Harvard FRG

Gabriel C. Drummond-Cole

August 25, 2009

1 D. Nadler I

I'll be talking at you a lot. Let me give you an overview of my plan. The breakdown of the three talks is roughly the following. Today I'd like to talk about perspectives on D-modules. Next time I'd like to talk about D-modules on flag varieties, and in the third one maybe I'll talk about geometric Langlands. So the second talk will have to do with three dimensional TFT, and the third with four dimensional. This talk is for no audience member left behind.

I'm not going to talk much about my work, but much of this is joint with David Ben-Zvi, and I've been very influenced by him.

Let's start with the basics. Let X be a smooth scheme over $\mathbb{C}$. The ring of differentials D_X is the enveloping algebra $U_{\mathscr{O}_X}\tau_X$ of vector fields, which is

$$\bigoplus_n \tau_X^{\otimes n}/vf - fv = v(f), vw - wv = [v, w]$$

Just to be specific, if $X = \mathbb{A}^n$, $D_{\mathbb{A}^n} = \mathbb{C}\langle x, \partial\rangle/\partial x - x\partial = 1$ Our main object of study will be D-modules, modules over this algebra.

A D module is a quasicoherent sheaf on X with a compatible action of D_X, where functions act as they would in the sheaf. If X is *Spec* R, then an quasicoherent sheaf is an R-module. You could have imagined, well, more complicated things, on any open set you take an R-module, and you localize it.

[Dennis wants more definitions]

A module is a module over coordinate functions, so you can forget quasicoherent sheaves to just think about this a little. It's important to know that locally this is something that looks like a module.

Where do these come from? Let me give you several perspectives on moduls. How to think about them? The first way I'll give you is analytic. Any time you have a linear PDE you have a D-module. Let me give you an example. Say we're on $\mathbb{A}^1$. One favorite is just $\partial f = 0$. This, you can think of as $D/D \cdot \partial$.

Once one has this, why is this a reasonable idea? If I have $P(f) = 0$ I get $D/D \cdot P$, then $Hom_D(M_P \to \mathscr{F})$ with $\mathscr{F}$ a function space like analytic functions, then these homomorphisms are solutions to P in $\mathscr{F}$.

So to see this, for a map $M_P \to \mathscr{F}$, I need to say where 1 goes, it goes to f. Then $P \mapsto P(f)$, but P is 0 in M_P, so $P(f) = 0$.

My more general example is $M = D/D(x)$, wel, $D/D \cdot \partial = \mathscr{O}_X$. Now $D/D(x) = \Delta_[$, the delta functions and derivatives at 0. This is the analytic starting point, trying to algebraicize linear PDEs.

Now some geometric perspective, recall that D_X is naturally filtered by the order of differential equation $\mathscr{O}_X \subset D_X^{\leq 1} \subset D_X^{\leq 2}$, and this filtration has an associated graded, and $gr\ D_X = \mathscr{O}_{T^*X}$. We take a commutative algebra, functions on T^*X, and deforming it to a noncommutative one. To think of a D-module, we can think of it as living on a noncommutative version of T^*X. So if we're bold enough, we can think, roughly, a subvariety of T^*X.

I'd like to make a little more precise dow to think of these. Let's assume for the moment, say a D-module is coherent if it's finitely generated. In other words, it's reasonable from some perspective. So given a coherent D-module, we can always choose a good filtration on M. Let me tell you what a *good* filtration is. It's a filtration by coherent submodules compatible with the filtration of D. So for example, let M^0 be generators over DX, and $M^{\leq i} = D^{\leq i} M^0$. It's extra structure, not canonical, but easy to find. We can then consider $gr\ M$ which is an $\mathscr{O}_{T^*X} = gr\ D_X$-module. This is not canonical, but what is canonical is the set theoretic support $gr\ M \subset T^*X$. This is called the singular support. Sheaves have support, but on the cotangent bundle, well, there are interesting codirections. I'll calculate examples in a moment.

