Notes on spectral networks

Andrew Neitzke

These are notes for a spring 2021 lecture series. They are a work in progress.

1. Introductory remarks

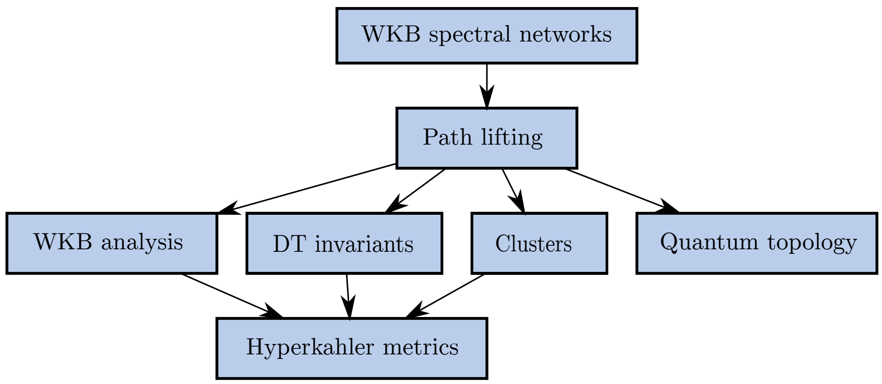

The rough plan is above. It is not necessarily the logical flow of the subject — really all arrows should point in both directions — but it is the path we will take. [actually not the path we really took, it seems]

The story I will explain follows along the lines of joint work with Gaiotto-Moore, Hollands, Yan, Freed. Related work by Fenyes. All closely connected with ideas of Kontsevich-Soibelman, Fock-Goncharov, Hitchin, Simpson, Voros, Ecalle, many others.

[thank A. Goncharov, B. Morrissey, T. Bridgeland, ... for suggesting to formulate on root cover; thank Ian Le for correcting a mistake]

[references]

2. WKB spectral networks

2.1. Data.

The data with which we work is:

- A compact Riemann surface C.
- An integer $N \geq 1$.
- A tuple $\vec{\phi} = (\phi_1, \ldots, \phi_N)$ where ϕ_i is a meromorphic section of $K_C^{\otimes i}$.

We will hold C and N fixed throughout and thus usually omit them from the notation, but it will be important sometimes to allow $\vec{\phi}$ to vary.

ϕ_1 does not play much of a role, so we often simplify by assuming $\phi_1 = 0$. (We should think that N is really standing in for $\mathfrak{gl}(N)$, and then getting rid of ϕ_1 is related to considering only traceless matrices, i.e. reducing from $\mathfrak{gl}(N)$ to $\mathfrak{sl}(N)$.)

The case $N = 1$ is more or less trivial at least for today. $N = 2$ is already interesting and considerably simpler than $N \geq 3$.

2.2. Spectral curve.

Definition 2.1 (Spectral curve). Define the spectral curve

$$\Sigma_{\vec{\phi}} = \left\{ \sum_{i=0}^{N} \phi_i \, y^{N-i} = 0 \right\} \subset T^*C, \tag{2.1}$$

where we define $\phi_0 = 1$.

When $\Sigma_{\vec{\phi}}$ is smooth reduced, the projection $\pi : \Sigma_{\vec{\phi}} \to C$ is a branched N-fold cover. For $N = 3$ a cartoon would be:

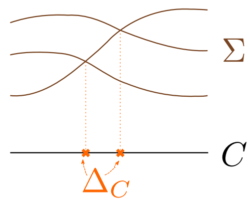

Definition 2.2 (Simple data). We say $\Sigma_{\vec{\phi}}$ is *simple* if it is smooth reduced and also has only simple ramification (ie all ramification points are of index 1) as illustrated above.

For what we will do later, the restriction to smooth reduced $\Sigma_{\vec{\phi}}$ will be essential, while the restriction to simple $\Sigma_{\vec{\phi}}$ will be merely convenient.

