four osculating planes of this curve, a fact you can verify directly.

16.12 Example. Take $g = 1$, V complete, $d = r+1$. This will give an imbedding of $C \hookrightarrow \mathbb{P}^r$ of degree $r+1$. Before we even try to remember the formula, let's ask what the inflectionary points are. The inflectionary points on this curve are points where there is a hyperplane vanishing to order at least $r+1$. But the curve has degree $r+1$. This means that there is a hyperplane that meets the curve *only* at one point.

If we have two such hyperplanes, the ratio of the linear forms is then a meromorphic function on the torus with a pole of order $r+1$ at one point and a zero of order $r+1$ at another point. So the inflectionary points are a union of cosets of torsion points of order $r+1$, if we fix an origin and get the elliptic curve structure. The number of inflectionary points is given by the above Plücker formula as $(r+1)^2$, and that is exactly the number of $r+1$-torsion points on an elliptic curves.

16.13 Example. If you use the Plücker formula with $r = 1$, then you recover Riemann-Hurwitz.

Lecture 17
11/2

§1 Review

Let C be a smooth projective curve of genus g. Let $\mathcal{L}$ be a line bundle of degree d on C. Let V be a $r+1$-dimensional space contained in $H^0(\mathcal{L})$, and let $p \in C$. Let $a_0 \leq a_1 \leq \cdots \leq a_r$, and consider the set $\{\mathrm{ord}_p(\sigma), \sigma \in V\} = \{a_0, \ldots, a_r\}$ (i.e. arrange the a_i in order). Let $\alpha_i = \alpha_i(V,p) = a_i - i$, and let $\alpha = \alpha(V,p) = \sum \alpha_i$ be the *total ramification of* V *at* p. This is the same notation as we had last time.

The basic idea is the picture in the plane that you're probably familiar with. In the plane, we say that a point on a curve is an *inflectionary point* if the tangent line has contact to order ≥ 3 to the curve.

If we have a curve $C \subset \mathbb{P}^3$, at a general point $p \in C$, the tangent line ℓ to the curve and look at the planes in $\mathbb{P}^3$ containing ℓ, all but one of them will lie on one side of the curve—except one. There will be one particular plane which bisects the curve. Alternatively, if you projected the curve from a point on the tangent line, then in the neighborhood of the image point, you'd see a cusp. There will be a unique line in the plane that cuts through the cusp, the tangent line through the cusp; the unique preimage of that in $\mathbb{P}^3$ is this plane, and is called the *osculating plane.* In general, the osculating plane will have contact of order three at that point. If the plane has contact of order four or more, then the point on the curve is *special* and is called a *stall.*

We saw last time that the above notion of total ramification for sections of a line bundle is a generalization of the above classical notions, when $\mathcal{L} = \mathcal{O}(1)$. But the notion for line bundles applies more generally, even for line bundles which aren't ample.

The formula which we discussed last time is called the Plücker formula. In modern language, it computes the total ramification:

$$\sum_{p \in C} \alpha(p, V) = (r+1)d + \binom{r+1}{2}(2g-2) = (r+1)(d + r(g-1)).$$

This basically follows because we write the $\alpha(p, V)$ as the zeros of a section of a line bundle, whose degree is the thing on the right.

Last time, we played around with this formula, and showed that a rational normal curve has no flexes, for instance. (This is also true because a rational normal curve is homogeneous.) When you project a rational normal curve, to get a rational non-normal curve, then the curve acquires inflectionary points. We saw that there were four inflectionary points for the projection of a rational normal curve in $\mathbb{P}^4$ to $\mathbb{P}^3$. The one last example was an elliptic normal curve, where we saw that the inflectionary points are exactly the torsion points of order two.

§2 The Gauss map

Let's now give another intepretation of this setup. This will introduce a couple of new ideas. Given a curve $C \subset \mathbb{P}^r$, or more generally a map $C \to C_0 \subset \mathbb{P}^r$ (what we're about to say applies to this more general situation, but thinking in terms of an imbedded curve helps with intuition), then consider: for every $p \in C$, we have a tangent line T_pC. There is naturally a *Gauss map* $\phi_1 : C \to \mathbb{G}(1, r)$ which sends a point to its tangent line. If the inclusion (or map) of C in (or to) $\mathbb{P}^r$ is given by a holomorphic vector-valued function $z \mapsto [v(z)]$, then ϕ_1 sends z to $[v(z) \wedge v'(z)]$ in the Grassmannian. This might not make sense at every point, but at least in characteristic zero, this map is generically nonzero. So at least, OK, we get a rational map from $C \dashrightarrow \mathbb{G}(1, r)$, and then you can extend to $\mathbb{G}(1, r)$ because the latter is proper.

More generally, if $k < r$, let's define the map ϕ_k which sends $z \in C$ to the k-plane spanned by $v(z), v'(z), \dots, v^{(k)}(z)$, so the map to the Grassmannian thus defined sends

$$\phi_k : z \mapsto [v(z) \wedge v'(z) \wedge \cdots \wedge v^{(k)}(z)].$$

This is true generically, and then we extend. Just as the Gauss map sends a point to its tangent line, we define:

17.1 Definition. The **osculating k-plane to C at a point** p is $\phi_k(p)$.

Just as in the case of the tangent line, the osculating k-plane at p is exactly the k-plane with maximal order of contact (at least $k + 1$).

If you want to see a picture of what goes on here, let's wonder when $v(z) \wedge v'(z) = 0$. What happens there is a cusp. If you want to see more, try some of the exercises in Batch C which aren't in the homework.

Here's one last thing to say.

17.2 Example. Take $k = r - 1$. In this case, the map ϕ_k associates to each point its *osculating hyperplane*, which means a map $\phi : C \to \mathbb{P}^{r*}$. We thus get a curve in the *dual projective space*, which we can call the **dual curve.**

Let's give an example. Take $C \subset \mathbb{P}^3$ a twisted cubic curve. What's the dual curve in this case? It's a twisted cubic. That's because the twisted cubic is homogeneous. So the dual curve is also homogeneous, and there's only one such example—the twisted cubic. This follows because if you have a homogeneous curve, then you can use the Plücker formula to compute the genus.

Let's state this as a remark:

Remark. The only irreducible nondegenerate curve in $\mathbb{P}^3$ which is homogeneous is the twisted cubic.

Or you could try to compute the dual curve explicitly, which is less fun.

Caution: When you have a variety $X \subset \mathbb{P}^r$ in general, then there's a notion of the *dual variety*, which is the locus of tangent hyperplanes (i.e. hyperplanes containing a tangent plane) to X. This is usually a hypersurface. This is *not* what we mean by the dual curve above; we just don't have a better name.

The next (closely related) topic after this, incidentally, is going to be the study of Weierstrass points, and we'll look at inflectionary behavior again.

§3 Plane curves

Consider $C \subset \mathbb{P}^2$. For notational ease, we're going to think of it as a curve *in a plane,* rather than a map from an abstract smooth curve to a plane. Let's assume C isn't a line.

Let C^* be the dual curve in $\mathbb{P}^2$, so it is the *image* of C under the Gauss map $C \to \mathbb{P}^{2*}$, which is a birational embedding (at least in characteristic zero). This is also the dual variety, incidentally. So, the genus of the normalization of C is the same as the genus of the normalization of C^*.

What would you like to know about the dual curve? You can ask lots of questions.

Question.

1. What's the degree of C^*?
2. What types of singularities does the dual curve C^* have?

A basic fact here is that the dual curve of the dual curve is the original curve itself. You can sort of see that, in fact. If you have a curve $C \subset \mathbb{P}^2$, and you look at the map that sends each point to its tangent line, and think of that as giving a map to $\mathbb{P}^{2*}$, then: if you pick two points $p, q \in C$ and take the limit of the intersection of the tangent lines ℓ_p, ℓ_q as $q \to p$ is just p itself. In the dual projective space, the secant line through the points $p^*, q^* \in \mathbb{P}^2$ corresponds to the intersection of ℓ_p, ℓ_q. If you think about it carefully, it follows.

Anyway, we just write this down.

17.3 Proposition. *$C^{**} = C$.*

Classically, people thought that there should be a theory of curves which applies *symmetrically* to a curve and its dual. What class of curves should we try to apply this to? Instead of arbitrary plane curves, we might look at a class of simpler curves—say, smooth curves. But that doesn't work. If I have a smooth curve, the dual curve is not necessarily smooth. If C is a smooth curve with a tangent line ℓ which is tangent at *both* p, q, then the Gauss map is going to send p, q to the *same point.* So the dual curve will have a node. And conversely: a node on the dual curve corresponds to a bitangent. Less visibly, if the curve you start with has a flex point, then there's a cusp on the dual. This can be computed in local coordinates.

In other words, we *can't just stick to smooth curves.* We wouldn't then get a theory that applies to the dual.

17.4 Example. There is only one curve in $\mathbb{P}^2$ which is both smooth and has a smooth dual: the plane conic.

This idea of having a symmetric (duality-preserved) class of curves motivates the next definition.

17.5 Definition. $C \subset \mathbb{P}^2$ has **traditional singularities** if C, C^* have only nodes and cusps.

This means that C can't have things like hyperflexes and tritangents. More precisely, the requirement is that C have only nodes or cusps as singularities, and only simple flexes and bitangents. Any line $\ell \subset \mathbb{P}^2$ can only meet C in one of the three types of intersections:

1. d distinct, simple points (for d the degree).
2. One double point and $d-2$ simple points.
3. One triple point and $d-3$ simple points.
4. Two double points and $d-4$ simple points.

This will ensure that C^* has only nodes and cusps.

Anyway, it's not clear—even now—why we choose this class of curves. It happens to be symmetric with respect to duality, but that doesn't explain it all. We're going, however, to apply certain reasoning to this class of curves. It'll simplify things.

Let C be a curve with traditional singularities. We're going to associate a whole bunch of numbers.

1. Let d be the degree of $C \subset \mathbb{P}^2$; let d^* be the degree of C^*.
2. Let g be the genus of C (or C^*); this means the genus of the normalization.
3. δ will be the number of nodes, and κ the number of cusps on C. b will be the number of bitangents, and f the number of flexes. We have

   $$\delta = b^*, \kappa = f^*, b = \delta^*, f = \kappa^*,$$

   where the starred variants of b, δ, κ, f are defined in the same way for C^*. That's because of the previous duality discussion.

Now we want to say what relations hold among these seven distinct invariants. These are called the *Plücker relations.* We're going to start deriving them, but let's first ask about the genus.

If the curve were smooth, then the genus would be $\binom{d-1}{2}$. But C has nodes and cusps; these, however, drop the genus by one. Let me just take a moment out and give a bit of justification for

17.6 Proposition. *The genus of C is $g = \binom{d-1}{2} - \delta - \kappa$.*

Proof. More generally, if $C \subset S$ is a curve on a surface with a double point at p, then we can *blow up* S at p, to get $\widetilde{S}$, and then the proper transform of C will be a smooth curve on $\widetilde{S}$ intersecting the exceptional divisor of the blow-up in two points. Note that we can compute canonical divisors, $K_{\widetilde{S}} = \pi^* K_S + E$ if E is the exceptional divisor and $\pi : \widetilde{S} \to S$. Likewise, $[\widetilde{C}] = \pi^*[C] - 2E$ for $\widetilde{C}$ the proper transform. Now apply adjunction. One finds that the arithmetic genus of $\widetilde{C}$ is one less than the arithmetic genus of C. So having the node drops the arithmetic genus by one. That's pretty much all we need for this calculation, although there's another case with the cusp. The formula easily follows. ▲

Of course, we also get

$$g = \binom{d^*-1}{2} - \delta^* - \kappa^* = \binom{d^*-1}{2} - b - f,$$

so if we knew everything else, we could get the degree d^* of the dual curve. But let's derive one more relation.

There are lots of ways of arriving at the degree of the dual. We could just apply the Plücker relations. But let's do it more concretely. Consider a *smooth* curve C. If C is the zero locus of a polynomial F, to ask for the degree of the dual C^* is to ask how many times a line in $\mathbb{P}^{2*}$ meets C^*; this is, in other words, the *number of tangent lines to* C that pass through a given point in $\mathbb{P}^2$. That is, we want to know how many lines starting from the point at ∞ are tangent lines. To get this, we need to find common zeros of $F, \frac{\partial F}{\partial x}$ (these are vertical tangent lines), and by Bezout, the number of zeros is $d(d-1)$. Thus, if C is smooth, the number of zeros is $d(d-1)$.

What happens when there are nodes? Here the Gauss map will send the two points above a node in C to the two tangent lines at that point. A general line will intersect the curve twice at that point. So we find:

$$d^* = d(d-1) - 2\delta - 3\kappa.$$

(This could also be done by applying Riemann-Hurwitz to a projection map.)

Let's now write down the relations:

1. $d^* = d(d-1) - 2\delta - 3\kappa$.
2. $d = d^*(d^*-1) - 2b - 3f$.
3. $g = \binom{d-1}{2} - \delta - \kappa$.
4. $g = \binom{d^*-1}{2} - b - f$.

So, any three of these quantities determine the rest.