For the time being, this is the best we can do for M, draw its singular support.

Exercise 1 *The singular support $SiS(M)$ is conical and coisotropic. The cotangent bundle is naturally Poisson. Locally, we used the symplectic pairing. That's the hint.*

So I should be drawing only conical pictures, like this: [Picture].

Dennis says, if you take $Y \subset X$ and look at distributions Δ_Y along Y, and $SiS(\Delta_Y) = T^*_Y(X)$.

So let's classify some D-modules based on pictures of their singular support. Let's classify D-modules on $\mathbb{A}^1$ with singular support in the union of the the 0 section and the cotangent bundle at 0, that is $T^*_{\mathbb{A}^1}\mathbb{A}^1 \cup T^*_0\mathbb{A}^1$.

Let's start with my favorite PDE, $D/D(\partial - 1)$. This is a test of how effective my lecture was, this corresponds to the exponential. I'd like to draw this diagonally but this is not conical. The singular support is just the zero sections. If your friend who lives at infinity calls you, he'll say this seems weird. The guy living at ∞ thinks it's quite frightening. So let's assume our D-modules extend to $\mathbb{P}^1 \supset \mathbb{A}^1$. That will rule out my favorite PDE. So we'll try to classify these on $\mathbb{P}^1$, with support in the zero section and the cotangents at 0.

So let's get to work, audience participation. Does anyone have a D-module living in this category? Let's take $\mathscr{O}_{\mathbb{P}^1}$, looking for things with singular support just in the zero section. That's perfectly good. Also Δ_0 has singular support $T_0^*\mathbb{P}^1$. So in coordinates, these look like $\mathbb{C}[x], \mathbb{C}[\delta]$. These are all you're going to get by fixing one or the other. Now let's draw the cross, try to draw something whose support is the union of these two, looking at the equation $x\xi = 0$. So $M_* = D/D(\partial \cdot x)$ or $M_! = D/D(x\partial)$. These are not the same, and both have singular support this cross.

Here's an exercise, a short exact sequence:

$$\mathscr{O}_{\mathbb{P}^1} \to M_* \to \Delta_0$$

$$\Delta_0 \to M_! \to \mathscr{O}_{\mathbb{P}^1}$$

Can anyone think of any others?

[What about D/Dx^2. You can find an isomorphism with a direct sum of these.

There is one more indecomposable. It can't fit into too interesting a family. Let me call these three and four, so that number five can be $\mathscr{T} = D/D(x\partial x)$. This one has singular support which can be upgraded to more than just a subset, everything, but with the vertical component doubled.

If you understand this exercise, you understand representations of sl_2.

So that was supposed to be, some general nonsense, some categorical background. That was a sort of intro to what I mean by a D-module. I will give other pictures later. I'm going to need to pass to derived categories of D-modules. I'm not sure if I should, but I'm going to apologize for this. Now $D(X)$ will be the derived category of D-modules. Dennis already gave the definition. These will be complexes of D-modules, with quasiisomorphisms inverted somehow. You can add boundedness, finiteness, all these games that one plays. For the cogniscenti, I want to work with an ∞-enhancement. I don't know if this level of subtlety will come up. Why do we work with derived categories? Because the functors we use are too silly until we derive. So for $f : X \to Y$ I'll define $f^\dagger$, the $\mathscr{O}$-module pullback. There's nothing fancy about it. This produces $\mathcal{D}$-modules from $\mathcal{D}$-modules, pull back $\mathscr{O}_Y$ and there will be something completely canonical. The pushforward f_* will be the coflat sections along the fibers. Normally you tensor rather than Hom so it's coflat. There are tons of books, but let me write that if Y is a point, then $f_*(M) = \underbrace{f_\dagger}_{\mathscr{O}-\text{module pushforward}}(\omega_X \otimes_{D_X} M)$.