Example 2.3 (Spectral curve for $N = 2$). If $N = 2$ and $\phi_1 = 0$, then $\Sigma_{\vec{\phi}} = \{y^2 + \phi_2 = 0\}$. $\Sigma_{\vec{\phi}}$ is smooth reduced just if ϕ_2 has all zeroes simple, and $\Sigma_{\vec{\phi}}$ is simple iff it is smooth reduced.

When $\vec{\phi}$ is fixed we sometimes write Σ for $\Sigma_{\vec{\phi}}$. We let $\Delta_C \subset C$ denote the branch locus of π.

For concrete pictures and computations we sometimes choose a local trivialization of Σ over some simply connected domain $U \subset C \setminus \Delta_C$, i.e. identify the components of $\pi^{-1}(U)$ with $\{1, \ldots, N\}$. In particular, we can take U to be the complement of a set of "branch cuts" emanating from the points of Δ_C. Then each branch cut is labeled by an element of S_N. Paths on Σ are represented by paths on U with a number indicating which preimage of the path we take; the number has to transform according to the permutation when we go across a cut.

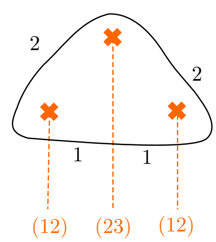

When Σ is simple, these elements are all transpositions (and thus we do not have to specify a co-orientation of the cuts).

2.3. **Foliations.** On a surface with a quadratic differential ϕ_2 there is a well-known foliation. We are going to construct a generalization of that.

Definition 2.4 (Oriented foliation for a holomorphic 1-form). For a Riemann surface S equipped with a nowhere vanishing holomorphic 1-form ρ, we have the distribution $\ker \operatorname{Im} \rho$, which integrates to give a foliation F_ρ of S. This foliation is oriented by $\operatorname{Re} \rho$. For example, if $S = \mathbb{C}$ and $\rho = \mathrm{d}z$, then F_ρ is the foliation by horizontal lines oriented to the right. Note $F_{-\rho}$ is the same foliation with the opposite orientation.

If ρ has an isolated zero then we make the same definition, now getting a singular foliation. Below are pictures of the foliation (up to conformal transformation) in a neighborhood of a generic point, a first-order zero, or a second-order zero of ρ. The pattern continues: for an

n-th order zero one gets $2(n+1)$ "wedges" coming together. We'll call such a singularity a $2(n+1)$-fold point.

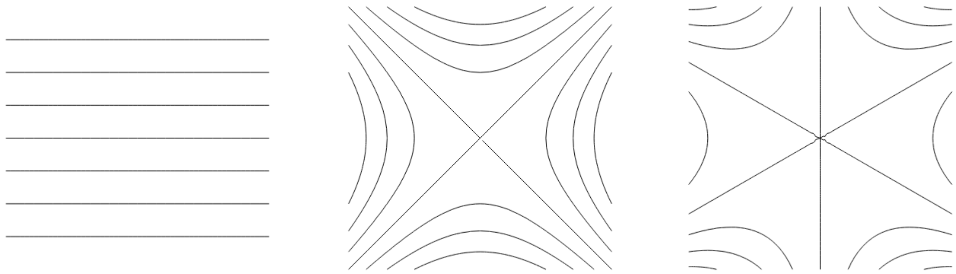

More generally, for $\vartheta \in \mathbb{R}/2\pi\mathbb{Z}$, we also define

$$F_\rho^\vartheta = F_{e^{i\vartheta}\rho}\,.$$

Our foliation is not going to be directly on C but rather on a slightly complicated-sounding covering (not so bad in practice):

Definition 2.5 (Root curve). Suppose Σ (as above) is smooth. The *root curve* is the closure of the fiber product with diagonal removed,

$$\widetilde{\Sigma} = \overline{\{(y,y') \in \Sigma \times \Sigma : \pi(y) = \pi(y'), y \neq y\}} \subset T^*C \times T^*C. \tag{2.2}$$

This is a smooth $(N^2 - N)$-sheeted branched covering of C, perhaps with some points deleted over the poles of $\vec{\phi}$.

