1 Preface

Counts of "special trajectories" of quadratic differentials (saddle points and closed trajectories) are a well-studied subject. Recently it has become clear that they are also examples of "generalized Donaldson-Thomas invariants."

That's interesting in itself: a nice computable example in that theory. But embedding them into this context has also led to several "external" developments:

- They obey a *wall-crossing formula* written down by Kontsevich and Soibelman, which governs how the special trajectories appear and disappear as the quadratic differential is varied;
- They are important ingredients in a systematic scheme for analyzing the asymptotics of differential equations (WKB);
- They are also key ingredients in a new construction of hyperkähler (Ricci-flat) metrics;
- Maybe most interesting, they admit a natural *generalization* to "higher-rank" invariants — attached to any Lie algebra of type ADE (quadratic differentials are the A_1 case);
- They are part of the physics of $\mathcal{N} = 2$ supersymmetric quantum field theory.

In these talks I'll try to describe all of this stuff from a sort of unified perspective. This perspective is work in progress with Davide Gaiotto and Greg Moore — an improvement of the approach we have described before. Some details can therefore be wrong but the basic picture is by now quite clear.

2 $\mathcal{S}$-walls

Fix a compact complex curve C. We are going to do a construction involving quadratic differentials on C:

$$\varphi_2(z) = f(z)\,\mathrm{d}z^2.$$

Any φ_2 determines a 1-parameter family of (singular) foliations $F(\varphi_2, \vartheta)$ of C. Leaves of $F(\varphi_2, \vartheta)$, or "trajectories", are paths along which $e^{-i\vartheta}\sqrt{\varphi_2}$ is a *real* 1-form. (In local coordinates: write $\varphi_2 = \mathrm{d}w^2$, then the leaves are straight lines of inclination ϑ in the w-coordinate.)

$F(\varphi_2, \vartheta)$ has singularities at the zeroes of φ_2. At *simple* zeroes, the singularity is 3-pronged. (Picture.) Assume for now that φ_2 has *only* simple zeroes. In essentially everything that follows, we will focus on the trajectories emerging from the zeroes. Call them "separating trajectories" or "$\mathcal{S}$-walls."

It may happen that an $\mathcal{S}$-wall has *both* ends on a zero. In that case we call it a "special trajectory." These can come in two flavors: either *saddle connections* or *closed trajectories*. (Picture.)

Our interest is in the question: how many special trajectories occur in $F(\vartheta, \varphi_2)$?

First observation: special trajectories can occur at most at countably many ϑ.

Why? φ_2 determines a double cover of C,

$$\Sigma(\varphi_2) = \{\lambda^2 - \varphi_2 = 0\} \subset T^*C.$$

Each special trajectory of φ_2 can be lifted in a canonical way to a 1-cycle on $\Sigma(\varphi_2)$; let $\gamma \in \Gamma = H_1(\Sigma, \mathbb{Z})$ denote its homology class. Call γ the "charge" of the trajectory.

Now, for *any* $\gamma \in \Gamma$ we can define

$$Z_\gamma = \oint_\gamma \lambda$$

with λ the tautological 1-form on T^*C. If γ is the lift of a special trajectory, then we must have $Z_\gamma \in e^{i\vartheta}\mathbb{R}_-$. But there are only countably many $\gamma \in \Gamma$, so this equation can be satisfied only for countably many ϑ. Moreover, once we fix γ, ϑ is determined.

3 Punctures

To reduce potential analytic hazards, fix $n > 0$ marked points $z_1, \ldots, z_n$ on C ("punctures"). Let $\mathcal{B}$ be the space of meromorphic quadratic differentials φ_2 on C, with double poles at all of the z_i. (I believe all of my main statements will be true even without these punctures, but at some moments I will rely on them to simplify the arguments; also, the simplest explicit examples are cases with punctures.) Let $\mathcal{B}' \subset \mathcal{B}$ consist of φ_2 with only *simple* zeroes.