Let me make a comment for experts. This will be a version of de Rham cohomology. If M is $\mathscr{O}$ it will be de Rham cohomology. If you push forward from an affine to a point, you don't lose information, but here you lose information, only get the de Rham cohomology. An unhappy representation theorist here is unhappy with this fact.

Let me finish with a definition and an exercise, we can go from schemes to Artin stacks. So if I give you a stack X I can think of it as being resolved by a simplicial scheme $X_0, X_1, \ldots$. Then I let $D(X) = \lim D(X_\cdot)$, so compatible collections of D-modules with respect to $f^\dagger$.

Exercise 2 *Let X be a point over G, so BG, resolved with a point, G, $G \times G$, and so on. Then $D(BG) = C_{-*}(G)$-modules. This doesn't recover.*

Next time we'll get to flag varieties.

2 Perspectives on Moduli Space

[For travel, talk to Robbie Miller in room 325. If you want your laptop on the internet, talk to me after. For suggestions for dinner, ask after the talk. Tomorrow morning the breakfast will be on this floor because it's the welcoming day for the first year graduate students. Tomorrow evening we have our own food for dinner. I'm very pleased to have Curt McMullen.]

I didn't know that I was giving this talk at four until 12:30 today. [That's your way of thanking the organizers?]

I have the fortune to be one of Dennis' early students. I'm going to give a talk summarizing what I learned in the hope that it is somehow relevant. This will be a survey of classical basic ideas regarding the moduli space of Remann surfaces. I'll start with a topological surface of genus $g \geq 2$. The algebraic topologists attach to this $H^1(\Sigma, \mathbb{Z}) = \mathbb{Z}^{2g}$ with the symplectic intersection pairing. But there's something that gets you more into the topology, the space of isotopy classes of simple closed loops on the surface, $\mathcal{S}$. That means that the curve can be represented without any crossings. It could be trivial in homology, but it need not be isotopic to the identity. On this space, you can put curves in minimal position; then the number of points in that intersection gives you a number $\mathcal{S} \times \mathcal{S} \to \mathbb{N}$, and this is without an orientation. So $\alpha \cdot \beta = 2$, even though α is homologically trivial.

Then there's the fundamental group $\pi_1 \Sigma$, which is a hyperbolic group Γ which has one relation. Scale the word metric to give it finite diameter. If you do it to this surface group you get the circle. This is hyperbolic because the relation involves many letters. This is part of small cancellation groups (or maybe the other way around?)

Now, these isotopy classes, in principle, can correspond to conjugacy classes in $\pi_1 \Sigma$, but figuring out which ones can be represented is hard. One thing that you can do is the following, you can make an embedding of the space of all curves into the space of functions on curves, $\mathbb{R}_+^{\mathcal{S}}$. So now we could take a collection of $a_i \alpha_i$, and just as we required that the α were simple, we can require that the components of the multicurve are disjoint (call these finite laminations), and then we can send this to the sum of the intersection numbers, a linear extension, which maps {finite laminations} $\to R_+^{\mathcal{S}}$, and then we can take the closure of the image, which is the measured laminations on Σ, $\mathcal{ML}$.

Now we can take a surface of genus two, and cut it in half, and then each half looks like a torus with a point removed. So you get a curve $\gamma_{p/q}$ that lives on one half. Rescale these by multiplying by $\frac{1}{q}$, let these go to ∞, and you'll get a lamination. So a very long simple closed curve with a small coefficient in front of it. Now $\mathcal{ML}$ is homeomorphic to $\mathbb{R}^{6g-6}$. If you projectivize you get a sphere of dimension $6g - 7$.

The next level is the automorphisms of these surfaces, so let me introduce the mapping class group Mod_g, the group of homeomorphisms that preserve orientation up to isotopy. The reason these objects I've introduced are interesting is that the mapping class group acts on them. It acts on homology, $Sp_{2g}\mathbb{Z}$. It acts on simple closed curves, the measured laminations, and then on S^{6g-7}. This is like the fundamental group acting on the circle at ∞.