There is a natural involution of $\widetilde{\Sigma}$ given by $\mu(y,y') = (y',y)$. Projection on the first and second factors gives maps $p_1, p_2 : \widetilde{\Sigma} \to \Sigma$, with $p_2 = p_1 \circ \mu$.

Example 2.6 (Root curve for $N=2$). For $N=2$, the root curve is just another language for talking about Σ and its standard covering involution: indeed the maps p_1, p_2 are both isomorphisms $\widetilde{\Sigma} \simeq \Sigma$, differing by composition with the covering involution. Under either of these isomorphisms, moreover, μ is taken to the covering involution.

Example 2.7 (Root curve for $N=3$). For $N=3$, the situation is already more interesting: $\widetilde{\Sigma}$ is a 6-fold covering of C, while Σ is a 3-fold covering, and the maps p_1, p_2 are branched double covers. Above a simple branch point $p \in \Delta_C$, $\widetilde{\Sigma}$ has 3 ramification points, all of index 1 (i.e. the 6 sheets come together in 3 pairs). But these 3 ramification points are different: one is fixed by the involution μ while the other two are exchanged.

If we trivialize Σ locally then there is an induced local trivialization of $\widetilde{\Sigma}$, labeling its sheets by ordered pairs (i,j) with $1 \leq i,j \leq N$ and $i \neq j$. So paths on $\widetilde{\Sigma}$ can be represented by pictures like the one below.

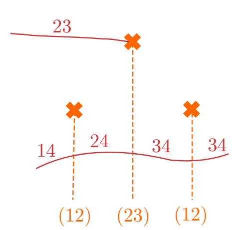

Definition 2.8 (Foliation of the root curve). Let λ denote the canonical (Liouville) 1-form on T^*C; its restriction is a holomorphic 1-form on Σ. Define

$$\rho = p_1^*\lambda - p_2^*\lambda. \tag{2.3}$$

ρ is a holomorphic 1-form on $\widetilde{\Sigma}$, with zeroes only over the branch points. Thus we have the foliation F_ρ^ϑ of $\widetilde{\Sigma}$. It has singularities only over the branch points.

Since ρ is the only 1-form we will use on $\widetilde{\Sigma}$ we will often drop it from the notation and just call the foliation F^ϑ.

We have $\mu^*\rho = -\rho$. It follows that μ preserves the leaves of F^ϑ but reverses their orientation.

Example 2.9 (Foliations for $N = 2$). The simplest case is $N = 2$. Then $\widetilde{\Sigma}$ is just a double cover of C, and μ is the covering involution. So the oriented foliation F^ϑ of $\widetilde{\Sigma}$ descends to an unoriented foliation of C. Around simple zeroes of ϕ_2 it has a three-pronged structure as shown.

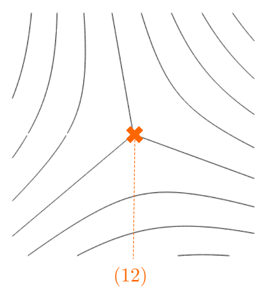

Around poles of ϕ_2 the behavior depends on the order of the pole as sketched below.

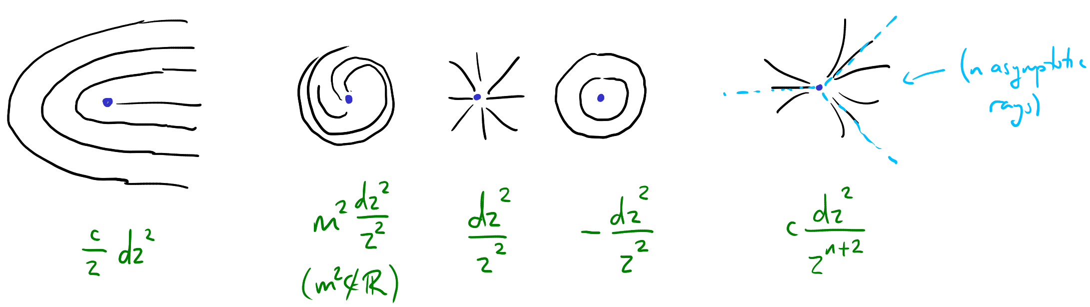

For ϕ_2 meromorphic with at least one pole of order ≥ 2, with only simple zeroes, and sufficiently generic, the global picture looks like:

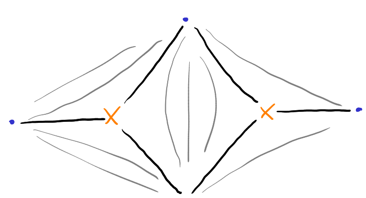

When ϕ_2 has *only* order 2 poles, this picture determines an ideal triangulation of C, as in the following picture.

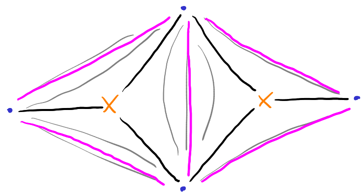

When poles of higher order are allowed, one gets a triangulation of a slightly different surface, where we replace each order-$(n+2)$ pole by a boundary circle with n marked points, a la Fock-Goncharov [add ref] — this reflects the behavior of the leaves near the singularity shown above. An extreme case is to take $C = \mathbb{CP}^1$ and $\phi_2 = P_m(z)\mathrm{d}z^2$ with P_m polynomial of degree m; then we have a pole of order $n+2 = m+4$ at $z = \infty$ and we get an ideal triangulation of a disc with $m+2$ marked points around the boundary, or equivalently a triangulation of an $(m+2)$-gon.

In this picture the leaves which end on the branch points play a particularly important role: they make up the *critical graph* $CG(\vartheta)$ of the quadratic differential ϕ_2. The notion of spectral network is a generalization of this structure to higher N.

Example 2.10 (Foliations for $N = 3$). The foliation F^ϑ for $N = 3$ looks considerably more complicated. Around a generic point the projection of F^ϑ to C looks (topologically, not conformally!) like:

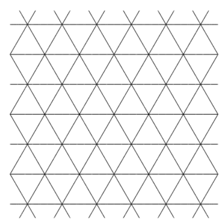

There are also some degenerate loci (caustics) where the three foliations become parallel.

Around a branch point it is hard to draw the full picture, but we can draw at least part of the projection: it has the same three-pronged structure that we saw in the $N = 2$ case above, Example 2.9.

2.4. **Solitons.** This subsection was actually omitted from the first lecture, but might be useful anyway.

If c is a 1-chain on $\widetilde{\Sigma}$, let $p(c)$ be a 1-chain on Σ given by

$$p(c) = (p_1)_* c - (p_2)_* c\,. \tag{2.4}$$

(So if in a local trivialization c carries the label ij then $p(c)$ consists of two components, one carrying label i and one carrying label j, with the second one oppositely oriented.) Note it has

$$\int_{p(c)} \lambda = \int_c \rho. \tag{2.5}$$

Definition 2.11 (Topological solitons). Fix $z \in C$ and $y, y' \in \pi^{-1}(z)$. A *topological soliton from y to y'* is a 1-chain c on $\widetilde{\Sigma}$, such that $\partial(p(c)) = y' - y$.

The projection of a topological soliton from $\widetilde{\Sigma}$ to C looks like a graph (generically trivalent), with one leaf at z and all other leaves at branch points. The labels on the edges of the graph are constrained: there is a balancing condition at each internal vertex, and also a condition at each branch point, coming from the condition that $\partial(p(c))$ doesn't contain any points lying over the internal vertices or branch points. Examples of projections of solitons obeying the conditions are shown below.