In case with punctures, Σ also has punctures: it is a double cover of $C \setminus \{z_1, \ldots, z_n\}$.

4 DT invariants

Now assume we are in the "generic" situation: all Z_γ are linearly independent over $\mathbb{R}$. (This is a condition on φ_2.) In that case, the possible phenomena are relatively limited. Either isolated saddle connections, or pairs of closed trajectories, bounding an annulus

[Strebel]. We define

$$\Omega(\gamma;\varphi_2) = \begin{cases} 1 & \text{if } F(\varphi_2, \vartheta = \arg -Z_\gamma) \text{ contains a saddle connection,} \\ -2 & \text{if } F(\varphi_2, \vartheta = \arg -Z_\gamma) \text{ contains a closed trajectory,} \\ 0 & \text{otherwise.} \end{cases}$$

So the $\Omega(\gamma;\varphi_2)$ are "counting" the special trajectories, while keeping track of their topological types.

5 Wall-crossing

As we vary the quadratic differential φ_2, the integers $\Omega(\gamma;\varphi_2)$ may change: special trajectories can appear/disappear. The changes occur at codimension-1 loci in the space $\mathcal{B}'$ of quadratic differentials — call these "walls." (Pictures: examples of 2-3 and 2-∞ wallcrossing.)

The problem of "wall-crossing" is: given the $\Omega(\gamma;\varphi_2)$ for one $\varphi_2 \in \mathcal{B}'$, to determine them at some other $\varphi_2 \in \mathcal{B}'$.

Kontsevich-Soibelman wrote a remarkable formula, in an *a priori* different context, which turns out to give a complete solution to this problem. The formula involves some surprising-looking ingredients. Let A be the field of fractions of the group ring $\mathbb{Z}[\Gamma]$. For any $\gamma \in \Gamma$, define a formal automorphism $\mathcal{K}_\gamma$ of A by

$$\mathcal{K}_\gamma(\gamma') = \gamma'(1 - \sigma(\gamma)\gamma)^{\langle\gamma,\gamma'\rangle}.$$

Here we had to throw in the annoying object

$$\sigma : H_1(\Sigma,\mathbb{Z}) \to \{\pm 1\}.$$

A quadratic refinement of the mod 2 pairing. I will not define it unless someone asks; all we will use of it in what follows is

$$\sigma(\gamma) = \begin{cases} -1 & \text{if there is a saddle conn. with charge } \gamma, \\ +1 & \text{if there is a closed loop with charge } \gamma. \end{cases}$$

Now, we draw a picture: vertical axis ϑ, horizontal axis any path in $\mathcal{B}'$. On the picture, put a curve ℓ_γ for each special trajectory, i.e. for each γ with $\Omega(\gamma) \neq 0$: $\ell_\gamma = \{e^{-i\vartheta} Z_\gamma \in \mathbb{R}_-\}$. Now, consider any small "rectangular" paths from (ϑ, u) to (ϑ', u') on this picture. Define

$$S(u) = \prod_{\gamma:\ \vartheta < \arg Z_\gamma < \vartheta'} \mathcal{K}_\gamma^{\Omega(\gamma;u)}. \tag{5.1}$$

The KSWCF says

$$S(u) = S(u'). \tag{5.2}$$

This equation is strong enough to determine all $\Omega(\gamma; u')$ given all $\Omega(\gamma; u)$!

Examples:

1. If $\langle \gamma_1, \gamma_2 \rangle = 1$ then

   $$\mathcal{K}_{\gamma_1} \mathcal{K}_{\gamma_2} = \mathcal{K}_{\gamma_2} \mathcal{K}_{\gamma_1+\gamma_2} \mathcal{K}_{\gamma_1}$$

   This one governs a situation where two saddle connections join into a third.

