8320 Spring 2010 Day one: Introduction to Riemann Surfaces

We will describe how Riemann used topology and complex analysis to study algebraic curves defined over the complex numbers. The main tools and results have analogs in arithmetic, which I hope are more easily understood after seeing the original versions. The idea is that an algebraic curve C, say in the projective plane, is the image by a holomorphic map of an abstract compact complex manifold, the Riemann surface X of the curve, where X has an intrinsic complex structure independent of its representation by C in the plane. Although the complex structure is inherited from the plane representation, it can be described in an intrinsic way and may be derived from many different plane representations. We will construct two important functors of an algebraic curve, the Riemann surface X and the Jacobian variety J(X), and natural transformations Abel:X^(d)--->J(X) from the "symmetric powers" X^(d) of X, to J(X).

The Riemann surface X

The fundamental construction is the Riemann surface of a plane curve:
{irreducible plane curves: f(x,y)=0} ---> {compact Riemann surfaces X}

The first step is to compactify the affine curve: f(x,y) =0 in A^2, the affine complex plane, by taking its closure C in the complex projective plane P^2. Then one separates branches at points where C intersects itself, and finally one smooths each of those branches, to obtain a smooth compact oriented surface X. X inherits a complex structure from the coordinate functions of the plane. If f is an irreducible polynomial, X will be connected. Then X will have a topological genus g, and a complex structure, and will be equipped with a holomorphic map *f*:X--->C of degree one. The map *f* will be an isomorphism except over points where the curve C is not smooth, e.g. where C crosses itself or has a pinch.

This analytic version X of the curve C retains algebraic information about C; e.g. the field M(X) of meromorphic functions on X is isomorphic to the field Rat(C) of rational functions on C = the quotient field k(x)[y]/(f), where k = complex number field. It turns out that two curves have isomorphic Riemann surfaces if and only if their fields of rational functions are isomorphic, if and only if the curves are equivalent under maps defined by mutually inverse pairs of rational functions. Since the map X--->C is determined by the meromorphic functions (x,y) on X, which generate the field Rat(C), classifying algebraic curves up to "birational equivalence" becomes the question of classifying these function fields and classifying pairs of generators for each field. Riemann's approach to this algebraic problem will be topological/analytic. We already can deduce that two curves are not birationally equivalent unless their Riemann surfaces have the same genus. This solves the problem that interested the Bernoullis of why most integrals of form dx/sqrt(cubic in x) cannot be "rationalized" by rational substitutions. In fact only curves of genus zero can be so rationalized and y^2 = (cubic in x) usually has genus one.

The symmetric powers X^(d)

To recover C from X, we seek to encode the map *f*:X--->C, equivalently *f*:X--->P^2, by intrinsic geometric data on X. If the polynomial f defining C has degree d, then each line

L in the plane P^2 meets C in d points, counted properly. Thus each line L defines an unordered d tuple L.C of points on C, possibly with repetitions, hence when pulled back via *f*, a d tuple on X called a positive "divisor" D = *f*^(-1)(L) of degree d. We have D = n1p1+...nk pk, where nj are positive integers, n1+...nk = d. Since lines L in the plane move in a linear space dual to the plane, and (if d ≥ 2) each line is spanned by the points where it meets C, we get an injection P^2*--->{unordered d tuples of points of X}, taking L to *f*^(-1)(L).

Let X^d be the d - fold Cartesian product of X, and Sym(d) is the symmetric group of permutations of d objects, and define X^(d) = X^d/Sym(d) = the "symmetric d-fold product" of X. Then the symmetric product X^(d) parametrizes unordered d tuples on X, and inherits a complex structure as well. Thus the map *f*:X--->C yields a holomorphic bijection P^2*--->∏ from the projective plane to a subspace ∏ of X^(d). Hence the map *f* determines a complex subvariety of X^(d) isomorphic to a linear space ∏ ≈ P^2*. Conversely, this "linear system" ∏ of divisors of degree d on X determines the map *f* back again as follows:

Define *f*:X--->∏* = P^2** =P^2, by setting *f*(p) = the line in ∏ consisting of those divisors D that contain p. This determines the point *f*(p) on C in P^2, because a point in the plane is determined by the lines through that point. Thus the problem of finding representations of X in P^2 becomes one of determining when the product X^(d) contains a holomorphic copy of P^2, or copies of P^n for models of X in other projective spaces.