That's the topology, topological perspective on the mapping class group.

The problem with the mapping class group is that it's populated by large Abelian subgroups. What about geometry. The next step is to change your perspective to X a hyperbolic Riemann surface, a quotient of $\mathscr{H}$ by a surface group Γ. This comes with a hyperbolic metric, a metric of constant negative curvature. Everything has been tautological. Let's consider the space T_g of all possible hyperbolic structures. So we choose not just X but a fixed homeomorphism $\Sigma \to X$, modding out by isotopy of the map φ and isometry in X. This is the Teichmueller space of genus g.

Why would we do this? Now the simple closed curves have canonical representatives. You have unique geodesics, and it turns out they realize the minimal intersection number. So now we have $\mathcal{S}$ as the space of simple closed geodesics. The topological boundary is now the boundary of hyperbolic space, and now the action on the boundary is by Mobius transformations.

Now you can take the unit tangent bundle, and once you have a vector, you can follow it on the geodesic flow. It looks like all the structure is very geometric, but it's not. You can construct this from the topological data, even with the leaves of the foliation by the geodesic flow. You see this because the tangent bundle is $S^1 \times S^1 \times S^1$ minus some diagonal with an action of $\pi_1\Sigma$. Two points at ∞ determine a geodesic, and then the third point on the boundary gives the tangent to it. So we get a tangent direction (along the line) and the point of intersection.

It's a kind of spectacular thing that the topology of the geodesic flow is independent of the choice of metric.

A little nuance here so that you don't find this too hard to believe. Just knowing the fundamental group you can build this. You can see the simple closed curves. What's not canonical is the projection of this bundle to X. There's a very interesting invariant. If you have three geodesics that intersect in a triangle, it could look one way or the other.

Okay, so, what is this Teichmueller space, now? It's the set of possible shapes for this surface of genus g. It's actually a nice object (this goes back to Fenchel-Nielsen). So you can choose a decomposition into pairs of pants by $3g-3$ simple closed curves. Up to homeomorphism, a pair of pants decomposition is the same thing as a trivalent graph with 2 loops. There are only finitely many pair of pants decompositions, then.

There are three numbers we can attach, then, ℓ_i, the length, but if you specify three legs arbitrarily, then there exists a unique surface that has those three lengths. I want to dwell on this. It's true for the same reason that you can build a spherical triangle when you pick three angles whose sum is greater than $180°$. So α, β, and γ determine a 3×3 matrix, which has ones on the diagonal, using cosines. Then this information is just saying that the matrix is positive definite.

Taking this hexagon [Picture], using geodesics at three proscribed distances, you can attach across this and get a pair of pants. There are many intuitive reasons this should exist. I wanted to point out one of them.

Now we just have to glue them together. There are points on each boundary component closest to every other. These might or might not match up. Specifying lengths and twists determines the surface. So $T_g \cong \mathbb{R}_+^{2g-3} \times \mathbb{R}^{3g-3}$.

There's a natural action of $\mathbb{Z}^{3g-3}$ under Dehn twists, and using this you get $(\mathbb{C}^*)^{3g-3}$ Incredibly, or, if you like, trivially, there is a symplectic structure on this Teichmueller space. So Scott Wolpert wrote it down, so you just take $d\ell_i \wedge d\tau_i$. It's a miracle that this is invariant under Dehn twists. It's invariant because it agrees with the obvious symplectic form. The tangent space is the same as the space of deformations of this representation, which is $H^1(\pi_1(G), sl_2\mathbb{R})_e)$. To be even more pedantic, [unintelligible].