2. If $\langle \gamma_1, \gamma_2 \rangle = 2$ then

$$\mathcal{K}_{\gamma_1} \mathcal{K}_{\gamma_2} = \left(\prod_{n=1}^{\infty} \mathcal{K}_{(n+1)\gamma_2 + n\gamma_1} \right) \mathcal{K}_{\gamma_1+\gamma_2}^{-2} \left(\prod_{n=\infty}^{1} \mathcal{K}_{(n+1)\gamma_1 + n\gamma_2} \right).$$

This one governs a pair of saddle connections joining into a closed loop plus an infinite tower of other saddle connections.

KSWCF as stated also has an evident interpretation in terms of going around *closed* loops in (ϑ, u) parameter space.

6 Path lifting

Now, let's try to explain *why* KSWCF is true.

We begin by introducing a strange-looking construction: a new "thing you can do with a quadratic differential."

Fix a pair (φ_2, ϑ). Recall the double cover $\Sigma \to C$, and the $\mathcal{S}$-walls on C.

To every open path $\mathcal{P}$ on C, we'll attach $L(\mathcal{P})$, a formal $\mathbb{Z}$-linear combination of open paths on Σ, in a way which is "compatible with concatenation", "twisted homotopy invariant."

First, suppose $\mathcal{P}$ does not cross any $\mathcal{S}$-walls. In this case, $F(\mathcal{P})$ is the formal sum of the 2 lifts of $\mathcal{P}$ to Σ:

$$L(\mathcal{P}) = \mathcal{P}^1 + \mathcal{P}^2.$$

Next, suppose $\mathcal{P}$ crosses exactly one $\mathcal{S}$-wall, at an intersection point z. In this case, $F(\mathcal{P})$ will involve *three* terms. Two are the naive lifts as before. The third is a path which "takes a detour". The *lift* of the $\mathcal{S}$-wall to Σ is an open path $S(z)$ running from say z^1 to z^2. We have

$$L(\mathcal{P}) = \mathcal{P}^1 + \mathcal{P}^2 + \mathcal{P}^1_+ S(z) \mathcal{P}^2_-$$

where the product means concatenation.

Finally, suppose $\mathcal{P}$ is a general path which misses the branch points: then $L(\mathcal{P})$ is constructed by breaking $\mathcal{P}$ into pieces and requiring $L(\mathcal{P}\mathcal{P}') = L(\mathcal{P})L(\mathcal{P}')$ (where we define the product of non-composable paths to be zero).

7 Homotopy invariance

We'd like to ask this to factor through homotopy, but that won't quite work. You can see that just by considering a closed loop $\mathcal{P}$ around a branch point.

Instead, pass to *twisted homotopy*: replace smooth paths by their lifts to the unit tangent bundles $\tilde{C}$, $\tilde{\Sigma}$. Identify any path which winds once around the fiber with -1. Then, claim: our construction factors through this "twisted homotopy." (Also, it can be extended to arbitrary paths on $\tilde{C}$, not just ones which arise as lifts of smooth paths on C.)

To check this homotopy property, two illustrative computations:

1. a path which crosses an $\mathcal{S}$-wall twice in opposite directions;
2. a loop around a branch point.

Show one part of the branch point computation: 2 terms cancelling. (NB, it wouldn't have worked without these detours.)

8 Lifting closed paths

In particular, we can consider $L(\mathcal{P})$ for $\mathcal{P}$ a path beginning and ending at the same z. $L(\mathcal{P})$ is a sum of paths beginning and ending at preimages z^i, some open, some closed. Define $T(\mathcal{P}) \in A$ as "trace" of $L(\mathcal{P})$: drop open paths, and replace simple closed curves by their homology classes.

9 Morphisms

We've defined a rule which assigns to each closed path $\mathcal{P}$ an element $T(\mathcal{P}) \in A$, a formal linear combination of classes in $H_1(\Sigma, \mathbb{Z})$. Now, we may ask: how does $T(\mathcal{P})$ change as we vary (φ_2, ϑ)?