The Jacobian variety J(X) and the Abel map X^(d)--->J(X)
To study this problem, Riemann introduced a second functor, the "Jacobian" variety J(X) = k^g/lattice, where k^g = complex g -dimensional space. J(X) is a compact g dimensional complex group, and there is a natural holomorphic map Abel:X^(d)--->J(X), defined by integrating a basis of the holomorphic differential forms on X over paths in X. Abel collapses each linear system ∏ ≈ P^n* to a point by the maximum principle, since the coordinate functions of k^g have a local maximum on the compact simply connected variety ∏. Conversely, each fiber of the Abel map is a linear system in X^(d).

Existence of linear systems ∏ on X: the Riemann - Roch theorem.
By dimension theory of holomorphic maps, every fiber of the Abel map X^(d)--->J(X) has dimension ≥ d-g. Hence every positive divisor D of degree d on X is contained in a maximal projective linear system |D| , where dim|D| ≥ d-g. This is called Riemann's inequality, or the "weak" Riemann Roch theorem.

The Roch part computes dim|D| more precisely by analyzing the relation between D and the divisor of a differential form. Note if D is the divisor cut by one line in the plane of C, and E is cut by another line, then E belongs to |D|, and the difference E-D is the divisor of the meromorphic function defined by the quotient of the linear equations for the two lines. If D is a not necessarily positive divisor, we define |D| to consist of those positive divisors E such that E-D is the divisor of a meromorphic function on X. If there are no such positive divisors, |D| is empty and has "dimension" equal to -1. If K is the divisor of zeroes of a holomorphic differential form on X, then the full Riemann Roch theorem says: dim|D| = d-g +1+dim|K-D|, where the right side = d-g when d > deg(K).

Day two: Definition of the Riemann surface of a non singular affine plane curve
We start at the beginning with curves and Riemann surfaces which are not necessarily compact. We want to define holomorphic functions on a plane curve C:{ f(x,y) = 0}. Of course we want the coordinate functions x and y to be holomorphic on C, hence also all holomorphic combinations of them. This amounts to saying a function on X is holomorphic if it is the restriction of a holomorphic function of x and y, i.e. if it extends to a holomorphic function on the plane k^2. But if we want to be able to expand a function on C in a power series about a point of C, we need a single local coordinate near each point, instead of the two coordinates x and y. This we do next.

First example: C is the graph of a holomorphic function y = g(x), i.e. a plane curve of form 0 = f(x,y) = y – g(x) . As always, the graph of a function is isomorphic to its domain, which allows us to use the single coordinate x on the graph. I.e. since x and y are both holomorphic on the plane, their restrictions to the graph C should be holomorphic. In particular projection (x,y)---> x should be holomorphic on C. Conversely, since g is holomorphic by hypothesis, the inverse map x--->(x,g(x)) is holomorphic. Thus if we assume that both coordinates x and y are holomorphic on C, it follows that the graph C of a function g(x) is holomorphically isomorphic to the x axis. Hence it follows that a function on C will be holomorphic if and only if it is holomorphic as a function of x. Thus we will get a definition of holomorphic function on X in terms of only one coordinate, by letting x be the coordinate, and saying that a function h on C is holomorphic if and only if it is holomorphic as a function of x, via the composition map x--->(x,g(x))--->h(x,g(x)).

Thus we know how to put holomorphic coordinates on the graph of a holomorphic function. But holomorphicity is a local notion, so we should be able to do the same on a curve which is only locally a graph. Detecting such curves is exactly what the implicit function theorem is designed for. Draw a picture of a curve in the plane and project it onto the x axis. Look at the points of the curve where the projection is not a local isomorphism I.e. recall a curve is a graph if it passes the vertical line test, i.e. on some (complex) neighborhood, all vertical lines meet the curve exactly once. Look at points where this fails. These are points where the tangent line is vertical as well as all points where two branches of the curve cross. The implicit function theorem says these points on f(x,y) = 0 can be detected by the vanishing of the partial derivatives of the function f. I.e. f = 0 is a level set of f, and gradf is perpendicular to its level sets. Thus at points where two branches of the level set cross each other, at least if they cross transversely, grad f must be perpendicular to two independent tangent vectors, one to each branch. Thus gradf must be identically zero at such crossing points, i.e. both partials $\partial f/\partial x$ and $\partial f/\partial y$ are zero. At a point where the curve f = 0 is smooth, i.e. with only one tangent line but that tangent line is vertical, gradf must be horizontal, i.e. $\partial f/\partial y$ must be zero. At all points other than these two types, i.e. at all points where $\partial f/\partial y \neq 0$, projection to the x axis is a local isomorphism. So we need to exclude precisely those points where $\partial f/\partial y = 0$, to have x be a local coordinate.