17.7 Example. Consider C is a degree d, and let's say C is smooth, with only traditional singularities. Then $\delta = \kappa = 0$. We get the genus $g = \binom{d-1}{2}$, which we knew already. Also, we get $d^* = d(d-1)$.

Let's ask: how many flexes and how many bitangents are there?

Notice that we could answer the number of flexes by applying the earlier Plücker formula we had on the board a half hour ago. Let's do it another way.

We also get $d = d^*(d^*-1) - 2b - 3f$. So we can determine $2b+3f$. That is,

$$2b + 3f = d^*(d^*-1) - d = (d^2-d)(d^2-d-1) - d = d^4 - 2d^3.$$

That gives us one relation between flexes and bitangents. Using another formula, we get

$$2b + 2f = (d^*-1)(d^*-2) - (d-1)(d-2).$$

If we plug in for $d^* = d^2 - d$, we get

$$2b + 2f = (d^2-d-1)(d^2-d-2) - (d-1)(d-2) = d^4 - 2d^3 - 3d^2 + 16.$$

Now we can just solve. If we want to find f, we can just subtract the second from the first. *The number of flexes is thus* $3d(d-2)$, which is also the number of points of intersection of the plane curve with its Hessian. *The number of bitangents is* $\frac{1}{2}(d^4 - 2d^3 - 9d^2 + 18d)$. This factors nicely as $\frac{1}{2}d(d-2)(d-3)(d+3)$.

Let me just say one more thing here. It'll be a subject of a nice exercise. This is going to do it for us when it comes to flexes and inflectionary points. As always, we've been talking about a fixed curve, or a map of a fixed curve to projective space. We haven't really talked about how curves behave in families, but it's a natural question to ask. A general curve in the plane is smooth, and it will have the number of flexes and bitangents given above. But there will also be curves with a simple node, or so on. What if we have a family of curves in the plane and they specialize to a curve with a single node? What happens to the flexes and the bitangents?

Let $\{C_t\}$ be a family of curves which are mostly smooth, and suppose C_0 has one simple node. How are the above formulas in the example altered? What is true is that there will be *six fewer flexes* at C_0 and the number of bitangents drops by $2(d^2-d-6)$. What happens to the missing six flexes and the missing bitangent lines? Where do they go? I'll just tell you the answer and then we'll call it quits for today. Think about it this way. We're talking about curves with traditional singularities; one way to characterize them is to look at the possibilities of intersections with various lines. A flex point of a curve is one where the tangent line "the flex line" meets the curve with multiplicity at least three. If you have a family of flex lines to a family of mostly smooth curves, then in the limit, then what might happen is that the flex line would "degenerate" to one of the tangent lines at a *node.* Missing flexes can disappear this way.

Let's ask the following question: *What about the case of C_0 with a cusp?*

Lecture 18
11/4

Today, we're going to talk about Weierstrass points. Next week, we'll talk about real algebraic curves. What we've been doing for the last week, and what we're going to continue doing for the next week, is to talk about various slightly isolated topics in the syllabus. After that, we'll get back to the main goal of this course: that is, the study of linear series on curves (or maps of curves in projective space). One crude form of the question is, What linear series exist on a curve? The most fundamental version of the question is whether a general curve of a given genus admits a map of a given degree to projective space. This is the Brill-Noether theorem. If there is such a map of a given degree (or linear series) to projective space on a general such curve, we can ask more: what is the geometry of that space of linear series on the Jacobian? So, after next week, we'll go into Brill-Noether theory.

We, as mathematicians, tend to something a little underhanded: when we can't solve a problem, we come up with some simpler variant of it and promise that when they've solved it, they can solve the real problem. When people first started talking about algebraic curves hundreds of years ago, they studied curves in $\mathbb{R}^2$. They asked questions about the number of connected components in terms of the degree. They had two separate notions of a node—one, when both branches are real, and one, where both branches are complex (e.g. $x^2+y^2=0$; over the real numbers, it's just a point). They had different names for the nodes $x^2+y^2=0, x^2-y^2=0$, and there were all these questions that were specific to curves over the real numbers.

In the nineteenth century, people discovered that it was *easier* to work with curves over the complex numbers. A curve over the real numbers gives a curve over the complex numbers, but one with a distinguished involution. The way to study real plane curves thus seems to be to study complex plane curves. Nonetheless, many things are easier:

18.1 Example. Smooth plane cubic curves in $\mathbb{R}^2$ may have two connected components; smooth plane cubic curves in $\mathbb{C}^2$ are connected.

We're going to relate questions over the complex numbers to questions over the real numbers, by proving Harnack's theorem.

After people discovered that working over $\mathbb{C}$ makes things easier, they pulled a bait-and-switch and stopped working with curves over $\mathbb{R}$. Still, for general culture, we're going to talk a little about it next Wednesday.

§1 Inflectionary points; one last remark

If we have a curve $C \subset \mathbb{P}^r$ (or more generally $C \to \mathbb{P}^r$ a map), we have a distinguished finite subset: the inflectionary points. These are special points on the curve *as it sits inside projective space.* This is an extrinsic invariant: it depends not just on the abstract curve, but on the particular choice of embedding in projective space. There is one exception to that: if we take the map to projective space to be the *canonical map*: that's a canonical choice of embedding. Inflectionary points for the *canonical model* are intrinsic. If I just give you an abstract curve C of any genus $g \geq 2$, there's a naturally defined distinguished finite subset.

You would think that on a curve, there are no distinguished points; this is the intuition you get from real manifold theory. This is *false.*

Remark. If C is hyperelliptic, then we get a map $C \to \mathbb{P}^{g-1}$ given by the canonical map, though it isn't an embedding. The inflectionary points are the ramification points of this map (note that the image curve, as a subset of projective space, is a rational normal curve and has no inflectionary points).

Similarly, inflectionary points of the *pluricanonical* models are intrinsic: that is, inflectionary points of the line bundles $K^{\otimes m}$. The number of inflectionary points (with multiplicities) will grow with m, by the Plücker formulas.

Remark. There is a reason why the canonical bundle is called *canonical*, other than it's canonical. Here we'll invoke the existence of a space that we haven't constructed. Let M_g be the set of isomorphism classes of smooth projective curves of genus g. We can also give $P_{d,g}$, which is the moduli space of pairs $(C, \mathcal{L})$ where C is a smooth projective curve of genus g, and $\mathcal{L}$ is a line bundle of degree d. Though we haven't proved it, both these have natural structures of varieties. There is a forgetful map $P_{d,g} \to M_g$; one might ask if there are sections to this map, or even rational sections. Such a section would be the data of assigning to a general curve, a line bundle of some degree of that curve.

One way of finding such an assignment, or such a section, is to take powers of the canonical bundle. So there are sections if d is a multiple of $2g-2$. It is a theorem, though well beyond the scope of this course, that the *only rational sections* of $P_{d,g} \to M_g$ are given by powers of the canonical bundle. The canonical bundle is not just canonical; it's *uniquely* canonical.

Remark. If we ask the same question of the previous question, but instead let $H_{n,g,r}$ be the space of "smooth" curves $C \subset \mathbb{P}^r$, nondegenerate of genus g and degree n. This is an open subscheme of the Hilbert scheme. We can form the bundle of Picard varieties over it; that is, we can look at the set of pairs $(C, \mathcal{L})$ where $C \in H_{n,d,r}$ and $\mathcal{L}$ is a line bundle of degree d on C. Let's call this P_H. We get a map $P_H \to H_{n,g,r}$ which is forgetful; you might ask whether there is a section.

There are clearly sections to this map: for instance, the canonical line bundle, and also $\mathcal{O}(1)$. So there is a rank two group of sections $H_{n,g,r} \to P_H$. The conjecture is that these are the only sections, for $r = 1, 2$ (on a given component of $H_{n,g,r}$, at least). For $r \geq 3$, this is *false*. Anyway, these are a lot of interesting classical questions.

Anyway, whenever you have an extrinsic invariant of a curve (one defined based on a map to projective space), the natural thing to get an intrinsic invariant is to apply it to the canonical (or pluricanonical) map. The above remark shows that this is not only natural, but just about the only shot.

§2 Weierstrass points

This is going to be relatively brief. The basic idea, in some sense, is something that is very natural. When we first talking about line bundles and divisors, the reason we did this was that compact Riemann surfaces don't have any nonconstant holomorphic functions: we have to look at meromorphic functions instead. But if we look at all meromorphic functions, there are way too many. So one strategy is to consider meromorphic functions with limited poles, by prescribing a worst possible divisor. Among those, one can just look for meromorphic functions which are holomorphic everywhere but a point.

Let C be a projective curve of genus g. Let's fix $p \in C$, and consider meromorphic functions on C which are holomorphic on $C \setminus \{p\}$. We want to know how many there are. If you allow a pole of high order, Riemann-Roch tells you that there will be a lot of such functions. In other words, you would be considering $H^0(\mathcal{L}(mp))$; when $m > 2g-2$, we know exactly what the dimension of this space should be. For low orders, it's not clear.

Given a curve and a point, what can we say about the *ring* of meromorphic functions holomorphic outside p? In particular, given m, when is there a meromorphic function in this ring whose order at p is precisely m?

The basic observation is as follows.

Say D is an effective divisor of degree d on C. We want to compare the space of sections $H^0(\mathcal{O}(D))$, meromorphic functions with poles on D, with $H^0(\mathcal{O}(D+p))$; in other words, we want to compare $r(D)$ and $r(D+p)$.

Riemann-Roch says that

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

What's the relation? The claim is that either the two quantities are equal, or the difference in dimensions is exactly one. That is,

$$r(D+p) = \begin{cases} r(D) \\ r(D)+1 \end{cases} .$$

We can tell the difference by looking at whether $h^0(K-D)$ and $h^0(K-D-p)$ are equal or not. Thus, we find that $r(D+p) = r(D)$ if p is not a base point of $K-D$; if $r(D+p) = r(D)+1$ if p is a base point. This is because p's being a base point of $K-D$ is equivalent to $h^0(K-D) = h^0(K-D-p)$.

We're essentially asking the following: does $D+p$ impose more conditions on holomorphic differentials than does D? This is what determines whether or not $r(D) = r(D+p)$, or whether the dimension of linear series is "bumped" or not.

Alright, now we want to apply this observation to the sequence of divisors

$$0, p, 2p, 3p, \dots, mp, \dots.$$

At each stage, we can look at $h^0(K) = g$, and then we look at

$$h^0(K-p), h^0(K-2p), \dots, h^0(K-mp), \dots;$$

at each stage, either the two successive numbers in the sequence are equal (in which case the dimension of $r(mp)$ jumps by one at that point) or the successive numbers stay the same (in which case the dimension of $r(mp)$ stays the same at that point).

Eventually, for m large, $h^0(K-mp) = 0$. When $m = 0$, this is g. There have to be exactly g jumps in the sequence of dimensions, as a result. Thus:

> There are exactly g values $m_1 < \cdots < m_g$ where $h^0(K-(m_i+1)p) < h^0(K-m_i p)$.

Again, to say that $h^0(K-mp) = h^0(K-(m-1)p)$ is to say that there exists a meromorphic function f on C such that $\operatorname{div}(f)_\infty = mp$; that is, a meromorphic function with exactly a pole of order m at p (and no other poles). The conclusion is:

18.2 Proposition. *There exist exactly g values $m \in \mathbb{N}$ such that there is no meromorphic function f with $\operatorname{div}(f)_\infty = mp$.*

Alternatively, let H be the set of $m \in \mathbb{N}$ such that there exists f with $\operatorname{div}(f)_\infty = mp$.[14] This is a *semigroup.* Then the proposition's claim is that

$$|\mathbb{N} \setminus H| = g.$$

18.3 Definition. The sequence $m_1 < \cdots < m_g$ enumerating the elements of $\mathbb{N} \setminus H$ is called the **Weierstrass gap sequence.** (This depends on the point p as well as the curve C.) Note that they are all at most $2g$, by Riemann-Roch.) Also, H is called the **Weierstrass semigroup.**

Let's state a couple of interesting theorems, and at some point it will become obvious why they are true.

[14] You can include zero or not as you like in the definition of $\mathbb{N}$.

18.4 Theorem. *For any C, then for all but finitely many $p \in C$, the Weierstrass gap sequence is just $\{1, 2, \ldots, g\}$.*

In other words, for a general $p \in C$, there won't be rational functions holomorphic everywhere with only a pole of order $\leq g$ at p.

Now some definitions.

18.5 Definition. We say that p is a **Weierstrass point** if the Weierstrass sequence is *not* $\{1, 2, \ldots, g\}$, or equivalently if the semigroup is not $[g + 1, \infty)$.

A Weierstrass point is a point that doesn't conform to the generic picture here. Equivalently, a point p is a Weierstrass point and only if

$$r(gp) > 0.$$

The last thing is the notion of the *weight* of a Weierstrass point.

18.6 Definition. We say that the **weight** of a point p, $w(p)$, is the extent to which the Weierstrass sequence is not the usual one,

$$w(p) = \sum m_i - \binom{g+1}{2}$$

where the $\{m_i\}$ is the Weierstrass sequence of p. So $w(p) > 0$ if and only if p is a Weierstrass point.