For "small" variations which don't change the topology of the $\mathcal{S}$-walls, $T(\mathcal{P})$ does not change (or better, varies continuously, as Σ varies). But when the $\mathcal{S}$-walls do change topology, $T(\mathcal{P})$ jumps.

Simplest example: two $\mathcal{S}$-walls crossing. At the moment when they cross we have a saddle connection, which lifts to some loop S. Compare any $L(\mathcal{P})$ before and after the crossing: they differ by a universal transformation, which can be described as an action directly on the paths on Σ. A path a which crosses S exactly once is split into two pieces a_1 and a_2 by S; after the crossing it is transformed by

$$a = a_1 a_2 \mapsto a_1 (1 + S)^{\langle a, s \rangle} a_2.$$

All $L(\mathcal{P})$ are simply modified by this transformation.

After tracing, this implies that $T(\mathcal{P})$ jumps by

$$\gamma \mapsto \gamma (1 + \gamma_S)^{\langle \gamma, \gamma_S \rangle}.$$

This is exactly the transformation we previously called $\mathcal{K}_{\gamma_S}$.

More interesting example: a tower of windings collapsing. At the moment of collapse we have a closed trajectory, which again lifts to some loop S. Compare any $L(\mathcal{P})$ before and after the crossing: they differ by a universal transformation, which can be described as an action directly on the paths on Σ. Namely: any path a which crosses S is split into two pieces a_1 and a_2 by S; after the crossing it is transformed by

$$a = a_1 a_2 \mapsto a_1 (1 - S)^{-\langle a, s \rangle} a_2.$$

Moreover, these closed loops come in *pairs*. After tracing, this implies that $T(\mathcal{P})$ jumps by

$$\gamma \mapsto \gamma(1-\gamma_S)^{-2\langle\gamma,\gamma_S\rangle}$$

This is exactly the transformation we previously called $\mathcal{K}_{\gamma_S}^{-2}$.

10 Proving the WCF

So far we have produced $T(\mathcal{P}) \in A$ for each path $\mathcal{P}$ on C, and shown that as we vary (φ_2, ϑ) along some path in $\mathcal{B}' \times S^1$, all $T(\mathcal{P})$ get transformed by the appropriate product of $\mathcal{K}_\gamma^{\Omega(\gamma)}$. If we vary along a *closed* path then the $T(\mathcal{P})$ must return to themselves. This would prove the desired KSWCF for the $\Omega(\gamma)$, *if* the $T(\mathcal{P})$ generate the whole A.

Indeed the $T(\mathcal{P})$ do generate A. (Essentially due to Fock-Goncharov). One way to understand this: "tropicalization" — let the "leading term" $M(\mathcal{P})$ be the γ appearing in $T(\mathcal{P})$ with greatest $\mathrm{Re}(e^{\mathrm{i}\vartheta} Z_\gamma)$. Then show that for any $\gamma \in \Gamma$ there exists a path $\mathcal{P}$ with $M(\mathcal{P}) = \gamma$.

11 Flat connections and Fock-Goncharov coordinates

In trying to understand the WCF we were led to the "path lifting" construction. This construction has other uses: as we will now see it gives a way of relating abelian ($GL(1)$) connections on Σ and non-abelian ($GL(2)$) connections on $C \setminus \{z_1, \ldots, z_n\}$.

First, recall a "naive" way of trying to relate the two. Suppose given a complex line bundle $\mathcal{L}$ with flat connection on Σ. The push-forward $E = \pi_* \mathcal{L}$ is a rank 2 bundle on Σ. Does it acquire a flat connection? Locally, *away from branch points*, E is just the direct sum of 2 line bundles $\mathcal{L}_1$ and $\mathcal{L}_2$, each with a flat connection, so E gets one too: the parallel transport along a path $\mathcal{P}$ is just the sum of the parallel transports along the lifts of $\mathcal{P}$.