Implicit function theorem: If f(x,y) is a polynomial in two complex variables x,y, and if p = (x0,y0) is a point where f(x0,y0) = 0 and also $\partial f/\partial y \neq 0$, then on some neighborhood of p, the level set f = 0 is the graph of a holomorphic function y = g(x).

As a corollary, if f(x,y) is a polynomial and C is the level set: {f=0}, then x is a local coordinate for C at every point p where $\partial f/\partial y \neq 0$. Similarly, y is a local coordinate for C near every point q where $\partial f/\partial x \neq 0$. And if p is a point where both partials $\partial f/x$ and $\partial f/\partial y$ are $\neq 0$, then by the implicit function theorem, y and x are holomorphic functions of each other, so they give equivalent holomorphic local coordinates. I.e. at such points, a function on C, is holomorphic as a function of x if and only if it is holomorphic as a function of y. The points where both partials are zero are not "Riemann surface" points. They can be modified to become such as we discuss later, but now we want to formalize the properties of these good points, in the definition of a Riemann surface.

Definition: A Riemann surface X is a Hausdorff topological space equipped with a cover of X by open sets Uj, and for all j, homeomorphisms fj:Uj--->Vj where Vj is an open subset of the complex numbers, such that for all i,j, the composition fjfi^-1 is holomorphic where defined, i.e. on the open set fi(UimeetUj) in the complex numbers.

Corollary of implicit function theorem: If f(x,y) is a complex polynomial, then the open subset of its level set f = 0 consisting of those points in k^2, k = complex field, where some partial, $\partial f/\partial x$ or $\partial f/\partial y$, is non zero, is a Riemann surface with coordinate functions given by restrictions of the functions x and y on appropriate sets.

A function g:X--->k on a Riemann surface is holomorphic if all compositions gf^-1 of g with coordinate functions are holomorphic where defined. The function g is holomorphic near p if this holds for one coordinate function defined at p. The derivative of such a g is zero at p if the derivative of gf^-1 is zero for one and hence for all charts f near p.

Definition: A level set f(x,y) = 0 in k^2 of a polynomial f(x,y) is called an affine plane curve. A point of f = 0 where both partials $\partial f/\partial x$ and $\partial f/\partial y$ are zero is called a singular point, and other points are called non singular points. A curve with no singular points is called a non singular curve. Thus a non singular plane curve is a Riemann surface.

Remark: The functions x and y are everywhere holomorphic functions on the Riemann surface of a non singular plane curve, but neither one is necessarily everywhere a local coordinate. The function x is a coordinate near all points where its derivative is non zero with respect to some local coordinate.

Example: If g(x) is a polynomial in one complex variable x with no repeated roots, then f(x,y) = y^n – g(x) = 0 defines a non singular plane curve in k^2 for $n \geq 1$, since if $\partial f/\partial y$ = ny = 0 at (x,y), then y = 0, so g(x) = 0, and $\partial f/\partial x$ = g'(x) is non zero by hypothesis on g.

Remark: Every Riemann surface is a smooth real 2 dimensional surface, and since multiplication by a complex number preserves orientation of the complex number line, the local coordinates induce on the Riemann surface an orientation consistent with the usual one on the complex numbers, via the local coordinate neighborhoods. Thus every Riemann surface is an orientable 2 dimensional real surface with a preferred orientation consistent with that given by the ordered real basis $\{1,i\}$ for the complex line k.

In fact every non singular affine plane curve $f(x,y) = 0$ is topologically a compact oriented 2 dimensional real surface with a finite set of points removed. Further, the non singular polynomial f is irreducible if and only if the compact Riemann surface is connected. We want to determine the genus of this compact surface, which will also have a complex structure of a Riemann surface.