The second result is:

18.7 Theorem. *The sum of the weights on C, $\sum_{p \in C} w(p) = g^3 - g$.*

Let's prove these theorems. These proofs will be groanworthy; it's very simple.

Proof. A given number m is, by what we just said, a gap for a given point $p \in C$ if and only if $h^0(K - mp) < h^0(K - (m - 1)p)$. That is the case if and only if there exists a holomorphic differential on the curve, $\omega \in H^0(C)$, which vanishes at the point p to *order exactly* $m - 1$. In other words, the conclusion is that if we look at the canonical linear series, then the gap sequence of a point is just the *vanishing sequence* of the canonical series K at that point. That is, the gap sequence is just $a(H^0(K), p) + 1$, where the a notation as is from last time.

Remember what we said last time. If we have an arbitrary linear series and a space of sections, there is a certain number of orders of vanishing at a point, and here they correspond to the Weierstrass gaps. I.e., p is not a Weierstrass point if and only if the Weierstrass sequence is exactly $\{1, 2, \ldots, g\}$, which is to say that $a(H^0(K), p) = \{0, 1, \ldots, g - 1\}$; thus, by the fact proved earlier that—*in characteristic zero*—a linear series is uninflected at a general point, we find that most points are not Weierstrass points.

Is this true in characteristic p? There do exist linear series in characteristic p where every point is inflectionary. Is this theorem on Weierstrass points, though, true in characteristic p? I don't know.

The second theorem follows by the Plücker formula applied to the canonical series. ▲

Here is a useful fact.

Remark. On a general curve, all Weierstrass points have weight one. (Weierstrass points of weight one are called **normal Weierstrass points.**)

§3 Examples

We don't have a whole lot of time, so we'll just do a few examples and then call it quits. A key point to remember is, in fact, today we introduced two new *intrinsic invariants* of a curve: the inflectionary points of the canonical series, and the Weierstrass points. And we have just seen (by Riemann-Roch) that those are the same. Incidentally, one can talk about inflectionary points of the pluricanonical series, but they don't have the same interpretation as do Weierstrass points.

18.8 Example. Take $g = 2$. There isn't much to say. There are two possible Weierstrass semigroups; there is $\{2, 4, 5, \dots\}$, and then there is $\{3, 4, 5, \dots\}$. The second one is the generic, non-Weierstrass possibility. The first one is the Weierstrass one. It occurs at the six ramification points of the canonical map $C \to \mathbb{P}^1$ (which is two-sheeted).

In genus two, all curves look similar as far as linear series are concerned.

18.9 Example. For genus three, there are four semigroups: $\{4, 5, 6, \dots\}$ (the non-Weierstrass points), $\{3, 5, 6, \dots\}$ (which is weight one), and $\{3, 4, 6, 7 \dots\}$ (which is weight two), and the hyperelliptic semigroup $\{2, 4, 6, 7, 8, \dots\}$ (weight three). The last case (which may or may not exist) naturally implies that the curve is hyperelliptic, because it posits the existence of a meromorphic function with exactly one double pole which is holomorphic elsewhere.

The total number of Weierstrass points (including weights) is 24. In the hyperelliptic case, these can be accounted for by the branch points of the hyperelliptic map: we get eight ramification points. There are no other Weierstrass points, since we've already accounted for them.

In the non-hyperelliptic case, the fourth case is ruled out, and we can think of the curve as a smooth plane quartic. The Weierstrass points of weight one and two correspond to ordinary flexes on the curve (a point on the curve where the tangent has contact of order three) or a hyperflex (a point where the tangent line has contact four).

If β is the number of hyperflexes and α is the number of ordinary flexes on the canonical curve, then we have said that

$$\alpha + 2\beta = 24$$

(Plücker formula), and a priori there are thirteen possibilities of (α, β). One might ask which possible pairs can occur for smooth plane quartics; it turns out that not all can.

In higher genera, there is a non-Weierstrass semigroup $\{g + 1, g + 2, \dots\}$, a *unique* semigroup of weight one that looks like $\{g, g + 2, g + 3, \dots\}$, and *two* semigroups of weight two, $\{g - 1, g + 2, g + 3, \dots\}$ and $\{g, g + 1, g + 3, \dots\}$. If you're interested how curves behave in families, there are two *subvarieties* in the moduli space of curves that correspond to curves with these different types of Weierstrass points.

The natural question is, which of these semigroups actually occur as Weierstrass semigroups? Does every semigroup actually occur? What I mean here is this—if we have a semigroup $H \subset \mathbb{N}$, of finite index in the natural numbers, i.e. the cardinality of the complement is finite, can we find a Riemann surface of genus $|\mathbb{N} \setminus H|$ and a point whose Weierstrass semigroup is exactly H?

Answer: **Not all occur.** This was resolved in the 1980's; you can exhibit a semigroup that doesn't occur. But it's not obvious, and it doesn't occur in low genus (genus 16 was the example). *But*, we still don't know which *do* occur.

Lecture 19
11/9

Today is going to be a relatively short lecture. As we said last time, we will be talking about real curves. There is no Friday this week, so next Wednesday we stop fooling around and get serious. I don't mean to say anything negative about real curves, though. It's a worthy object of study. But you see now that the theory of algebraic curves, or our understanding of them—as with algebraic geometry in general—is sort of a mix of fairly specific and concrete questions that nonetheless fit together into an overall understanding of the subject.

There is also a sort of broader theory of algebraic curves, and if it's not too late already, I'd like to balance the two: to acquaint you with both the more specialized, local questions that one can ask about algebraic curves and about the general topic of Brill-Noether theory. I really would like to talk about theta characteristics at some point, but eventually we should get back to this general theory. The content of next Wednesday and later will be chapters four and five of the text: we will try to describe the varieties, or more properly the schemes, that parametrize linear series on a curve. We'll then use that description to prove theorems about linear series.

§1 Real algebraic curves

Let's start talking about the following. C will be a **real algebraic curve.** Back in the days when people first started doing algebraic geometry, when people started thinking about polynomials in multiple variables, naturally they looked at polynomials over the *reals.* The first thing they see is that the zero locus of a two-variable polynomial is more than a finite set of points—it has a *geometry,* related to the properties of the polynomial.

People eventually passed from working over $\mathbb{R}^2$ to $\mathbb{RP}^2$, and then to $\mathbb{CP}^2$. But today—just today—we'd like to think about real polynomials. We'll take an intermediate position. We're interested in curves in $\mathbb{RP}^2$. (Working in $\mathbb{R}^2$ complicates things.)

19.1 Definition. A **real algebraic curve in** $\mathbb{RP}^2$ is the zero locus $V(F)$ of a homogeneous polynomial $F \in \mathbb{R}[X, Y, Z]$.

Of course, associated to such a curve, is the set $C_{\mathbb{R}}$ of $\mathbb{R}$-valued points. (What we're more familiar with is the space $C_{\mathbb{C}}$ of complex-valued points.) The picture in this case is fundamentally different from the complexes. If we look at all curves $C \subset \mathbb{P}^2$ over

$\mathbb{C}$ of degree d, then they are just parametrized some large projective space, and the singular ones form a hypersurface. In a complex manifold, when you remove a proper submanifold, you get something connected. Thus the family of *smooth complex curves* is connected, and from that one can see that there is exactly one homomeomorphism type of a smooth complex curve of a given degree.

However, if you take the space of all real curves of degree d, then again they are parametrized by some $\mathbb{RP}^N$, and the singular ones are parametrized by a hypersurface (corresponding to the discriminant). The singular ones *do* disconnect the space. So there are many different topological types of plane curves of a given degree. (Nobody knows how many.)

One of the many reasons that things are easier over $\mathbb{C}$ is that the discriminant hypersurface (corresponding to the singular curves) doesn't disconnect the space of all curves.

OK, pick a real curve C. If C is smooth,[15] then $C_{\mathbb{R}}$ is a compact smooth submanifold of $\mathbb{RP}^2$,[16] and in particular it's a disjoint union of S^1's—there aren't a whole lot of 1-dimensional manifolds. The first question you might ask:

> How many copies? I.e., how many connected components does $C_{\mathbb{R}}$ have?

19.2 Definition. The connected components of $C_{\mathbb{R}}$ are called **ovals** or **circuits**.

I don't know why the word "oval" got into it. This is all old language. An oval is just an embedded copy of S^1 inside $\mathbb{RP}^2$. The first thing to do is to make a distinction between the two cases.

If we have $\gamma \subset \mathbb{RP}^2$, an imbedded curve, there are two possibilities:

1. The simpler one, which you would draw, is one homologous to zero in $H_1(\mathbb{RP}^2, \mathbb{Z}) = \pi_1(\mathbb{RP}^2) = H_1(\mathbb{RP}^2, \mathbb{Z}/2)$. This is the case exactly when the complement is disconnected. In this case, one component of the complement is homeomorphic to a disk (called the interior), and one component of the complement is homeomorphic to a Mobius band (called the exterior). These are called **even ovals.**

2. Or, γ is not homologous to zero, so γ is the nonzero element of $H_1(\mathbb{RP}^2, \mathbb{Z})$. This means that the complement is connected. Those are things like a *line* in $\mathbb{RP}^2$. This is a classic example of a cycle not homologous to zero. These are called **odd ovals.**

The one observation to make in connection with this is that there's an intersection pairing mod 2 on H_1, which is nondegenerate (i.e. nonzero). The self-intersection of the nonzero element is one. The conclusion is:

19.3 Proposition. *Any two odd ovals in $\mathbb{RP}^2$ must meet.*

This means in particular that:

[15] This means that C is smooth over $\mathbb{C}$.

[16] Indeed, to say that C is smooth is to say that the defining polynomial equation in affine coordinates doesn't simultaneously vanish with its derivatives. So you use the implicit function theorem to see that $C_{\mathbb{R}}$ is smooth. The converse is certainly false. If $C_{\mathbb{R}}$ is a manifold, C needn't be smooth.

19.4 Proposition. *If C is smooth, then $C_{\mathbb{R}}$ has at most one odd oval, and possibly many even ovals.*

Assume C is smooth. Then, by the above, if the degree of C is even, then C has *only* even ovals, while if the degree of C is odd, then C has *exactly one* odd oval.

§2 Singularities

We should now say something about singularities. One topic that I'd like to talk about in this course, which hasn't been done yet, is the classification of singularities of curves. It's a topic where you can actually answer a lot of questions, so that it's satisfying intellectually. But there are also potentially fertile areas of investigation.

The situation over $\mathbb{R}$ is, as always, is more complicated. There are two different types of singularities over $\mathbb{R}$ that become nodes over $\mathbb{C}$.

19.5 Definition. There are **crunodes**, which look like two lines meeting, and **acnodes**, where there is just a point.

In both cases, there are two complex arcs crossing transversely. In the second case, though, the complex arcs don't go through the real plane.

In the first case, the completion of the local ring is $\mathbb{R}[[x, y]]/(x^2 - y^2)$. In the second case, the completed local ring is $\mathbb{R}[[x, y]]/(x^2 + y^2)$.

§3 Harnack's theorem

The main theorem we want to prove today is Harnack's theorem. Let me give you a very somewhat unnecessarily constrained version of the theorem, and an unnecessarily complicated proof.

19.6 Theorem (Harnack)**.** *Let C be a real, smooth algebraic curve in the plane of degree d. Then the number of ovals of $C_{\mathbb{R}} \subset \mathbb{RP}^2$ is at most $\binom{d-1}{2} + 1$, and this is sharp.*

For example, let's just look at some cases.

19.7 Example. When $d = 2$, the number of ovals could be zero: there are conics with no points. Or you could have one oval, and this is more typical: you'll have an ellipse (or a hyperbola—in the affine plane this is disconnected but the two branches connect in projective space, so this has one oval[17]).

Remark. If you have a curve, say $f(x, y) = 0$, then you can perturb it to get curves $f(x, y) = \epsilon$ that hug the shape of the initial curve and which tend to that in the limit as $\epsilon \to 0$. That may give rise to different topological types.

Let's continue with the examples.

[17] This again is an illustration of the fact that working over the affine plane is harder.

19.8 Example. Consider plane cubics, $d = 3$. By Harnack's theorem, they have at most 2 connected components. They must have at least one, since there is at least one odd oval. The number of connected components is one or two. You have seen pictures of both these types of cubics. When you have something of the form $y^2 = f(x)$, when f has one real root or three real roots, you get different things.

19.9 Example. This will be something possibly new. $d = 4$. Here, the number of connected components is a priori anything between 0 and 4, according to Harnack's theorem. Here 0 occurs, e.g. $X^4 + Y^4 + Z^4 = 0$. Can we find one with four connected components? Yes, because Harnack's theorem is sharp.

How do we find one with four connected components? The strategy is to take a singular curve and deform it. Consider the product of two ellipses with equations f, g. The union of these curves has equation fg. Consider the perturbed curve $fg \pm \epsilon$. Then, you might get four components.

§4 Nesting

One further question you can ask, and we'll come back to this in a moment, is the following. If you have more than one even oval, you can still ask about whether ovals are *nested* or not.