Now you can start to see things that are really incredible. When you have a space with a symplectic flow, it cries out for Hamiltonian flows. If you take a loop, you can consider its length. That's a function on Teichmueller space. Each $\alpha \in \mathcal{S}$ gives a function $\tau_g \to \mathbb{R}$ given by $X \mapsto L(X, \alpha)$. Associated to this there is a Hamiltonian flow and vector field $Tw_{t\alpha}\tau_g \to \tau_g$. The flow associated to x is just shearing. This is a map of Teichmueller space preserves lengths of the specified curve. You cut the surface open, twist it a little, nad glue it in. This is, essentially, $d\tau_\alpha$.

The amount of time you have to flow depends on the surface you started with. That's already interesting. What I want to point out, you don't have to be ambitious, you can take for a lamination λ the same kind of length function L_λ which is real analytic. You can still cut open, twist, and reglue. You're taking a limit of Hamiltonian flows. Those paths are called earthquake paths, and they connect any two points on Teichmueller space. Namely

Theorem 1 *(Kirkhoff (?), Thurston)*
There exists a unique earthquake $\lambda \in \mathcal{ML}$ so that $Tw_\lambda)X = Y$.

One corollary is that any finite group $H \subset Mod_g$ has a fixed point on Teichmueller space. Now I realize I forgot to mention that Mod_g acts on τ_g, you can modify the map to X by precomposing with an automorphism of Σ. This modifies, and the quotient of this action is $\mathcal{M}_g = \{RS^{hyp}X$ of genus $g\}$ up to isometry, is the quotient of Teichmueller space by Mod_g. This quotient is not a covering map, it's a description of an orbifold, shouwing that $\mathcal{M}_g$ is a $K(\pi, 1)$ in the orbifold sense, there's an orbifold universal cover that's contractible.

[unintelligible]. There's a solution using length functions and earthquakes to show this fixed point. So you take a combination of length functions, as a convex function, and then show [unintelligible].

Now what were these laminations introduced for in the first place? They're great for compactifying Teichmueller space. Thurston observed that you can glue onto τ_c the space $\mathbb{P}\mathcal{ML}$. This space is a sphere of just the right dimension to form the boundary of this ball, and the action of the mapping class group extends. Thurston introduced this because if you want to analyze the mapping class group, this is a ball, the Brouwer fixed point theorem tells you that you have a fixed point. You can give an analysis by saying what happens if the fixed point is in the interior or the boundary.

What does it mean for X_n in Teichmueller space to converge to $[\lambda] \in \mathbb{P}\mathcal{ML}$. Well, $[\ell(X_n, \alpha)] \to [i(\lambda, \alpha)]$ when [unintelligible]. The lengths of curves are just intersections with some lamination, so then you say the curve is that lamination. Takee X_n where α has length $\frac{1}{n}$. The area is constant by Gauss-Benet.

If you try to predict the length of a curve is approximately a constant times the number of times it crosses α. This is a collection of copies that converge to the class of $[\alpha]$ in δTg. Now you can also take $\tau_\alpha^n(X_0) = X_n$, where we let the markings change, with n Dehn twists. Every time that β crosses α, it gets wrapped by α some large number of times. So this is about $\ell(\alpha, X_0)i(\alpha, \beta)n$. There's a joke that there are two ways of killing a Riemann surface. You can strangle it or wring its neck. You can go straight toward α or go toward along a horocycle. The horocycle is a level set of the length function.

[If I want to mod out the action, in this example I want to restrict to rational points in the circle. Is there a generalization of that?]

I might need five more... days.

We're getting to very good questions. I want to emphasize the central role of laminations. They come up all the time and turn out to be one of the main things you have to appreciate. But they're attached to Teichmueller space, not moduli space. The question, what is the boundary of moduli space, it's not as exotic as $\mathbb{P}\mathcal{ML}$.

Let's ask, what does moduli space actually look like? Here's a first approximation [cusp]. It's not compact in a very simple way. There's only one way to go to infinity, and it's via strangulation. If you look at $X \subset \mathcal{M}_g$ so that the length of the shortest closed geodesic is at least $r > 0$, then this space is compact.