But this flat connection in E cannot possibly extend over the branch points: it has *monodromy* (permutation matrix).

Now, our "improved" method. We'll construct the corrected bundle by building its *sheaf of flat sections*. By definition, a flat section of the improved bundle will be a section of E which is invariant under the improved parallel transport: i.e. under the *abelian* parallel transport along the paths given by $F(\mathcal{P})$. (So it's discontinuous as a section of E, but it will be continuous as a section of the new glued bundle.)

Our homotopy invariance property means this is indeed a (twisted) *flat* connection. (Could get rid of the twisting by choosing spin-structures on C and Σ, but let's not.) So we get a "non-abelianization" map $\nabla^{ab} \mapsto \nabla$, from the moduli space of flat $GL(1)$-connections over Σ to the moduli space $\mathcal{M}$ of flat $GL(2)$-connections over C. (More precisely, flat $GL(2)$-connections over C with the extra data of a flag at each puncture.)

In this picture $T(\mathcal{P})$ has a particularly concrete meaning: it is giving the trace of the holonomy of ∇ around $\mathcal{P}$, as a function of the holonomies $\mathcal{X}_\gamma$ of ∇^{ab} around loops γ in Σ.

Do we get all $GL(2)$-connections ∇ this way? Almost: this "non-abelianization" map is actually an isomorphism onto an open dense patch of $\mathcal{M}$. This is basically a result of Fock-Goncharov: strictly speaking they studied $SL(2)$-connections, but the overall $GL(1)$ part goes through trivially (I hope; could be some $\mathbb{Z}_2$ subtleties here to fuss with).

Their proof goes by constructing the explicit inverse of our map: "abelianization." Since a $GL(1)$-connection is specified by its $\mathbb{C}^\times$-valued holonomies, concretely this amounts to specifying an open dense coordinate patch on the space of flat $GL(2)$-connections. The $SL(2)$ part is the interesting part. Fock-Goncharov build these coordinates by taking cross-ratios of flat sections. (Picture.)

More precisely, this is one coordinate system for every $\mathcal{S}$-wall network; when the $\mathcal{S}$-wall network changes topology, the coordinate system jumps. The different coordinate systems are related by "cluster transformations": a concrete instantiation of the $\mathcal{K}_\gamma$ we wrote before, now acting on actual *functions* rather than formal variables. Quite interesting structure, for reasons I'm not fully competent to explain.

12 Higher rank

The story seems to have a natural *generalization* to "higher rank."

Starting point: replace the quadratic differential φ_2 by a tuple $(\varphi_2, \ldots, \varphi_K)$ where φ_i is a section of K^i.

The special trajectories we studied before could be understood as loci where the $\mathcal{S}$-wall network jumped. So: what is the appropriate generalization of the $\mathcal{S}$-walls here?

As before, we can define a *spectral curve* by

$$\Sigma = \{\lambda^K + \sum_{n=2}^{K} \varphi_n \lambda^{K-n} = 0\} \subset T^*C. \tag{12.1}$$

A K-fold cover of Σ.

For any choice of a labeling of sheets of Σ (locally defined), we thus have K 1-forms on C, $\lambda_1, \ldots, \lambda_K$. We define an ij-trajectory to be one along which the 1-form $\lambda_i - \lambda_j$ is real (and positive). Our $\mathcal{S}$-wall network will be built out of these ij-trajectories.

Moreover, using our $\mathcal{S}$-wall network we want to be able to build a path-lifting rule, with the same kind of twisted homotopy invariance as we had in the $K = 2$ case.