19.10 Definition. If we have two even ovals $\gamma, \gamma' \subset \mathbb{RP}^2$, we say they are **nested** if γ lies in the interior of γ' or the other way around. They are **not nested** if each lies in the exterior of the other.

If you deform the product of two conics, you can get two nested ovals, or four mutually non-nested ovals. You might ask what the relations on nesting that can occur.

Question: Describe the nesting of the ovals in a smooth curve.

There are a number of restrictions on this. For quartics, you have up to four ovals. If you have two ovals, there are two types of nesting that can occur. However, if you have three ovals in a quartic, then the claim is that no two are nested. This follows by Bezout and parity. If two are nested and the other isn't, then by picking a point in the middle of the inner one and a point inside the non-nested oval, and drawing a line, you would get at least six intersections.

In general, the whole question to describe the nesting of ovals in a smooth curve is **wide open.**

19.11 Example. Take sextics. These can have up to eleven ovals; let's suppose you have all the possible eleven. You can have nesting. You can't have two nested pairs, by the Bezout theorem, or a triple nesting if there are eleven ovals.

Let's assume there is a nested pair, γ, γ' such that $\gamma \subset \mathrm{Int}(\gamma')$. Then all the other curves either lie in $\mathrm{Int}(\gamma')$ or outside that. You might ask what possible combinations occur. How many lie outside and how many lie inside $\mathrm{Int}(\gamma')$? Not all possible combinations occur. There is a complicated, messy, theorem that the difference of these two numbers (the number of curves in $\mathrm{Int}(\gamma') - \mathrm{Int}(\gamma)$ and $\mathrm{Ext}(\gamma')$) is divisible by four.

§5 Proof of Harnack's theorem

We won't prove the sharpness of the bound. The sharpness is proved by deforming products of curves of smaller degree. In Coolidge's book on "Plane Algebraic Curves," the only math book I've seen with a fold-out, there is a proof of the existence of real curves with lots of ovals.

Let's prove the theorem. It's a *very simple* proof.

Proof. Suppose C is a smooth plane curve of degree d. Suppose there are m ovals on the real points $C_{\mathbb{R}}$. Let's assume, moreover, that $m > \binom{d-1}{2} + 1$. Call these ovals $\gamma_1, \dots, \gamma_m$. At most one of these is an odd oval, so let's say that $\gamma_1, \dots, \gamma_{m-1}$ are all even ovals. We don't know about γ_m, which could be either even or odd.

Here's the basic idea. We're going to just choose a point $p_i \in \gamma_i$, on the oval γ_i. I want to find $Y \subset \mathbb{P}^2$ a general plane curve of degree $d-2$ (you'll see the reason in just a moment) passing through the points $p_1, \dots, p_{\binom{d-1}{2}+1}$. These are points on the first $m-1$ ovals because of the hypothesis on m. The reason we can find such a plane curve is that the dimension of the space of curves of $d-2$ in projective space is $\binom{d}{2}$. So we can find a degree $d-2$ curve Y passing through the $\binom{d-1}{2} + 1$ points $p_1, \dots, p_{\binom{d-1}{2}+1}$.

But we have some degrees of freedom left over. Choose $q_1, \dots, q_{d-3} \in \gamma_m$, and we can arrange that our $d-2$-degree curve Y passes through all these. We can do this because there are a "lot" of curves of degree $d-2$, as we saw above.

Now the observation is that if a curve passes through a point on an even oval, it comes in—and it has to go out. (Unless there is a tangency, but then there will be intersection multiplicity ≥ 2.) So the intersection $Y \cap C$, with multiplicity, is at least $2(\binom{d-1}{2} + 1) + d - 3 = d(d-2) + 1$. The factor of two comes because Y intersects each $\gamma_1, \dots, \gamma_{\binom{d-1}{2}+1}$, and those it must thus intersect twice. This contradicts Bezout's theorem. ▲

Let me very briefly mention a few problems related to plane curves. One is *convexity.* One aspect of the real numbers, which $\mathbb{C}$ doesn't share, is the total order. One can talk about the *convexity* of an oval. Here's a question.

> In addition to specifying the number of ovals, and their nesting properties, one can ask further how many of them are convex.

There are some restrictions based on Bezout's theorem. Given a quartic with two nested ovals, the inner one had better be convex, or there would be a line intersecting it at four points and the outer one as well. But the exterior oval might be very non-convex.

One last bit of extrinsic geometry. You can ask about other phenomena like bitangents and flexes. For instance, for a smooth quartic, we know over $\mathbb{C}$, there are 28 bitangents by Plücker's theorem. They can all be real, it turns out. However, in higher degree, I don't know whether all the bitangents can all be real. It's an open problem. However, for the issue of flexes of a real plane curve, there's a bizarre theorem: at most *one third* of the flexes over the complex numbers of a smooth plane curve can be real.

Here's another bizarre question. For a moment, forget curves. Just think of points in $\mathbb{RP}^2$.

> Can I have a finite configuration of points such that the line through any two contains a third? (Besides the trivial case of all the points on a line.)

The answer is that you can't. This is easy to see intuitively if you try to write this down. But it can also be proved.

However, you know such a configuration. Let's just look at plane cubics. If you look at the flexes of a plane cubic, they correspond to torsion points of order three on an elliptic curve, if one of them is chosen at the origin. These form a subgroup of the points on the elliptic curve, and if you think about how the addition law on the elliptic curve works, the 3-torsion points are *precisely such a configuration of points:* the line through any two contains a third. It follows that all the flexes of a plane cubic cannot be real.

I just had to mention that.

Lecture 20
11/16

I want to start of this class with a quals problem. So, let's let $M_{a,b}$ be the space of a-by-b matrices, mod scalars. So that's a projective space $\mathbb{P}^{ab-1}$. Sitting inside here, I want to look at the subvariety $M^k_{a,b}$ of rank $\leq k$. The two statements are:

1. The codimension of $M^k_{a,b}$ is $(a-k)(b-k)$. (This you can do by the standard incidence correspondence.)
2. The singular locus of $M^k_{a,b}$ is contained in $M^{k-1}_{a,b}$. (Here you can use homogeneity.)
3. If $A \in M^k_{a,b}$ is a smooth point (i.e. not in $M^{k-1}_{a,b}$), then the projective tangent space at the point $T_A(M^k)$ is given by

$$T_A(M^k) = \{\phi : \phi(\ker A) \subset \mathrm{Im} A\}\,.$$

One way to remember all this is to note that the normal space to M^k in M at A is just $\mathrm{Hom}(\ker A, \mathrm{coker} A)$. In particular, the dimension of the normal bundle is equal to the dimension of the kernel times the dimension of the cokernel. That's the quals problem.

§1 Basic notions

Now let's get to work. There are very few lectures left. We will meet on Friday. If there is no "Basic Notions" talk, there will a class on Monday. Then, after Thanksgiving, we have one more week of classes. We'll probably have two classes that week. I really have to talk about the basic facts of life now, when it comes to curves.

The main object of study, when you talk about curves, are linear systems on curves. To describe linear systems on curves, we'll introduce some basic objects, most of which we've already encountered.

As always, let C be a smooth projective curve of genus g. We are interested in:

1. C_d is the set of effective divisors of degree d on C. We can think of this as C^d/S_d, and this is a smooth, d-dimensional variety.

2. We also have the variety $\mathrm{Pic}^d(C)$, which is just the set of simply line bundles of degree d on C. This is isomorphic to the Jacobian of C, which we described as $H^0(K)^\vee/H_1(C,\mathbb{Z})$. This is not obvious, but is given by the Abel-Jacobi map $u : C_d \to J(C)$ which sends a divisor $\sum a_i p_i$ to the sum of the integrals $\sum a_i \int_p^{p_i}$ (where p is some fixed basepoint). This defines a linear functional on $H^0(K)^\vee$, but because there might be many paths the linear functional is defined only mod $H_1(C,\mathbb{Z})$.

3. What we want to know now is the following. How many divisors are there on the curve that move in linear series? We can introduce C_d^r as the set of divisors D on the curve of degree D such that $r(D) \geq r$. This is a subvariety of C_d.

4. W_d^r is the collection of line bundles $\mathcal{L}$ such that $h^0(\mathcal{L}) \geq r+1$ of degree d, so this sits inside $\mathrm{Pic}^d(C)$.

5. Obviously C_d^r maps to W_d^r under the Abel map u. By Abel's theorem, again, the *fibers* of u are exactly the complete linear series of a given degree. So, in other words, we can think of W_d^r as the set of line bundles $\mathcal{L}$ such that the pre-image $u^{-1}(\mathcal{L}) \subset C_d$ has dimension at least r. By the upper semi-continuity of dimension, this is a closed subset of $\mathrm{Pic}^d(C)$. The pre-image C_d^r is thus closed as well.

Most of the questions we've asked so far are about existence: when are these empty or nonempty? Now we want to talk about the irreducibility, singularity, smoothness, dimension, etc. of these varieties. Note that these spaces W_d^r, C_d^r have the structures of *varieties* as closed subvarieties of $C_d, \mathrm{Pic}^d(C)$.

§2 The differential of u

We're eventually going to see that there is, further, a *scheme* structure on W_d^r, C_d^r, after which we can say that these objects actually represent suitable functors. This scheme structure is not always reduced. That's our goal.

So, again, we're starting with the map $u : C_d \to J(C)$ which sends a divisor $D = p_1 + \cdots + p_d$ up to $\sum \int_p^{p_i}$. The key thing is to understand the *derivative* of this map. Luckily, we can write it down fairly explicitly. Here's the point. Let's suppose for the moment the $\{p_i\}$ are distinct points. Then that means at least étale locally, the symmetric product looks like the ordinary product—there's not a fixed point. So, the tangent space $T_D C_d = \bigoplus T_{p_i} C$. If you want to vary the divisor, then you vary each point individually. Of course, the tangent space at $u(D)$ of the Jacobian is the same as the tangent space everywhere. The Jacobian is the quotient of a vector space mod a lattice, so its tangent space everywhere is $H^0(K)^\vee$.

So, now, we want to know: **what is the differential $du : T_D C_D \to T_{u(D)} J$?**

20.1 Proposition. *du sends a tangent vector v_i at $p_i \in C$ to the linear functional $H^0(K) \to \mathbb{C}$ that sends a $\omega \mapsto \omega(p)(v_i)$.*

If I have a tangent vector to the curve C at the point p_i, which corresponds to direction in which to vary the point p_i, then du sends that to the value of the differential

ω at p, evaluated at v_i. Think of it this way: if we take the transpose of the differential, the codifferential, that's a map $H^0(K) \overset{\simeq}{\to} \bigoplus T^\vee_{p_i}C$, and the latter is the sum of the canonical bundles at each p_i. So alternatively, it's the natural map

$$H^0(K) \to H^0(K/K(-D)).$$

The map, when thought of as a codifferential, is just the tautological evaluation map.

We'll have to generalize these constructions to the case where the points are not all distinct. Suppose the $\{p_i\}$ in D are not necessarily all distinct. We then need a better way of describing the tangent space $T_D(C_d)$. For that, I want to remind you of a basic identity. If $p \in C$ is a point, then the cotangent space $T^\vee_p C = \mathfrak{m}_p/\mathfrak{m}_p^2$. If I want to think of this as a sheaf, it's just the skyscraper sheaf supported at a point. So a cotangent vector is just a function vanishing at a point modulo functions vanishing to order two. The claim is that we can dualize this and describe $T_p(C) = H^0(\mathcal{O}_C(p)/\mathcal{O}_C)$; that is, functions near p with at most a simple pole of p modulo functions which are holomorphic at p. There is an obvious pairing between this space and $\mathfrak{m}_p/\mathfrak{m}_p^2$ given by multiplication and evaluation. So we have a way of describing the tangent space in a nice way.

Consequently, we can write for D a divisor of distinct points

$$T_D(C_d) = H^0(\mathcal{O}_C(D)/\mathcal{O}_C).$$

This is *in general* given in this way by the previous paragraph (when the points are distinct, by what we saw earlier about the tangent space splitting as $\bigoplus T_{p_i}C$), and then you can extend this identity over all divisors.

20.2 Proposition. *For any divisor D,*

$$T_D(C_d) = H^0(\mathcal{O}_C(D)/\mathcal{O}_C).$$

Moreover,

$$T^\vee_D(C_d) = H^0(K/K(-D)).$$

In this setting, the pairing between the tangent and cotangent spaces is given by the sum of the residues of the product.

So now we have a description of the tangent and cotangent spaces to the symmetric product. Again, the codifferential of the map u is exactly as we described.

Remark. It is not obvious that C_d is smooth. The only varieties whose symmetric products are smooth are smooth curves.

The symmetric product of surfaces is necessarily singular. If we have a surface S, and take $S \times S/\Sigma_2$, then the fixed point locus of that involution is the diagonal, which is codimension two. If the quotient were smooth, then the branch divisor of $S \times S \to \mathrm{Sym}^2 S$ would have to be codimension one. (That's the purity theorem.)

For a curve, you can see smoothness for $\mathbb{P}^1$: $\mathrm{Sym}^d\mathbb{P}^1 = \mathbb{P}^d$. The construction of the symmetric powers is a local construction. So that's how you can prove it in general.