An abstract example (not a plane curve): a "complex torus"
If a "lattice" is a Z submodule of the complex numbers k generated by two complex numbers which are linearly independent over the reals, then the quotient group T = k/lattice, with its natural quotient topology, is called a "torus".
Ex: There is a unique structure of Riemann surface on such a torus T, so that the natural quotient map k—> T is a local holomorphic isomorphism. (A function on an open subset of T is holomorphic iff its composition is holomorphic on an open subset of k.)
Ex: Every such torus is holomorphically isomorphic to one whose lattice has generators $\{1,t\}$ where $\mathrm{Im}(t) > 0$.
Ex: All such tori are diffeomorphic.
***Ex:** Not all such tori are holomorphically isomorphic.

Day three: Visualizing examples of Riemann surfaces of affine plane curves
Since every compact orientable surface has a genus, it is instructive to try to compute the genus in some examples. Even without compactifying our affine curves, the genus is determined and may be visible, since the affine curve only differs from the compact Riemann surface by the removal of a finite number of points.

1) C:{y^2 = x}, where x and y are complex coordinates on k^2, k = complex numbers. If we project onto the y axis, then over each y there is a unique x, namely x = y^2. So this projection is a bijection and this curve is isomorphic to the y axis, with inverse the map from the y axis to the curve taking y--->(y^2,y).

It is useful also to try to see the structure of the surface by projecting on the x axis. We have a map C--->k which is 2:1 over every x except x=0, where there is only the one point y = 0 over 0. If we think of the x axis as a large disc or rectangle, it is not easy to visualize how a smooth surface can lie directly over a rectangle, with two points over every point except the center, over which there is only one point. We know however since this is a holomorphic map, that locally it has the same structure as some power map z^k, which must be squaring, i.e. k must = 2.

Indeed we know the curve C is isomorphic to the y axis, and the composition of the isomorphism y--->(y^2,y) with projection on x, is just the squaring map y--->y^2. If we recall the behavior of the squaring map, it doubles angles about 0, so this map sends the top half of the y axis onto the whole x axis, with the horizontal line dividing the top half of the y axis from the bottom, going doubly onto the positive real line in the complex x axis. Similarly, the bottom half of the complex y axis also maps onto the whole complex x axis with the real line of the complex y axis mapping again doubly onto the positive real x axis. Thus we can visualize the way the curve C lies over the complex x axis as follows.

If we slit the complex x axis along the positive real line, there are two copies of this slit rectangle lying over it on C. These two copies are isomorphic to the upper and lower halves of the complex y axis. Now the upper and lower halves of the complex y axis are joined along their common border, and crossing the real line takes us from one half to the other half. So these two slit rectangles up on C are joined cross ways along their two slits, so that crossing the slit takes one from the upper copy to the lower one and vice versa. Thus we can visualize constructing the curve C from its image the complex x axis, without knowing in advance that the curve C was isomorphic to the complex y axis. I.e. we slit the x axis along its positive real line, take two copies of this slit rectangle, and glue the edges of their slits crossways. This is impossible in 3 space while keeping each point over its image without passing the glued slits through each other of course.

Thus if we want to look at the curve C as a rectangle, we have to visualize the branched 2:1 map defined on that rectangle as wrapping around the center and doubling angles. If we want to have the branched 2:1 map to be just vertical projection from the curve C down onto the x axis, then we have more work to do to visualize the way the curve C lies over the x axis. Viewed in 3 space it would then intersect itelf along the preimage of the positive x axis. Of course C actually lies in k^2 or real 4 space where this self intersection does not occur. So either we have a simple view of the curve C and a somewhat complicated view of the map from C (squaring), or we have a simple view of the map from C as just vertical projection, but then we have a harder task to visualize how the curve lies over its image. Anyway, since the affine curve C is isomorphic to the complex numbers, i.e. a copy of the real plane, the only possible compact surface that can be obtained by adding in a finite set of points is the sphere, so the genus is zero.