Branch points are labeled by transpositions (ij). To get the homotopy invariance around each (ij) branch point, we will need to have 3 $\mathcal{S}$-walls emerging. (Draw the picture.) But now we have a new problem: the $\mathcal{S}$-walls might *collide*. Suppose an (ij) and a (jk) $\mathcal{S}$-wall collide. In this case we will have failure of homotopy invariance (a loop around the collision point is not equivalent to a trivial one). The way to fix it is to add a *new* (ik) $\mathcal{S}$-wall emerging from the branch point. This new $\mathcal{S}$-wall then evolves along with the rest. We build up a rather complicated, but controlled, structure. (NB, it is also possible for $\mathcal{S}$-walls to die.) If there are punctures, with each φ_i having a pole of order i, then all $\mathcal{S}$-walls eventually wind up at the punctures.

Using this new $\mathcal{S}$-wall network we can define a path-lifting rule $\mathcal{P} \mapsto L(\mathcal{P})$, the straightforward generalization of what we did in the $K = 2$ case; take traces to get $T(\mathcal{P})$. As before, the crucial question is: when does $T(\mathcal{P})$ jump discontinuously? Answer: whenever two $\mathcal{S}$-walls collide head-to-head.

The most obvious way for this to happen is to have a saddle connection, like before. But there are also more interesting possibilities. (Show examples.) Whenever the $\mathcal{S}$-walls collide, there is a corresponding *finite subnetwork*. Its lift to Σ defines a charge

$\gamma \in \Gamma = H_1(\Sigma, \mathbb{Z})$.

The analysis of $T(\mathcal{P})$ at the special loci $e^{-i\vartheta} Z_\gamma \in \mathbb{R}_-$ goes much like before: they jump by an automorphism $\mathcal{K}_\gamma^c$ where c depends on the topology of the subnetwork. Simple examples: three-pronged network gives $c = +1$, loop with attached edge gives -1. Conjecture: every network gives ± 1. At any rate, it's in principle straightforward to determine the contribution from any particular network. So we will obtain invariants $\Omega(\gamma)$ like before; and the same argument we used would be expected to prove KSWCF in this setting too (if there are "enough" $T(\mathcal{P})$.)

All the usual questions about special trajectories of quadratic differentials should be interesting for these finite subnetworks, too. (e.g. how many of them with length $\leq L$?)

Our path-lifting construction leads to "non-abelianization" map relating $GL(1)$-connections on Σ to $GL(K)$-connections on C. Conjecture: as before, this map is onto an open dense subset of the moduli space $\mathcal{M}$ of such connections. So each $\mathcal{S}$-wall network would give a set of "Fock-Goncharov-like" coordinates on $\mathcal{M}$. If we take $\varphi_3, \ldots, \varphi_K$ to be very small and arranged in a particular way, we can actually identify them with the honest Fock-Goncharov coordinates for higher rank. (Show picture of spin-lift and the higher-rank flip.)

13 WKB

There is another well-known approach to "abelianizing" a connection, or more precisely a *family* of connections: WKB. Suppose given a family of $GL(K)$-connections of the form

$$\nabla = \varphi/\zeta + D(\zeta) \tag{13.1}$$

where φ is a $gl(K)$-valued matrix and D a connection, regular at $\zeta = 0$. One often wants (e.g. in quantum mechanics) to study the flat sections ($\nabla\psi = 0$) in the limit $\zeta \to 0$. WKB approximation says: just *diagonalize* φ,

$$\varphi = diag(\lambda_i) \tag{13.2}$$

and then construct formal solutions in the form

$$\psi_i^{WKB} = \exp\left[\int \lambda_i/\zeta\right] e_i(\zeta) \tag{13.3}$$

where $e_i(\zeta)$ is a power series in ζ, determined by iteratively plugging into the flatness equation. The $\psi_i^{WKB}(\zeta)$ then look like they define an *abelian* connection over Σ of the form

$$\nabla^{ab,WKB} = \lambda/\zeta + D^{ab,WKB}(\zeta) \tag{13.4}$$

whose pushforward would be ∇.