In any case, what I'm going to suggest is that we simply focus on the case where the points are distinct. At the end of the day, we can check that whatever we do extends over the diagonal, the locus of non-reduced divisors.

All right. What's the point of all this? It tells me a lot right off the bat. The differential du is, as we saw, an evaluation map. At a point (divisor) $D = p_1 + \cdots + p_d$, then we get a map $\bigoplus T_{p_i} C \to H^0(K)^\vee$. What is the image? Just those linear functionals on $H^0(K)$ that are obtained by *evaluation* at points of the divisor. The kernel of these linear functionals consists of forms which vanish at the points of the divisor. It follows that the image of du_D consists of the *annihilator* of this kernel, so in other words,

$$\mathrm{Im} du_D = \mathrm{Ann}(H^0(K-D) \subset H^0(K)) \subset H^0(K)^\vee.$$

The rank of the image is thus given by $g - H^0(K-D)$.

If we consider the dimension of the *kernel* of u_D, then that will have dimension exactly equal to $d-(g-h^0(K-D))$. By geometric Riemann-Roch, exactly the dimension of the fiber. In other words, even when the map u_D is not smooth, where it has a fiber, then the *kernel* of the differential is *exactly* the tangent space to the fiber. We know that set-theoretically, the fibers are projective spaces—complete linear systems. In fact, this is true scheme-theoretically too. The differential at each point has *maximal rank given the fiber dimension.*

This is going to be an important part of the following discussion. Again, the important point is, the dimension of the fiber is equal to the dimension of the kernel of the differential. One consequence of this is:

Before, we characterized C_d^r is the space of divisors such that the fiber $u^{-1}(u(D))$ has dimension $\geq r$. This is a closed subset. But now, thanks to this equality, this is also the set of D such that the *rank* of the differential du_D is $\leq d-r$; that is, there's at least an r-dimensional kernel. The differential, however, is not just a map of vector spaces; it's a map of *vector bundles.* Given the Abel-Jacobi map $u : C_d \to J$, then there is a map

$$du : TC_d \to u^* TJ$$

and the space C_d^r is the space of points where the kernel of this map of vector bundles which is of a sufficiently large rank. As a result, we can give C_d^r a *scheme structure.*

Remark. In general, when you have a map $E \to E'$ of vector bundles, then you can define a subscheme of the base on where the rank of the map on fibers $E_x \to E'_x$ is bounded. This can be defined using linear algebra.

This observation that we can characterize C_d^r by the rank of the map on vector bundles, rather than simply the rank of the fiber, while really a trivialty, will be very useful.

Let's move further afield. What I want to do next (and this is going to have to wait until next time) is to describe dimension estimates on W_d^r and C_d^r, and also describe the tangent spaces to these. That's why I brought up the quals problem. If we think of the C_d^r as given by the locus of points where a certain matrix has rank $d-r$ or less, then we can think of this locally as a map to the matrix varieties described earlier. That's how we will describe the tangent spaces.

§3 Marten's theorem

We are now going to continue with questions of the form, what is the dimension of W_d^r? How many line bundles of a given degree and dimension can there be? Let me just give you the statement.

20.3 Theorem. *The dimension* $\dim W_d^r$ *is at most* $d-2r$ *with equality only if* $r=0$ *or* C *hyperelliptic.*

In some sense, this is a version of Clifford's theorem for families of line bundles. This will emerge on Friday as a consequence of the description of the tangent spaces. But we also want to approach it from a naive point of view. We're going to use the statement proved earlier about the uniform position lemma to deduce this.

Proof. We're looking at divisors. C_d^r is the locus of divisors $D \subset C$ that lie in a $\mathbb{P}^{d-r-1}$ under the canonical map (by geometric Riemann-Roch). In other words, they fail to impose independent conditions on the canonical series by r. Let's assume first that C is not hyperelliptic.

If we have such a family of divisors, the natural thing to do is to relate them to something we know about—hyperplanes in the canonical space. Consider the incidence correspondence $\Sigma = \{(H, D) : D \subset H \cap C\}$ where H is a hyperplane in $\mathbb{P}^{(g-1}$. So we have

$$\Sigma \subset (\mathbb{P}^{g-1})^* \times C_d^r.$$

We have natural projection maps $\Sigma \to \mathbb{P}^{g-1*}$ and $\Sigma \to C_d^r$, where the latter sits as a generically projective bundle over W_d^r. The general fiber of $C_d^r \to W_d^r$ is isomorphic to a $\mathbb{P}^r$.

This is something that should have been said earlier, but I'll say it now. That is, $W_d^{r+1} \subsetneq W_d^r$, and in fact $W_d^r \setminus W_d^{r+1}$ is dense in W_d^r. This fact is based on the following observation. If we have some D, and we want to compare D to $D-p$, then $r(D)$ and $r(D-p)$ will be either equal or differ by one. The first situation, when they're equal, happens for *finitely many points*, namely the basepoints of the linear series. If D is an effective divisor, then for all but finitely many points $p \in C$, we have $r(D) = r(D-p)+1$. The exceptions are simply the base points of the linear series, if there are any. You can also apply the same thing dually to $K-D$, if it's effective. It follows that for all but finitely many points p, we also have $r(D+p) = r(D)$. (Notice that this is a little different than when you subtract a point.) Thus, given a $D \in C_d^r$ (let's say $0 \le r < d-g$), for a *general* $p, q \in C$ we have

$$r(D+p-q) = r(D) - 1.$$

So we can always find nearby divisors where the value of r drops. This is why the general fiber of $C_d^r \to W_d^r$ is generically a $\mathbb{P}^r$.

OK, let's go back to the incidence correspondence. The general point in C_d^r is a point moving in an exactly r-dimensional series, so lying in exactly a $\mathbb{P}^{d-r-1}$-plane by geometric-Riemann-Roch. It follows that the generic fiber of $\Sigma \to C_d^r$ is a $\mathbb{P}^{g-d+r-1}$. (How many hyperplanes contain a $\mathbb{P}^{d-r-1}$...?) (If the fibers were strictly larger, then they'd only help the rest of the argument. So this was actually unnecessary.)

Anyhow, one finds

$$\dim \Sigma = \dim W^r_d + r + g - d + r - 1.$$

This is from the multiple fiberings. Now you go the other way. The second projection, $\Sigma \to (\mathbb{P}^{g-1})^*$. Take a pair, forget the divisor, and just look at the hyperplane. But if you take a hyperplane, there are only finitely many divisors which can be supported in that hyperplane. So $\Sigma \to (\mathbb{P}^{g-1})^*$ is a finite map. It's also not surjective; a general hyperplane will intersect the curve in points which are in linear general position. So a general point in $(\mathbb{P}^{g-1})^*$ is not in the image of Σ. It follows that

$$\dim \Sigma \leq g - 2.$$

It follows that

$$\dim W^r_d \leq d - 2r - 1.$$

This is a bit stronger than the statement in the theorem; when C is hyperelliptic, we can get up to $d - 2r$. But that's because we know what the special linear series are on a hyperelliptic curve. ▲

Lecture 21
11/18

§1 Setting things up again

Today we will talk about the various theorems that are known about the behavior of linear series on curves. What I'd like to do in the last week is to take a sample of the body of theorems (a fairly large body) of this nature, and indicate proofs. Let's again start with the set-up. We have a smooth projective curve C of genus g, etc. We want to look at the diagram of objects as follows.

1. The dth symmetric product C_d, which maps via $u : C_d \to \mathrm{Pic}^d(C) \simeq J(C)$.

2. Inside C_d, we have C^r_d, the variety of divisors which moves in an r-dimensional pencil. This maps to W^r_d, the family of line bundles in $\mathrm{Pic}^d(C)$ which have at least $r + 1$ independent global sections.

3. We also have the variety G^r_d of pairs $(\mathcal{L}, V)$ where $\mathcal{L} \in \mathrm{Pic}^d(C)$ and $V \subset H^0(\mathcal{L})$ has dimension r. This maps to W^r_d, and essentially it "blows up" along W^{r+1}_d (the map $G^r_d \to W^r_d$ is birational, as follows from the fact that we proved last time that $W^{r+1}_d \subsetneq W^r_d$ is nowhere dense).

4. We recall that if $D \in C_d$, then we can realize

$$T_D C_d = H^0(\mathcal{O}_C(D)/\mathcal{O}_C), \quad T^*_D C_d = H^0(K/K(-D))$$

and the duality comes from Serre/Grothendieck/Kodaira duality or simply via residues.

5. The tangent space to the Jacobian $J(C)$ is $H^0(K)^*$. The cotangent is $H^0(K)$.
6. The codifferential[18] of the Abel-Jacobi map $u : C_d \to J(C)$ gives a map
$$H^0(K) \to T_D^* C_d = H^0(K/K(-D))$$
and it is given by restriction.

These are the basic objects we'll be talking about, and now I want to make the basic construction, which I talked about last time. Then we can start stating theorems.

The idea is, if $D = p_1 + \cdots + p_d$ with the p_i distinct (we're in the complement of the diagonals), then near D, the locus C_d^r is just the locus where the *rank* of a certain matrix drops; that matrix is the evaluation matrix on differentials. That is, if $\omega_1, \dots, \omega_g$ are a basis for $H^0(K)$, then it is the matrix

$$\begin{bmatrix} \omega_1(p_1) & \dots & \omega_1(p_d) \\ \dots & \dots & \dots \\ \omega_g(p_1) & \dots & \omega_g(p_d) \end{bmatrix}.$$

The entries of this matrix make no sense, but the rows are *homogeneous*: they are well-defined up to multiplication by a scalar. They're all elements of a one-dimensional complex vector space, which just isn't canonically identified with $\mathbb{C}$. So the rank condition is independent of any choices we might make. Alternatively, we could just write this out in local coordinates. The condition is that the rank of this matrix have rank $\leq d - r$, if $D \in C_d^r$. This is Riemann-Roch. It is simply saying that d points move in an r-dimensional linear series exactly when they fail by r to impose independent conditions on the canonical series.

More generally, if we allow the divisor D to be arbitrary, this is a *matrix representative* of the evaluation map of vector spaces

$$H^0(K) \to H^0(K/K(-D))$$

and both of these are naturally the fibers of a vector bundle on C_d. The first comes from a trivial vector bundle, and the second comes from the (generally nontrivial) vector bundle of rank d: the cotangent bundle T^*C_d. This matrix is then a representation of the fiber of the above map of vector bundles

$$H^0(K) \times C_d \to T^*C_d.$$

The point is simply that we get, locally, a map from an open subset $U \subset C_d$ to the space $M_{d,g}$ of d-by-g matrices. As we said last time, sitting inside $M_{d,g}$ is the subspace $M_{d,g}^{d-r}$ of matrices of rank $\leq d - r$ or less. The pre-image of that is given by C_d^r, by what we've said. But $M_{d,g}^{d-r}$ has codimension $r(g - d + r)$ in $M_{d,g}$ and pulling back will only shrink the dimension.

The conclusion is:

21.1 Proposition. *The dimension of C_d^r is at least $d - r(g - d + r)$ everywhere (if it is nonempty).*

[18] This sounds like it should be a word.

Similarly, we find by applying the fiber dimenison theorem to $C_d^r \to W_d^r$ (which has generic fiber of dimension r)

21.2 Proposition. $\dim W_d^r = \dim C_d^r - r \geq g - (r+1)(g-d+r)$ *(if it is nonempty).*

That last number, $g-(r+1)(g-d+r)$, is the **Brill-Noether number.** This is the expected dimension of the variety of linear series of degree d and dimension at least r; we've seen that the dimension of this variety is at least the Brill-Noether number.

Remember what we said last time about the tangent spaces to these varieties. Given $M_{d,g}^{d-r}$, at a general point A (given by a matrix), the tangent space is given by those linear maps that carry $\ker A$ into $\mathrm{Im} A$.

Let's set some notation.

1. $\alpha : H^0(K) \to H^0(K/K(-D))$ is the straightforward evaluation map.
2. $\mu : H^0(D) \otimes H^0(K-D) \to H^0(K)$ is the multiplication map.

What we can conclude, in these terms, from the description of the tangent space to this matrix space, is the following.

21.3 Proposition. *If $D \in C_d^r \setminus C_d^{r+1}$ (so the corresponding Brill-Noether matrix has rank exactly $d-r$, not less), then the tangent space*[19] *to C_d^r at D, is given by the annihilator of $\alpha \circ \mu$.*

Similarly, $T_L W_d^r$ at a line bundle $L \in W_d^r - W_d^{r+1}$ is the annihilator of $\mathrm{Im}\mu$.

What does this mean? So μ is the multiplication map $H^0(D) \otimes H^0(K-D) \to H^0(K)$, and α sends this to $H^0(K/K(-D)) = T_D C_d$. The annihilator of the image of this is the tangent space to C_d^r. In the second case, recall that $T_L J(C)$ is $H^0(K)^*$ and the annihilator of $\mathrm{Im}\mu$ sits in here.

Remark. One can do a "reality check" here by counting the dimensions.