2) Let C: {y^2 = x^2 -1}. Now we do not know at once what the curve C looks like by using the first method above, since neither projection onto x, nor projection onto y is an isomorphism. So we adopt the second method above, choosing as map the vertical projection from C onto the x axis and trying to visualize how the curve C is constructed from gluing copies of the image x axis. Here again the map C--->k is usually 2:1 but there is only one point over x = -1 and one over x = 1. Moreover we know that locally near each of these points, the map looks like squaring, by the local structure theorem for holomorphic functions. So if we cut the complex x line by a vertical line separating the two points 0 and 1, i.e. if we cut the complex x line into two rectangles, then by example 1, the preimage of each rectangle is again a rectangle up in C. Further since every point of the line we cut along in the x plane has two disjoint preimages in C, we are joining our two rectangles in C by gluing two disjoint intervals in each rectangle.

Thus we are choosing two disjoint edges of each rectangle and gluing them together, each edge from one rectangle to one edge of the other rectangle. This gives either a cylinder or a twisted cylinder (Mobius strip), but since C is orientable it must be a cylinder. Now the only compact surface obtained by adding a finite number of points to a cylinder is a sphere, so this curve again has genus zero. Note that last time we only added one point to the complex line k to get a sphere, and this time we add two points to a cylinder, which is homeomorphic to k-{0}, to get a sphere. So the compact Riemann surfaces are the same, but the affine ones we started from were different.

To see this construction using a different cut in the x plane, make a slit along the interval joining the two branch points -1 and 1. A closed loop in the plane which goes around both branch points deforms to the union of two closed circles, one going around 1 and the other going around -1, with the circles having one common point at 0. Up on the curve C, this means the curve lying over these loops has passed from one preimage of 0 to the other and back again, so after going around both branch points the preimage of 0 is the same at both ends of the full closed curve. This means every closed loop in the complex x line which goes around both branch points once has as preimage a curve that is closed on C. Thus every closed loop in the x line that misses the interval between the two branch points must have a closed preimage on C. Hence if we take two disjoint copies of the x line, both cut along the interval from -1 to 1, then C is obtained from these two copies by cross gluing them along opposite edges of the cuts. Since the preimage on C of the slit betwen -1 and 1 is a circle, we thus form C by gluing in a circle to join up two slit planes, obtaining a cylinder. The gluing must be done across the slits to retain the orientation of C.

In fact the first method above also works if we first change coordinates, setting s = x+y and t = x-y. Since C has equation y^2 = x^2 - 1, or equivalently 1 = x^2-y^2, this gives as equation just 1 = st, or equivalently t = 1/s, for s≠0. Thus projection of this curve onto the s - axis, is an isomorphism between C and the set of non zero complex numbers, which is again seen to be topologically a cylinder. In the original coordinates this map takes (x,y) to x+y = s. This is an isomorphism from the curve to the non zero complex s line. I.e. 0 ≠ s = x+y implies that on C, 1/s = t = x-y, so we can recover x as [s+(1/s)]/2, and y from [s-(1/s)]/2. I.e. s--->([s + 1/s]/2, [s – (1/s)]/2) = (x,y) is inverse to the map (x,y)--->x+y = s, between the curve C and the set {s≠0} on the s line.

3) C: {y^2 = x(x^2-1)}. Here there are three branch points, so we can view C as obtained by gluing three rectangles. We get C topologically by gluing two opposite sides of the center rectangle to two opposite side of the left rectangle, and then gluing the remaining two opposing sides of the center rectangle to two opposite sides of the right rectangle. This gives a cylinder with a strip glued from part of the border of one hole to part of the border of the other hole, like two links in a paper chain. I claim the compact surface has genus one, after adding one additional point. (What is the boundary curve?)

4) C: {y^2 = (x^2 - 1)(x^2 - 4)}. Show this has genus one, after adding two points.

5) C:{y^2 = x(x^2 - 1)(x^2 - 4)}. Show the genus is two, after adding one point.

6) C:{y^2 = (x^2 - 1)(x^2 - 4)(x^2 - 9)}. ????

7) C: $\{y^2 = (x-a_1)\ldots(x-a_n)\}$, all a_j distinct. ????