But as we know, you can't really construct ∇ this way (if you could, ∇ would have monodromy around branch points). So what goes wrong? The point is that the series defining ψ_i^{WKB} typically is not a *convergent* series: it only allows us to abelianize the connection in a *formal neighborhood* of $\zeta = 0$.

14 Comparing our story with WKB

We constructed a "de-abelianization" map, using the additional datum of a pair $(\vartheta, \varphi_2, \ldots, \varphi_K)$. Conjecture (true for $K = 2$): it's invertible, so gives "abelianization" map (defined on dense open subset).

Now suppose as above that $\nabla = \varphi/\zeta + D(\zeta)$, and take φ_i to be the coefficients of the characteristic polynomial of φ. Apply the abelianization map.

This in particular provides *actual* flat sections ψ_i on the complement of the $\mathcal{S}$-walls. (Concretely, for $K = 2$, the exponentially-smaller monodromy eigensections at the "nearest" puncture.) These ψ_i *jump* at the $\mathcal{S}$-walls.

Conjecture (true for $K = 2$): this construction is *compatible* with the WKB method, in the sense that the actual flat sections ψ_i have asymptotic expansion given by ψ_i^{WKB}, as $\zeta \to 0$. (Although they are not continuous!)

This WKB property is vital for some applications. It wouldn't have worked if we chose a "random" network; depends on using the network that's really defined by $(\varphi_2, \ldots, \varphi_K)$.

The jumps of ψ_i^{WKB} are related to "WKB connection formula." So this is a re-telling of a somehow familiar story (Ecalle, Voros etc.) The part involving closed geodesics may be new, also the higher rank story.

15 Hyperkahler metrics

One application of this WKB analysis is a new way of thinking about the Hitchin system.

An amazing fact [Hitchin, Simpson, Corlette, Donaldson]. Consider the space $\mathcal{M}$ of flat $GL(K, \mathbb{C})$-connections. Given any ∇ (subject to some "stability" condition, automatically satisfied in our case with generic punctures) you can find a decomposition

$$\nabla = \varphi + D + \bar{\varphi}, \tag{15.1}$$

where D is unitary and $\bar{\varphi}$ is adjoint of φ (with respect to some metric), and we have

$$F_D + [\varphi, \bar{\varphi}] = 0, \tag{15.2}$$

$$\bar{D}\varphi = 0. \tag{15.3}$$

So, now, let's try starting from the pair (D, φ). We can build ∇ from them, but in fact we can build a 1-parameter *family* of flat connections:

$$\nabla^{(\zeta)} = \varphi/\zeta + D + \bar{\varphi}\zeta. \tag{15.4}$$

So $\mathcal{M}$ is identified with the complex manifold of flat $GL(K, \mathbb{C})$-connections in many different ways. i.e. $\mathcal{M}$ has many different complex structures $J^{(\zeta)}$, $\zeta \in \mathbb{C}^\times$. In particular, if you fix $J_1 = J^{(\zeta=1)}$, $J_2 = J^{(\zeta=i)}$, $J_3 = J^{(\zeta=0)}$, then $J_1 J_2 = J_3$ and cyclic permutations: quaternion algebra.

$\mathcal{M}$ also has a holomorphic symplectic form: defined in terms of symplectic quotient starting from

$$\varpi^{(\zeta)} = \int_C \operatorname{Tr} \delta\nabla^{(\zeta)} \wedge \delta\nabla^{(\zeta)}. \tag{15.5}$$

Expands explicitly as

$$\varpi^{(\zeta)} = \frac{\omega_1 + i\omega_2}{\zeta} + \omega_3 + (\omega_1 - i\omega_2)\zeta \tag{15.6}$$

where ω_1, ω_2, ω_3 are three *real* symplectic forms. In fact, they are Kähler forms with respect to the three complex structures J_1, J_2, J_3, determining a *single* Riemannian metric g. g is thus called a hyperkähler metric. In particular it's Ricci-flat.