§2 Theorems

Now let's state some theorems. Fix g, r, d. The notation is the same as before.

21.4 Definition. We let ρ be the Brill-Noether number $g-(r+1)(g-d+r)$.

21.5 Theorem (Existence)**.** *If $\rho \geq 0$, then W_d^r is nonempty.*

The way this is proved is to go back to that map of vector bundles over C_d, and to consider the locus of points where the rank is at most $d-r$, as before. This locus corresponds to points of C_d^r. We showed that the dimension of C_d^r (or W_d^r) at each point was at least the Brill-Noether number. Of course, one has to show that there is actually such a point! Using the Porteus formula, one can compute the class in the Chow ring of this space, and show that it is nonempty. This is due to Kempf and Kleiman.

21.6 Theorem (Nonexistence)**.** *If $\rho < 0$, then for a general curve C, we have $W_d^r = \emptyset$. For a general curve C, we have* $\dim W_d^r = \rho$.

[19] Recall that $T_D^* C_d = H^0(K/K(-D))$.

Let us now state a result valid for *general* linear series on *general* curves. This will be a corollary of the proof of the non-existence theorem above.

21.7 Theorem. *Suppose* $\rho > 0$.

If $r \geq 3$*, then for a general curve and a general divisor in* C^r_d *on it, the associated map* $C \to \mathbb{P}^r$ *is an imbedding. When* $r = 2$*, the map is birational onto its image, which has only nodes. When* $r = 1$*, the map is simply branched.*

21.8 Example. For example, if C is a general curve of genus g, then the smallest degree of a nonconstant meromorphic function on C (i.e. the smallest degree of a map $C \to \mathbb{P}^1$ which isn't constant) is exactly $\frac{g+1}{2} + 1$ (the ceiling of that). If the genus is one or two, this gives you 2; if the genus is three or four, this gives you 3.

The smallest degree of a plane curve birational to C is $\frac{2}{3}g + 2$. The smallest degree of an imbedding is $\frac{3}{4}g + 3$ (again, the ceiling) when $g \geq 4$.

We next have a strong result, which implies much of what's on the board.

21.9 Theorem (Giserker-Petri)**.** *For a general curve* C *and* $\mathcal{L}$ *an* ***arbitrary line bundle*** *on* C*, the map*

$$\mu : H^0(\mathcal{L}) \otimes H^0(K - \mathcal{L}) \to H^0(K)$$

is injective.

This implies the emptiness of W^d_r when $\rho < 0$, because in that case the dimension of $H^0(\mathcal{L}) \otimes H^0(K - \mathcal{L})$ would be too large. We're not going to prove this theorem. There are two outstanding proofs of this, and I can describe them to you.

21.10 Example. Here is a consequence of the Giserker-Petri theorem.

If we look at the singular locus of W^r_d, it is W^{r+1}_d for a general curve C. Similarly, for a general C, $(C^r_d)_{sing} = C^{r+1}_d$. That follows by the description of the Zariski tangent space as the annihilator of these multiplication maps.

One can in fact use this to show that G^r_d is smooth for a general curve C.

21.11 Example. In particular, suppose $\rho = 0$, then for a general curve C, W^r_d has dimension zero. The previous example shows that W^r_d is *smooth*. In particular, it is a disjoint union of reduced points; there are no fuzzy points. Using Portier's formula, one can calculate the number of g^d_r's on the curve. One finds the answer:

$$g! \prod_{\alpha=0}^{r} \frac{\alpha!}{(g - d + r + \alpha)!}.$$

Remark. There is a duality

$$W^r_d = K - W^{g-d+r-1}_{2g-2-d}.$$

This comes from Riemann-Roch.

21.12 Example. If you want to express a general curve of genus four as a branched cover of $\mathbb{P}^1$, you need at least three sheets. That comes out of the case $g = 4, r = 1, d = 3$. The result one gets is that there are (generically) two g^1_3's, and that can be seen geometrically by representing the curve on a quadric surface and taking rulings. (For some special curves, the associated W^r_d turns out to be one non-reduced point, e.g. when there is one g^1_3.)

21.13 Example. Take $g = 6, r = 1, g = 4$. There are five ways to express a general curve as a four-sheeted cover of $\mathbb{P}^1$. This can be seen by writing down an explicit model for the curve.

21.14 Example. In genus 8, there are *fourteen* ways of expressing a general curve as a five-sheeted cover of $\mathbb{P}^1$. I don't know a good way of seeing this geometrically.

Since I'm stating a whole bunch of theorems, let me give you one more.

21.15 Theorem (Fulton-Lazarfeld)**.** *For C general, and $\rho > 0$, the scheme W^r_d is irreducible.*

This requires non-trivial theorems. Fulton and Lazarsfeld proved a generalization of the classical geometric fact that a positive-dimensional complete intersection in projective space is connected, for ample vector bundles. They were thus able to prove W^r_d connected, which easily implied irreducibility. Anyway, the basic idea is that this involves a theorem on degeneracy loci on vector bundles, which I'm not get into here.

One last statement. This is going to be a mess. I just warn you here, but this mess is nonetheless going to be a component of the proof we discuss later.

21.16 Theorem (Inflection)**.** *If C is a curve and $p_1 \dots p_\delta \in C$, consider linear series with specified ramification sequences at the p_i. That is, fix a sequence $\alpha^1_i, \dots, \alpha^r_i$ for each i, and consider the variety*

$$G^r_d(p_1, \dots, p_\delta) = \{(\mathcal{L}, V) \in G^r_d : \textit{the ramification sequence of the linear series } V \textit{ at } p_i \textit{ is at least } \alpha^\bullet_i.\}.$$

Then, for a general curve, the dimension of $G^r_d(p, \alpha)$ is $\rho - \sum \alpha^i_k$.

Although it is a mess, this is a pretty natural question to ask. As we'll see later, this statement is better-suited for an inductive approach.

21.17 Example (Genus four)**.** Let's just take a look at one example (which will arise on the last homework assignment). The claim is that the *scheme* of g^1_3's on a non-hyperelliptic genus four curve C, i.e. W^1_3, is degree two. If the curve is hyperelliptic, this is positive-dimensional (take the g^1_2 plus a random basepoint), so we've ruled out that case. By Marten's theorem, in this case, W^1_3 is finite.

You've seen the picture. The curve C sits inside a quadric surface, and it is the intersection of a quadric and a cubic, in the canonical model in $\mathbb{P}^3$. A triple of collinear points on the curve is a g^1_3, and we get two from the two rulings of the quadric. There are no other g^1_3's, because a line can intersect a quadric surface only at at most two points.

Can we explicitly verify that $H^0(\mathcal{L}) \otimes H^0(K - \mathcal{L}) \to H^0(K)$ is injective when $\mathcal{L}$ is in W^1_3? By counting dimensions, one expects this to be an isomorphism. Here's why it is an isomorphism. You don't have to think of them as tensor products, but rather as divisors. You have a $\mathbb{P}^1$ of divisors coming from $\mathcal{L}$, and a $\mathbb{P}^1$ divisors of $K - \mathcal{L}$, and taking the span of the pairwise sums. Is that surjective?

It is, by geometric arguments, when the quadric is smooth. When the quadric is not smooth, it's not. (*Some argument about pairwise sums of divisors and plane sections of curves and quadrics.*)

Lecture 22
11/30

What I promised you is a complete and self-contained proof of the Brill-Noether theorem, at least the nonexistence half. I said in my email that we were going to do this using the general theory of limit linear series, but that's sort of an overstatement. In the interest of giving a complete proof in a short period of time, I would like to keep the definitions to a minimum. In particular, I will not use the word "limit linear series," but if you listen to the next two lectures, you'll essentially know what one is.

§1 Recap

We do have to establish (and remind ourselves of) some notational conventions introduced earlier in the course.

1. C is *always* going to be a smooth projective curve.

2. A g^r_d on C consists of a pair $(\mathcal{L}, V)$ where $\mathcal{L} \in \mathrm{Pic}^d(C)$ and $V \subset H^0(\mathcal{L})$ is $r+1$-dimensional. (This is standard notation.) We just write V for short.

3. For $p \in C$, we define a sequence $\{a_0(V,p), \ldots, a_r(V,p)\} = \{\mathrm{ord}_p(\sigma), \sigma \in V^*\}$; this is the *vanishing sequence* of V at p.

4. We define $\alpha_i(V,p) = a_i(V,p) - i$ to be the *ramification sequence*, and set the total ramification to be $\alpha(V,p) = \sum_i \alpha_i(V,p)$.

5. We recall the *Plücker relation.* For any linear series V, then
$$\sum_p \alpha(V,p) = (r+1)(d + r(g-1))$$
(where V is a g^r_d, as usual).

6. Given (r,d), we define the **Brill-Noether number** (which also depends on the genus)
$$\rho(r,d) = g - (r+1)(g - d + r);$$
this is the expected dimension of the variety of linear series of type (r,d) (i.e. the g^r_d's). Given a linear series V which is a g^r_d, then we also write $\rho(V) = \rho(r,d)$.

We state a form of the Brill-Noether theorem.

22.1 Theorem (Brill-Noether)**.** *When the Brill-Noether number $\rho(r,d) < 0$, then there are no linear series on a general curve C of genus g. Alternatively, if C is general, then for any linear system on C, then $\rho(V) \geq 0$.*

The word "general" implicitly invokes the fact that the *set* of all curves of genus g has the structure of a variety (the moduli space). Thus we can talk about "general."

Remark. The condition that every linear series on a curve have non-negative Brill-Noether number is an open condition in the moduli space. As a result, to prove the Brill-Noether theorem, one can just exhibit a *single* curve such that every linear series on it has non-negative ρ, because the moduli space is irreducible. However, nobody has ever done this. All you have to do is write down a curve without any exceptional linear series, and you can do this in low genera (we've implicitly done this up to genus 5), and that's all fine, but in higher genera, nobody has yet written down the equations of a single curve which is general in this sense.

We want to prove something stronger. We want to consider not just linear series, but linear series plus ramification. Suppose we have now a curve C, always smooth and projective of genus g, as always, together with a specified collection of points $p_1, \dots, p_m \in C$. We're going to consider linear series *with specified ramification* at these points. We want to generalize the definition of ρ, as a result.

22.2 Definition. If V is a g^r_d on C, we define the **adjusted Brill-Noether number** with respect to the points $p_1, \dots, p_m$ to be

$$\rho(V, p_1, \dots, p_m) = \rho(V) - \sum_{k=1}^{m} \alpha(V, p_k).$$

We are going to prove the stronger version:

22.3 Theorem (Brill-Noether)**.** *If C and the points $p_1, \dots, p_m$ are general,*[20] *and V any linear series on C, then*

$$\rho(V, p_1, \dots, p_m) \geq 0.$$

There are various reasons why it's easier to prove this stronger version than the original version. Note also that if you know this result, then you can deduce the dimension statement. Suppose we consider the dimension of the variety of g^d_r's on a general curve, and it's bigger than ρ. Then we can impose conditions of simple ramification and the above theorem states that eventually it will become empty, not too slowly. Then that gives a bound on the dimension of the variety.

Remark. We won't invoke the existence of a moduli space of pointed curves.

§2 Families

A couple of things. This is a problem with a variational component. You need the "general" part of the theorem for it to be curve; there exist lots of plane curves,

[20] I.e., C general, and $p_1, \dots, p_m$ general on C.

hyperelliptic curves, etc., and all these have special linear series. So plane, hyperelliptic, etc. curves are not general. Anyway, we're going to need these notions in families.

Consider the following situation. Let $\pi : \mathcal{C} \to B$ (where in practice B will be a disk in the complex analytic world, or say the Spec of a discrete valuation ring) be a "family of curves," by which we mean that the fibers are smooth projective curves of genus g. We can think of this as a family of curves C_t, given by the fibers for $t \in B$. Note that the parametrizing space B is always one-dimensional.

What does a family of line bundles mean?

22.4 Definition. A **family of line bundles** on the family of curves C_b (of a given degree) $\mathcal{L}_t \in \operatorname{Pic}^d(C_t)$ is going to be a collection of restrictions $\mathcal{L}|_{C_t}$ where $\mathcal{L}$ is a line bundle on the total space $\mathcal{C}$. (The fact that the $\mathcal{L}_t$ come from the total space means that the $\mathcal{L}_t$ fit together continuously, in some sense.)

Next we need the notion of a family of linear series. This should be a choice of spaces of sections that fit together.

22.5 Definition. A **family of linear series** on $\{C_t\}$ will be a collection of subvector spaces $V_t \subset H^0(\mathcal{L}_t)$ which form a vector bundle in $\pi_*\mathcal{L}$. Note that $\pi_*\mathcal{L}$ is a vector bundle on B (because it's torsion-free and torsion-free on a smooth curve means locally free).

To say that the V_t form a sub-bundle of $\pi_*\mathcal{L}$ is the local trivialty condition we desired.