8) C: $\{ y^3 = 1 - x^2\}$. This one can be done by projecting on the y axis as above, since this is a 2:1 map branched over the three cube roots of 1, by using the equation in the form $x^2 = 1-y^3$. But if we project onto x, we have a triple cover of the x axis, branched at the 2 square roots of 1, so by separating again those 2 branch points by the imaginary line in the x plane, we have two disjoint rectangles, each covered by a triple cover with on branch point. Since that map is equivalent to $y \longrightarrow y^3$, its domain up in C is a disc, or rectangle, or let's say a hexagon. Thus we form C by gluing two disjoint hexagons along the preimage curves of the imaginary line in the x plane. These preimage curve in each hexagon, are three disjoint edges of the hexagon, but when we join these images from one hexagon to the other, the result is a single closed curve in C, because of the way the edges are joined cyclically over each rectangle. If one chooses three disjoint edges in each hexagon and glues in the simplest way, one obtains a sphere with three disjoint discs removed, not a single connected boundary curve. This puzzled me for some time, as this would give genus zero, the wrong answer. But it is possible to visualize how to join the edges of the hexagons to get a connected boundary curve as follows. Take two hexagons and label their edges as if looking at a clock face, noon, 2 o'clock, 4 o'clock, 6 o'clock, 8 o'clock, and 10 o'clock. We want to identify three of these pairs of edges, the ones at 2 o'clock, 6 o'clock, and 10 o'clock on the first hexagon, with those at 4 o'clock, 8 o'clock, and noon on the second hexagon. So first identify 6 o'clock on the first hexagon with noon on the second, and then identify 2 o'clock on the first with 8 o'clock on the second, and finally identify 2 o'clock on the first with 4 o'clock on the second. This system of cross identifications leaves a simple connected boundary curve and is homeomorphic to a genus one surface with one point (or disc) removed.

9) C: $\{ y^3 = 1 - x^3\}$. ???? I have not visualized this one or the next one, but this one should be obtainable by identifying edges on three dodecagons, since for the center hexagon we need two pairs of disjoint edges and another unidentified edge in between each of those 6 edges.

10) C:$\{ y^4 = 1 - x^4\}$. ???? Maybe one should begin on $\{ y^4 = 1 - x^2\}$ where we know the answer, as we did in problem 8.

Obviously these are getting harder to visualize, so we need some more powerful tools to even compute the topological genus of a plane curve. Soon we will introduce the method of "degeneration", or variation of curves in families, to help us make more computations.

Preview: We will define the compact Riemann surface associated to a not necessarily non singular affine plane curve, in two steps. First we will compactify the affine plane curve to obtain a projective plane curve. Then we will desingularize the projective curve by removing its singular points and replacing each singular point by a finite number of new points which will become non singular on the Riemann surface. I.e. at a non singular point, projection on a suitably chosen axis is a local isomorphism from a neighborhood of the non singular point to a disc. It can be shown that projection of a neighborhood of a singular point onto a suitably chosen line, gives a finite covering space from a punctured neighborhood of the singular point to a punctured disc.

Such a finite covering space is always a finite disjoint union of punctured discs, on each of which the covering map is isomorphic to z^k for some k. On the curve, these punctured discs all share the same center point, the singular point of the curve. Then the Riemann surface is obtained by giving a different center to each of these punctured discs, so that we get a disjoint union of discs. The number of disjoint discs in a neighborhood of a singular point is called the number of local analytic branches at that singularity. Thus a singular point is replaced by as many non singular points as there are branches at the singularity. However, even when there is only one branch at a singularity, and the Riemann surface has only one point corresponding to the singularity, the map from the disc neighborhood of that non singular point on the Riemann surface, to a neighborhood of the singular point of the curve, is never a local isomorphism even if it may be a local homeomorphism. I.e. at a singular point with one branch, although there is a neighborhood of the singularity which is homeomorphic to a disc, that disc is pinched, and there is no projection to any axis which is a local coordinate. There is always a unique holomorphic map from the Riemann surface of the projective curve to the projective curve, but it is either not one to one over a singularity, or has derivative zero at some preimage of the singularity, as a holomorphic map to the plane. E.g. k--->k^2 defined by (x,y) = (t^2,t^3) mapping the t line onto the curve y^2 = x^3, is a homeomorphism but has derivative zero at t = 0. Then k is the (non compact) Riemann surface of the affine plane curve y^2 = x^3 with a singularity at (0,0).