Now, suppose we want to actually *construct* this metric in some concrete terms. It's enough to construct $\varpi^{(\zeta)}$. But, the symplectic structure looks rather complicated. Good news: the abelianization map we have discussed is actually also a symplectomorphism! And the symplectic structure on the space of $GL(1)$-connections is very simple: just $\langle d \log \mathcal{X}, d \log \mathcal{X} \rangle$.

Why does this help? After all, $\mathcal{X}$ is still a complicated function on the original $\mathcal{M}$. But, we know a lot about it. WKB determines its leading asymptotic as $\zeta \to 0, \infty$ as $\mathcal{X} \sim \exp[Z_\gamma/\zeta]$, $\mathcal{X} \sim \exp[\bar{Z}_\gamma \zeta]$ respectively. And we know where its discontinuities are. Thus we have a *Riemann-Hilbert problem* which can be solved by an explicit integral equation.

That gives a recipe for constructing the actual ϖ and hence the hyperkähler metric g.

16 Some physics

Now, what is the meaning of all this for physicists?

In many (all?) cases, DT invariants can be understood in terms of 4-dimensional supersymmetric quantum field theory (QFT). I won't say what a QFT is: suffice to say that it is supposed to have a Hilbert space $\mathcal{H}$, with subspace $\mathcal{H}^1$ (space of "1-particle states") which forms a unitary representation of a super extension of the Poincare group $ISO(3,1)$.

$ISO(3,1)$ has a Casimir operator "M^2" which, acting on $\mathcal{H}_1$, tells us the mass-squared of the particles. Our extension also has a second Casimir operator Z, complex-valued even in unitary representations. All unitary representations have $M \geq |Z|$. Moreover the representations with $M = |Z|$ are special ("short"). States in these representations are called "BPS states." There is an index Ω which counts the multiplicity of such representations, which cannot change under continuous deformations of the representation $\mathcal{H}^1$ (designed so that it vanishes for "long" representations).

The theories we consider are actually (in the IR) *abelian gauge theories*. Like electromagnetism. In such a theory, $\mathcal{H}$ has a decomposition into "charge sectors" labeled by a lattice Γ of electromagnetic charges,

$$\mathcal{H} = \bigoplus_{\gamma} \mathcal{H}_{\gamma}.$$

Moreover the Casimir operator Z acts as a scalar Z_γ in each sector.

We can compute the index Ω in each $\mathcal{H}_\gamma$ separately: get $\Omega(\gamma) \in \mathbb{Z}$. Now, we may ask, what happens when we vary the parameters of the theory? For small variation, just get some small variation of each $\mathcal{H}^1_\gamma$, so $\Omega(\gamma)$ is invariant. But exactly when different Z_γ become aligned, $\mathcal{H}^1_\gamma$ actually "mixes with the continuum" and $\Omega(\gamma)$ becomes ill-defined. Thus we have the possibility of wall-crossing. Indeed, there is a nice semiclassical picture of a "bound state" which decays [Denef].

How to get such an $\mathcal{N} = 2$ supersymmetric QFT? One attractive option: string theory on Calabi-Yau threefold. IIB: BPS states come from D3-branes wrapped around special Lagrangian 3-cycles. But there is also a second way of getting an $\mathcal{N} = 2$ supersymmetric QFT. Namely, there is a somewhat mysterious 6-dimensional QFT ("theory $\mathcal{X}$"), or more precisely one $\mathcal{X}_{\mathfrak{g}}$ for each ADE algebra $\mathfrak{g}$ (plus more trivial abelian ones). Formulating this theory where we take spacetime to be $C \times \mathbb{R}^{3,1}$ we get the theory we have been (implicitly) discussing.

The two pictures are not unrelated: one way of realizing theory

$\mathcal{X}$ is as the Type IIB string theory near an ADE singularity. Then by wrapping D3-branes around the collapsing 2-cycles of the ADE singularity, we get effective "strings."