Now we want to do this in the context in the setting of linear series with *ramification*. So we're going to consider a **family of curves with marked points.** This is going to be a family of curves $\pi : \mathcal{C} \to B$, together with sections $\sigma_1, \dots, \sigma_m : B \to \mathcal{C}$. The claim is:

22.6 Proposition. *Consider a family $\{V_t\}$ of linear series. The adjusted Brill-Noether number $\rho(V_t, \sigma_1(t), \dots, \sigma_m(t))$ is lower semi-continuous in $t \in B$.*

This is just saying that the ramification index is upper semi-continuous, because the order of vanishing only increases when you specialize. That is, $\alpha(V_t, \sigma_k(t))$ is upper semi-continuous in t.

§3 The basic construction

With that said, let's introduce the basic construction. This will introduce one of the basic tools of algebraic geometry: specialization arguments.

As stated earlier, if we managed to exhibit a curve satisfying the Brill-Noether conditions, then we'd be done. But we can't. The strategy is going to be to specialize a general curve to a *singular* curve.

Consider a family $\mathcal{C} \to \Delta$ (as before, a morphism), where $\mathcal{C}$ is a *smooth surface*, and Δ is the disk in the complex plane. The fibers C_t are going to be smooth projective curves for $t \neq 0$, *but* the fiber over zero is going to be the union of smooth curves $X \cup Y$. That is, $C_0 = X \cup Y$ where X, Y are smooth, and $X \cap Y$ is a point where there is a node of C_0. So this is a degeneration of a family of smooth curves to a curve with a node.

We want to try to relate the Brill-Noether theorem on the general family C_t to the Brill-Noether theorem on the *components* X, Y on C_0. That is, we are going to show that if Brill-Noether holds on the two components X, Y, then we'll show that it holds on a general fiber.

We want to introduce sections: marked points (or families thereof).

Remark. Every section of this family $\mathcal{C} \to \Delta$ has to meet the special fiber at a smooth point (by counting intersection numbers).

Let $\sigma_1, \dots, \sigma_m$ be sections of $\mathcal{C} \to \Delta$. Then let's suppose: $\sigma_k(0) \in X$ for $k = 1, \dots, \delta$ but $\sigma_k(0) \in Y$ for $k = \delta + 1, \dots, m$; here, again, X, Y are the two smooth components of C_0. We can do this by the remark.

Here is the basic argument.

22.7 Theorem. *If the adjusted Brill-Noether theorem holds for* $(X, \sigma_1(0), \dots, \sigma_\delta(0), p)$ *and similarly for* $(Y, \sigma_{\delta+1}(0), \dots, \sigma_m(0), p)$*, then that implies that the adjusted Brill-Noether theorem for* $(C_t, \sigma_1(0), \dots, \sigma_m(0))$ *for* $0 < |t| < \epsilon$*.*

We want to deduce the adjusted Brill-Noether statement for the general fiber of this family *from* the statement for the two components. That, in effect, will enable me to break down the given curve into smaller and smaller pieces (X, Y will have smaller genera). We will still have to check this, at the end of the day, for curves of genus zero and one, but there at least we know something about linear series.

§4 Constructing such families

You might ask how we know such families exist. We have to show that such families can be constructed, and with such a choice of sections as well. Here's a brief outline of the idea.

To construct such families, we start with X, Y general and a point $p \in X, p \in Y$ (also general) (which we'll identify to form a single node). In order to construct this family, we'll remove a small neighborhood of the point p; call that U. The picture is that we've picked two points on the Riemann surfaces X, Y, and removed small neighborhoods in X, Y (of p); those neighborhoods are U. We'll consider $X \cup Y - U$; here we've taken out the disk. Consider the product $(X \cup Y - U) \times \Delta$. This looks like a single Riemann surface of genus $g_X + g_Y$ with a band removed, producted with Δ.

Let $\mathfrak{n} = \{(z, w, t) : zw = t, |z|, |w|, |t| < 1\}$. This is a family of conics, parametrized by t, specializing to a union of two lines. For $t \neq 0$, the family is a cylinder. So topologically it looks like a family of cylinders whose waist keeps shrinking until it becomes the union of two disks meeting at a point. Then, we remove a neighborhood of the origin, and we simply get a family of pairs of annuli specializing again to a pair of annuli. We want to glue this $\mathfrak{n}$ construction into $X \cup Y \setminus U$ by identifying annuli in U with $\mathfrak{n}$.

So we're doing a bit of surgery. We take annuli around p, in U, and glue in this family we've just described. The details are specified in the paper on the course website.

Remark. If you like algebra, you can invoke algebraic deformation theory to get this family.

But we want to do more. We want to start with $\sigma_1(0), \ldots, \sigma_\delta(0) \in X \backslash \{p\}, \sigma_{\delta+1}(0), \ldots, \sigma_m(0) \in Y \setminus \{p\}$, and then extend these to sections in *some neighborhood* of zero. So, starting with marked points in X, and marked points in Y, we get a family of smooth curves specializing to a reducible curve, and a family of sections of this family extending the general configuration of points on X, Y that we started with.

§5 Specializing linear series

Let's keep the same notation as before.

Now, suppose we have a family of linear series on $\pi : \mathcal{C}^* \to \Delta^*$. Here $\Delta^* = \Delta \setminus \{0\}$ and $\mathcal{C}^* = \pi^{-1}(\Delta^*)$. What do we get in the limit? This is what the game of specialization is all about. We let these linear series degenerate at zero, and see what we might get.

OK, so a family of linear series on $\pi : \mathcal{C}^* \to \Delta^*$ is the pair:

1. $\mathcal{L}^*$ on $\mathcal{C}^*$ (a line bundle).
2. A vector subbundle $V^* \subset \pi_* \mathcal{L}^*$.

The first question to ask is:

1. Can we extend $\mathcal{L}^*$ to a line bundle $\mathcal{L}$ on $\mathcal{C}$?
2. Is this extension unique?

The answer is yes to the first question. Why? It's easy; just take any global rational section of $\mathcal{L}^*$, take its divisor, (a divisor on $\mathcal{C}^*$), and take the closure to get a divisor on $\mathcal{C}$. So we can always extend a *line* bundle on $\mathcal{C}^*$ to $\mathcal{C}$. Unfortunately, the answer to the second question is no: the extension of the line bundle is not unique.

Yet this non-uniqueness of the extension is a feature, not a bug. For instance, the trivial line bundle on $\mathcal{C}^*$ extends in non-trivial manners to $\mathcal{C}$; this is in effect the question one must construct. So, let's ask:

> How does the trivial line bundle (the structure sheaf) on $\mathcal{C}^*$ extend to a line bundle on $\mathcal{C}$? What line bundles $\mathcal{L}$ on $\mathcal{C}$ are there such that $\mathcal{L}|_{\mathcal{C}^*}$ are trivial?

If $\mathcal{L}|_{\mathcal{C}^*}$ is trivial, then there is a map

$$\mathcal{O}_{\mathcal{C}^*} \to \mathcal{L}_{\mathcal{C}^*}$$

which is everywhere nonzero and holomorphic, and we get a *rational map*

$$\mathcal{O}_{\mathcal{C}} \dashrightarrow \mathcal{L},$$

which is the same thing as a rational section of $\mathcal{L}$. Along the special fiber, there might be zeros and poles, though. If it has a zero, then the zeros of that section are codimension one, so the zeros of the section have to consist of either or both of the components X, Y.

That is, if $\mathcal{L}|_{\mathcal{C}^*} \simeq \mathcal{O}_{\mathcal{C}^*}$ (i.e. $\mathcal{L}$ a line bundle whose restriction to $\mathcal{C}^*$ is trivial), then we can take a rational section σ of $\mathcal{L}$ with poles only along X, Y. Thus $\mathcal{L}$ is the line bundle associated to the divisor $\alpha X + \beta Y$, and precisely it is these that arise in this way.

That is:

22.8 Proposition. *The line bundles $\mathcal{O}_{\mathcal{C}}(\alpha X + \beta Y)$ are precisely the line bundles on $\mathcal{C}$ which are trivial on $\mathcal{C}^*$.*

Well, not all of these are nontrivial, though.

Remark. The line bundle $\mathcal{O}_{\mathcal{C}}(X + Y)$ is trivial (take $1/t$ as a section). However, $\mathcal{O}_{\mathcal{C}}(X)$ is not. To see this, restrict to the curve Y (to get $\mathcal{O}_Y(p)$, which is nontrivial); similarly, $\mathcal{O}_{\mathcal{C}}(mX)$ is nontrivial for $m \neq 0$.

Similarly, one can check that $\mathcal{O}_C(mX)|_X = \mathcal{O}_C(-mY)|_X = \mathcal{O}_X(-mp)$, which is nontrivial.

So the ambiguity in extending a line bundle $\mathcal{L}^*$ over $\mathcal{C}^*$ to $\mathcal{C}$ is precisely the ambiguity of twisting by (some multiple of) X or Y.

The upshot is the following. Given a line bundle $\mathcal{L}^*$ on $\mathcal{C}^*$ and any $\alpha \in \mathbb{Z}$, then there is a *unique* extension $\mathcal{L}$ of $\mathcal{L}^*$ such that $\deg(\mathcal{L}|_X) = \alpha$.

22.9 Definition. We write $\mathcal{L}_\alpha$ for the above construction. This can be thought of as a limit of the $\mathcal{L}_t^*, t \neq 0$.

If we try to describe the geometry of the linear series on the general fibers, we want to say something about them by looking at the linear series on the limiting special fiber. But there are many different possible limits; we have to deal with many of them. At least, two of them. We will focus on the extensions $\mathcal{L}_d, \mathcal{L}_0$: the first has all its degree on X and nothing on Y, while $\mathcal{L}_0$ has all its degree on Y and nothing on X. What do we get when we do these?

Think about it this way. If we have a family of linear series on $C_t, t \neq 0$, i.e. a line bundle $\mathcal{L}^*$ on $\mathcal{C}^*$ and a subbundle $V^* \subset \pi_* \mathcal{L}^*$, we can choose an extension $\mathcal{L}_\alpha$ of $\mathcal{L}^*$ to all of $\mathcal{C}$, then we can likewise extend the subbundle V^* to a subvector bundle $V \subset \pi_* \mathcal{L}$.

Remark. If we have a vector bundle on a disk, and a subbundle on the disk minus a point, then we can just take its closure to get a subbundle over the whole disk of the same rank. You can do this algebraically, too. The point is that any map from a punctured disk to the Grassmannian extends to the full disk.

So, in addition to the vector spaces $V_t, t \neq 0$, we get a subspace $V_0 \subset H^0(\mathcal{L}_\alpha|_{C_0})$. To emphasize the dependence on α, we write $(V_\alpha)_0$. Let's take $\alpha = d$; then $(V_d)_0 \subset H^0(\mathcal{L}_d|_{C_0})$. There is an injection

$$H^0(\mathcal{L}_d|_{C_0}) \to H^0(\mathcal{L}_d|_X)$$

(a section over Y is determined by the value on p). So we get a g^r_d on X.

Similarly, if we take $(V_0)_0$ (the case where the extension $\mathcal{L}_0$ is chosen), we get a g^r_d on Y.

Thus we can get linear series on both X and Y obtained by degeneration. The crux of the proof is to determine what the relation between the linear series on X, Y has to be (they're not arbitrary if they fit into such a family). Next (and last) time, the sequence we'll do is the following. We'll describe the relation between these two g^r_d's on X, Y, and deduce from this a relation between the adjusted Brill-Noether numbers of the linear series on the general fiber (i.e. $\rho(V_t, \sigma_1(t), \dots, \sigma_m(t))$) and the adjusted Brill-Noether numbers of these g^r_d's on X, Y (with respect to the marked points there). That will allow us to deduce the Brill-Noether theorem for the general fiber.

Lecture 23
12/2

Previously, we were in the following setup. As always, C is a smooth projective curve, and $p_1, \dots, p_m$ are points on C. V will be our notation for a g^r_d on C; this means the data of a pair $(\mathcal{L}, V)$ where $\mathcal{L}$ is a line bundle and $V \subset H^0(\mathcal{L})$ is dimension $r+1$.

For any p, we define $V(-ap)$ to be the collection $\{\sigma \in V : \mathrm{ord}_p(\sigma) \geq a\}$. These consist of sections of V vanishing to degree $\geq a$ at p, and we can also view $V(-ap)$ as a subset of $H^0(\mathcal{L}(-ap))$. We will use this abuse of notation.

Given the situation of a curve and a point of points on it (C and $p_1, \dots, p_m$) we defined the *adjusted Brill-Noether number* $\rho(V, p_1, \dots, p_m) = g - (r+1)(g-d+r) - \sum_{k=1}^m \alpha(V, p_k)$. We're looking at the locus of curves and linear series with specified ramification at certain points. We want to show that this adjusted Brill-Noether number is in fact the dimension of the space of such linear series. That is, we will prove:

23.1 Theorem. *If C and the points $p_1, \dots, p_m$ are general, and V an arbitrary linear series on the curve, then $\rho(V, p_1, \dots, p_m) \geq 0$.*

The classical Brill-Noether theorem without ramification is just the special case when $m = 0$. We're going to get into the basic construction introduced last time in just a moment, but what I want to first say is the following.

§1 Two special cases

There are two special cases where we know the answer.