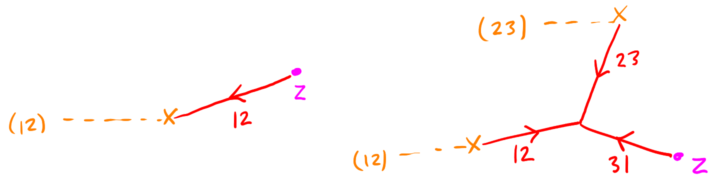

Definition 2.12 (WKB solitons). Fix $z \in C$ and $y, y' \in \pi^{-1}(C) \subset \Sigma$. A *WKB soliton from* y *to* y' *with phase* ϑ is a topological soliton from y to y' made up of paths which travel along the foliation F^ϑ (permitting turns at the singularities).

So the projection of a WKB soliton to C is a graph as before, but now with some rigidity: the edges are required to be leaves of the foliation with matching labels.

Definition 2.13 (WKB spectral network). The *WKB spectral network* $\mathcal{W}(\vartheta)$ is the set of all $(y, y') \in \widetilde{\Sigma}$ such that there exists a WKB soliton from y to y' with phase ϑ.

Example 2.14 (WKB spectral network when $N = 2$). When $N = 2$ and $\phi_1 = 0$, all WKB solitons are just segments whose projection runs from z to a branch point. It follows that the projection of $\mathcal{W}(\vartheta)$ to C is the critical graph $CG(\vartheta)$.

2.5. **Computing the WKB spectral network.** In the actual lecture, the algorithm below became the definition.

Proposition 2.15 (Algorithm for computing $\mathcal{W}(\vartheta)$). If $\Sigma_{\vec{\phi}}$ is simple, then we can compute $\mathcal{W}(\vartheta)$ as follows. For each branch point $p \in \Delta_C$, there is a unique μ-fixed ramification point $r_p \in \widetilde{\Sigma}$ lying over p. ρ has a second-order zero at r_p, so there are three leaves of F^ϑ passing through r_p. Let $\mathcal{C}$ be the smallest set of *half-leaves* of F^ϑ such that:

- $\mathcal{C}$ must include, for each p, the three half-leaves beginning at r_p which are oriented away from r_p.

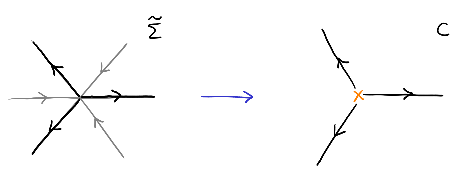

- Suppose $\ell_1, \ell_2 \in \mathcal{C}$ are such that $p_1(\ell_1)$ intersects $p_2(\ell_2)$ transversely at a point $y \in \Sigma$, with preimages $(y, y_1) \in \ell_1$ and $(y_2, y) \in \ell_2$. Then $\mathcal{C}$ must also include a third half-leaf beginning at (y_2, y_1). The projection of these three half-leaves to the base C then looks like the picture below in a neighborhood of $\pi(y)$:

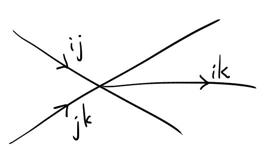

We emphasize that the intersection point $\pi(y) \in C$ is a regular point, neither a branch point of $\Sigma \to C$ nor a singularity of $\vec{\phi}$; the only thing that is special about this point is that it happens to be the place where $p_1(\ell_1)$ and $p_2(\ell_2)$ intersect.
- (This one is technical: in the generic case it will not be needed.) Suppose $\ell \in \mathcal{C}$ is a half-leaf which runs into a singularity of F^θ at a point $r \in \Sigma$, which is not μ-fixed. (In our local notation on C, this means a leaf labeled ij running into a branch point of type jk, for $k \neq i$.) Then ρ has a first-order zero at r, so the foliation F_ϑ has a 4-fold point at r. In this case $\mathcal{C}$ must include both of the half-leaves beginning at r which are oriented away from r.