This should remind you of the fact, when it's thin, there's an obvious short loop in moduli space, the Dehn twist around the short loop. Then let me postulate the Teichmueller metric, so that moduli space has finite volume and the injectivity radius goes to zero.

[Could there be many?]

I'll answer this and end. You could add in limits. One thing you can do is glue on $\delta\mathcal{M}_g$, the Deligne Mumford compactification, where you allow the length of a curve to become zero. Here's a surface I want to describe to you. The length of this geodesic is L, of this one is L', and of this one is 0. What you get is a stable curve, a degeneration where there are two components. Fenschel Nielsen coordinates provide beautiful coordinates near any your point. There's a natural stratification from how many curves have been pinched. The symplectic structure is just the same.

That's one way of compactifying moduli space, which gives you a look at the boundary.

The last thing, let me say what the curve complex is, and state a theorem. Remember we had the set $\mathcal{S}$ of simple closed curves. Take these to be the zero simplices, and then say a collection of simple closed curves forms a simplex if and only if you can draw them so that they are disjoint (and not parallel). This is a simplicial complex of dimension $3g-4$ with an action of the mapping class group. One reason people don't like this is that if your collection of curves don't fill the surface, how many one cells are attached to this vertex? Infinitely many. But there is a beautiful theorem of Minsky and Mazur that says that $\mathscr{C}_g$ is a hyperbolic metric space, there is a canonical way to interpolate between them. That's one remark. You can also take $\mathscr{C}_g/Mod_g = \bar{\mathscr{C}_g}$. The top cells are pair of pants decompositions. The number of top cells are the number of trivalent graphs with first Betti number equal to the genus. There are two top dimensional cells in $g = 2$, the theta graph and the barbell. The vertices of the theta graph are the nonseperating curves. This space in the case of genus two, well, these edges are identified, This is something you can get your hands on.

What does moduli space look like? This has big differences from a hyperbolic manifold. A cusp is fundamentally Abelian. This is not true in moduli space. A neighborhood of the end of moduli space maps onto the fundamental group of moduli space. To force yourself to stay in the end, you can hold one curve short somewhere else. This is easy to verify, and quite surprising. The end has all the combinatorial complexity of moduli space. But it has finite volume. There is a known way to answer this. Take moduli space, and multiply the metric by ϵ. Take the limit as $\epsilon \to 0$. This exists in the Gromov-Haussdorff [unintelligible]. The final theorem is that $\epsilon\mathcal{M}$ goes to the cone over $\bar{\mathscr{C}_g}$, not just topologically, but as a metric space, using the Teichmueller metric on $\mathcal{M}_g$.

[Can you say a few words about how this, the foliation given by Dehn twists, it passes to the moduli space?]

No, it doesn't, the curve is not preserved. The foliation are the horocycles, circles tangent to the axis at that point. It's invariant to the parabolic subgroup of Dehn twists, but not the whole group. We had to choose a particular lamination to define this flow. You can take $\mathcal{ML} \times \tau$ and let Mod_g act on both factors.

Does the last theorem work for punctured Riemann surfaces? Almost all the results have natural analogues there.

If you take a periodic surface and scale it, it converges to the real axis. If you take one that admits a homothety, it would be invariant under scaling. The really cool example is if you take hyperbolic space, multiply by ϵ, and the geodesics joining give very thin triangles. So the ideal triangles go to cones on three points. So this doesn't converge, but a subsequence goes to an $\mathbb{R}$-tree.

There are geodesics in Teichmueller space staying a nonnegative distance apart. The curve complex is dual to Deligne Mumford. Pairs of pants give 0-cells of Deligne Mumford, and top cells of the curve complex.

[If you take the homotopy quotient on $\mathscr{C}_g$, do you recover Deligne Mumford?]

You have the notation $\mathscr{C}_g//Mod_g$, where you make this free by crossing with something free. If you take a compactification. Cut off ∞ to make it a manifold with boundary, then $\mathscr{C}_c//Mod_g \cong \delta M_g$