23.2 Example (Genus zero)**.** Take $g = 0$, in which case $C = \mathbb{P}^1$. The adjusted Brill-Noether theorem states that the Brill-Noether number, which is $-(r+1)(-d+r) - \sum \alpha(V, p_k) = (r+1)(d-r) - \sum \alpha(V, p_k)$, should be nonnegative when $p_1, \dots, p_k$ is general.

This may sound pretty sick, but by the Plücker formula, if we consider

$$\sum_{p \in C} \alpha(V, p) = (r+1)(d-r)$$

where the sum is over *all* points in C, then we're done. Clearly the sum $\sum \alpha(V, p_k)$ is at most $\sum_{p \in C} \alpha(V, p)$, but we've seen that the total ramification is exactly $(r+1_(d-r)$.

23.3 Example (Genus one)**.** Let $C = E$ be an elliptic curve, and suppose $m = 1$, so there is only one point. Write $p_1 = p$.

Let $V \subset H^0(\mathcal{L})$ where is $\mathcal{L}$ is a line bundle over degree d. What are the orders of vanishing of sections of V at p? The top order of vanishing $a_r(V, p) \leq d$, because a section can have at most d zeros. Moreover, $a_{r-1}(V, p) \leq d-2$; for if $a_{r-1}(V, p) = d-1$, then $\dim V(-(d-1)p) \geq 2$. That's a contradiction because $\mathcal{L}(-(d-1)p)$ is degree one.

Similarly, $a_{r-2}(V, p) \leq d-3$, and so on. We find:

$$\alpha_r(V, p) \leq d - r, \quad \alpha$$

One concludes that:

$$\alpha(V,p) \leq (r+1)(d-r-1)+1. \tag{7}$$

The conclusion of the adjusted Brill-Noether theorem is that (under generality hypotheses)

$$1-(r+1)(1-d+r)-\alpha(V,p) \geq 0,$$

but that's clear from this inequality (**??**) So the Brill-Noether theorem is sort of trivial for a genus one curve with only one point. (Note that when you work with one point, it is automatically general. Any two points on a genus one curve look alike.)

If you work it out, the adjusted Brill-Noether theorem fails for an elliptic curve and two points if and only if they differ by appropriate torsion.

§2 The inductive step

Now we want a way of proving the adjusted Brill-Noether theorem by some kind of induction on the genus and the number of marked points. The claim is:

> **Inductive step.** ABN (adjusted Brill-Noether) for a given genus g and a given number of marked points m_1+1 and ABN for a given genus h and m_2+1 marked points implies ABN for genus $g+h$ and m_1+m_2 marked points. That is,
>
> $$\mathrm{ABN}_{g,m_1+1} \wedge \mathrm{ABN}_{h,m_2+1} \implies \mathrm{ABN}_{g+h,m_1+m_2}.$$

Once we start with $\mathrm{ABN}_{0,m}$ for any m and $\mathrm{ABN}_{1,1}$, we can use this inductive step to get ABN for any genus and any number of points by induction. For instance, we can get ABN for an elliptic curve with any number of marked points. So, once we've proved the inductive step, the theorem is proved.

§3 The basic construction

Here is the construction. What I want to do is to start with a pair of curves of genera g,h and a collection of marked points on each of them. So, let X,Y be general curves of genus g,h, and let $p,p_1,\dots,p_{m_1} \in X$ are general points, and $p,p_{m_1+1},\dots,p_{m_1+m_2}$ are general points on Y. We use the same letter p for one of the points on X,Y, which is bad notation; but I'm about to fix this by identifying the two points. We let $C_0 = X \sqcup Y/p$: that is, we glue the two curves at p to form a nodal curve with a node at p. So C_0 is a nodal curve containing X,Y as components.

The basic construction, as last time, is to smooth it. That is, we can find a smooth, complex analytic surface $\mathcal{C}$, fibered over a disk Δ via $\pi : \mathcal{C} \to \Delta$, such that the fibers $C_t, t \in \Delta^*$, are smooth curves of genus $g+h$, but the special fiber C_0 is the singular curve that we called C_0 as before. You can carry this out complex analytically by throwing away of a neighborhood of the point p, taking a complement of the constant family, and gluing in a family of the form $xy=t$. You can see this worked out in "Moduli of curves" or in the original paper. If you're more algebraically inclined, you can invoke the deformation theory of nodal curves. Then Δ is replaced by $\mathrm{Spec} R$ for R a discrete valuation ring, and there are only two fibers.

What did we do with all this? Suppose we had a family of linear series on the smooth curves in this family. Suppose $V_t \subset H^0(\mathcal{L}_t)$ (for $t \neq 0$) is a family of g^r_d;s on $C_t, t \neq 0$; this means that we have a line bundle $\mathcal{L}^*$ on the total space $\mathcal{C}^* = \pi^{-1}(\Delta^*)$ whose restrictions to the fibers are the $\mathcal{L}_t$, and V^* is a vector bundle over Δ^* contained in $\pi_*\mathcal{L}^*$. This was the notation we set up last time for families of linear series.

The basic observation we made last time is that there are a lot of different ways of extending this line bundle $\mathcal{L}^*$ to the total space $\mathcal{C}$. For any integer a, there exists a unique extension $\mathcal{L}_a$ of $\mathcal{L}^*$ to all of $\mathcal{C}$, with the following property: the degree of $\mathcal{L}_a|_X$ is a, and on $\mathcal{L}_a|_Y$ is $d - a$.

Remark. Many details are being swept under the rug here. Also, many pictures were drawn on the board, which do not appear in these notes.

So again, let's try to wonder what the "limit" linear series of the V_t will be: this should be a bunch of linear series on X, Y.

23.4 Definition. For any a, we can construct an extension $\mathcal{L}_a$ of $\mathcal{L}^*$ to $\mathcal{C}$ to have degree a on X and degree $d-a$ on Y, as last time. (The way to do this is to start with any line bundle extending $\mathcal{L}_a$, and then twist by X or Y appropriately.)

In this case, we let V_a be the closure or saturation of V^* in $\pi_*\mathcal{L}_a$. (This is because away from 0, we have a well-defined subbundle of $\pi_*\mathcal{L}_a$, and we just extend it.) We let $\mathrm{V}_a = (V_a)_0$ be the fiber over zero, so it sits inside $H^0(\mathcal{L}_a|_{C_0})$.

We can focus on the two extremes $\mathcal{L}_d, \mathcal{L}_0 \in \mathrm{Pic}(\mathcal{C})$ such that on the special fiber, all the degrees are concentrated in one component. These are the "limit" linear series. Let $\mathcal{L} = \mathcal{L}_d|_X, \mu = \mathcal{L}_0|_Y$.

Note that $\mathrm{V}_d \subset H^0(\mathcal{L}_d|_{C_0}) \hookrightarrow H^0(\mathcal{L}_d|_X)$ because the restriction of $\mathcal{L}_d$ to Y is trivial, and a section is determined over Y by its value on p. Let W be the image of V_d in $H^0(\mathcal{L})$, so this is a g^r_d on X. We can do the same thing for $\mathcal{L}_0$; we have that $\mathrm{V}_0 \subset H^0(Y, \mu)$, as before. We call the image U.

So we have linear series W, U on X, Y obtained by this limiting process.

Now we're going to choose sections $\sigma_1, \dots, \sigma_{m_1+m_2}$ of π which extend the p_k, i.e. such that $\sigma_k(0) = p_k$. We can thus think of the family of curves C_t as having marked points $\sigma_k(t)$ on them.

The basic observation is that the ramification sequence in t is upper semicontinuous. If we have a family of linear series, then the ramification sequences at $\sigma_k(t)$, that is $\alpha_i(V_t, \sigma_k(t))$ (for any i), is upper semicontinuous. When you specialize something, the ramification can only go up.

The strategy is now:

1. Assume ABN for $(X, p, p_1, \dots, p_{m_1})$ and $(Y, p, p_{m_1+1}, \dots, p_{m_1+m_2})$.

2. Deduce it for the nearby curves with the nearby collections of marked points.

So, OK, we've seen upper semicontinuity of the ramification sequence.

Remark. We are using the fact that ABN is an open condition. It is enough to deduce ABN for a *single* curve of genus $g + h$ and a single choice of the appropriate number of marked points. (We use the fact that the moduli space of curves with marked points is irreducible.)

§4 A relation between the ramifications of U, W

There is one more thing to say. The key now is to understand the relation between these two linear series on X, Y: that is, U, W.

For each a, we've defined $\mathcal{L}_a$ such that $\deg \mathcal{L}_a|_X = a$. We also know that

$$\mathcal{L}_a = \mathcal{L}_d((-d-a)Y) = \mathcal{L}_0(-aX).$$

If we look at $V_a \subset H^0(\mathcal{L}_a|_{C_0})$, we see that we have a line bundle on this reducible curve, and a vector space of sections of dimension $r+1$. We can restrict this in turn to each component respectively; that is, there is a map

$$H^0(\mathcal{L}_a|_{C_0}) \subset H^0(\mathcal{L}_a|_X) \oplus H^0(\mathcal{L}_a|_Y)$$

and we get the set of pairs of sections which agree on p. So the map is either an isomorphism (if every section on X and on Y vanishes on p) or is such that the image has codimension one. If we look at V_a, then by virtue of this expression, we find that

$$V_a|_X \subset W(-(d-a)p), \quad V_a|_Y \subset U(-ap).$$

The dimension of these subspaces must add up to at least $r+1$, by combining these observations. It follows that:

$$\dim W(-(d-a)p) + \dim U(-ap) \geq r+1, \quad \forall a. \tag{8}$$

Now what are these dimensions? Well, $\dim U(-ap) = r+1-$the number of i such that $a_i(U,p) < a$. So this inequality is saying is that for any a, the number of i such that $a_i(U,p) < a$ plus the number of j such that $a_j(W,p) < d-a$ adds up to $\leq r+1$. If you think about this, it follows that[21]

$$a_i(W,p) + a_{r-1}(U,p) \geq d. \tag{9}$$

What this is saying is that the vanishing sequences of these linear series U, W have to be somehow complementary. Incidentally, a complementary way of saying this is that

$$\alpha_i(W,p) + \alpha_{r-i}(U,p) \geq d-r,$$

which means that these linear series are pretty highly inflected at p. One conclusion is that the *total ramification indices* satisfy

$$\boxed{\alpha(W,p) + \alpha(U,p) \geq (r+1)(d-r),}$$

which is the key thing we're going to use.

Let me try to say this again. What I claimed initially is (9), for which we can give an alternative proof.

Proof. Suppose $a_i(W,p) = b$ and $a_{r-i}(U,p) = a$ with $a+b < d$. Then, before this, there are $a_0, \ldots, a_{r-i-1} < a$ where U has ramification. When we require to vanish to order a, we kill that many sections, so

$$\dim U(-ap) = i+1,$$

[21] I'm not really following this lecture, so these notes are going to be sketchy.

and there are no base points because by hypothesis there is a section vanishing exactly to order a. At the same time,

$$\dim W(-(d-a)p) \leq \dim W(-(b+1)p) \leq r-i$$

by the same logic. But that just directly contradicts the assertion earlier (8), which was a part of a fairly deep story. If you have a pair of limiting linear series in this way, they have to be fairly inflected. ▲

All I want to do from this analysis is to take the inequality:

23.5 Proposition. $\alpha(W,p)+\alpha(U,p) \geq (r+1)(d-r)$.

Now we'll finish this in the next five minutes, though you've heard that before surely.

§5 Finishing

Now we want to compare the adjusted Brill-Noether numbers of W on X, U on Y, and the general member of this family on C_t. Let's write down what we have.

$$\rho(W,p,p_1,\dots,p_{m_1}) = g-(r+1)(g-d+r)-\sum_1^{m_1}\alpha(W,p_k)-\alpha(W,p).$$

Likewise (this is a waste of chalk):

$$\rho(U,p,p_{m_1+1},\dots,p_{m_1+m_2}) = h-(r+1)(h-d+r)-\sum_{m_1+1}^{m_1+m_2}\alpha(U,p_k)-\alpha(U,p).$$

This is just the definition of ρ. What happens when we add these? We get

$$g+h-(r+1)(g+h-d+r)+(r+1)(d-r)-\sum_1^{m_1+m_2}\alpha(W\text{or } U,p_k)-\alpha(W,p)-\alpha(U,p),$$

and by upper semicontinuity and the proposition, this is at most (for $t \simeq 0$, not zero)

$$g+h-(r+1)(g+h-d+r)+(r+1)(d-r)-\sum_1^{m_1+m_2}\alpha(V_t,\sigma_k(t))-(r+1)(d-r).$$

This of course is

$$\rho(V_t,\sigma_1(t),\dots,\sigma_{m_1+m_2}(t)).$$

So we get

$$\rho(W,p,p_1,\dots,p_{m_1})+\rho(U,p,p_{m_1+1},\dots,p_{m_1+m_2}) \leq \rho(V_t,\sigma_1(t),\dots,\sigma_{m_1+m_2}(t))$$

And that's it. The sum of the adjusted Brill-Noether numbers of the two components of the special fiber is at most that of the general fiber.

Remark. One has to make some argument to fit any vector bundle on a nearby fiber to the whole surface.