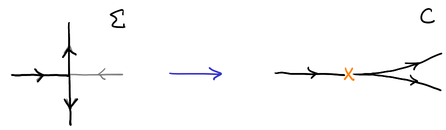

Then $\mathcal{W}(\vartheta)$ is the union of all half-leaves in $\mathcal{C}$.

Remarks:
- When the ϕ_i are holomorphic, $\mathcal{W}(\vartheta)$ is usually dense on $\widetilde{\Sigma}$.
- When the ϕ_i have poles, sometimes $\mathcal{W}(\vartheta)$ is not dense (e.g. this happens when $N = 2$, ϕ_2 has at least one double pole, and ϑ is generic.)
- The *complement* of $\mathcal{W}(\vartheta)$ is always dense.

Some simple examples of $\mathcal{W}(\vartheta = 0)$ with various $\vec{\phi}$ and $N = 2, 3, 4$:

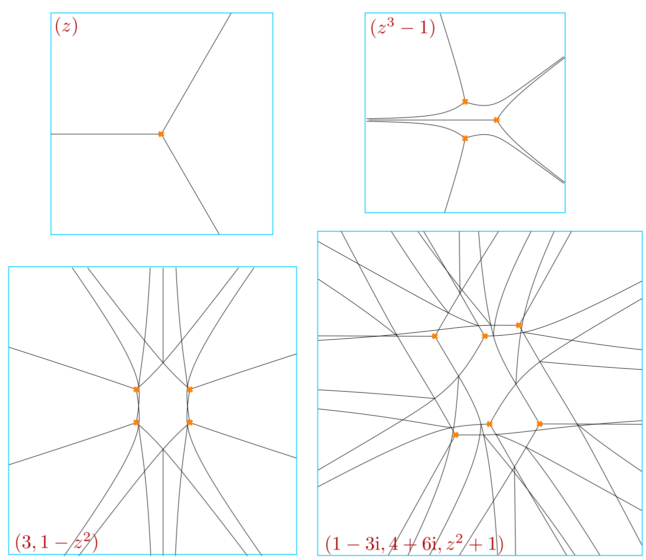

These examples are misleadingly simple: the generic behavior is probably that $\mathcal{W}(\vartheta)$ is dense.

Exercise 2.1. Prove that when $N = 2$ and $\phi_2 = f(z)\mathrm{d}z^2$, with f a polynomial of degree d, the projection of $\mathcal{W}(\vartheta)$ to the base $\mathbb{C}$ consists of curves which are asymptotic to $d + 2$ asymptotic directions in the plane. How do those directions depend on ϑ?

Question 2.1. Experimentally, $\mathcal{W}(\vartheta)$ is finite in simple enough cases. Can it be proven? For example, suppose we take $N = 3$, some integer d, and $\vec{\phi} = (\phi_2, \phi_3)$ where $\phi_2 = f(z)\mathrm{d}z^2$,

$\phi_3 = (z^d + g(z))\mathrm{d}z^3$, with $\deg g < d$ and $\deg f < \frac{2d}{3}$. Then, relations to cluster algebras lead to the wild guess that for $d \leq 5$ the network $\mathcal{W}(\vartheta)$ only involves finitely many trajectories, while for $d \geq 6$ it could involve infinitely many. Numerical experiments suggest the finiteness is really true at least for $d = 2, 3$. How to prove it?

2.6. **BPS phases.** Suppose that, for some ϑ_c, $\mathcal{W}(\vartheta_c)$ contains a leaf which runs into one of the μ-fixed singularities (or, in our local notation on C, an ij-leaf runs into a branch point labeled (ij)). In this case we say that ϑ is a *BPS phase*. These phases are especially important; in particular, $\mathcal{W}(\vartheta)$ changes discontinuously as ϑ moves across a BPS phase ϑ_c. A related point is that the BPS phases are exactly the phases for which the projections of two walls can overlap on a whole segment.

Example 2.16 (BPS phases for $N = 2$). Here is an example of a BPS phase for $N = 2$.

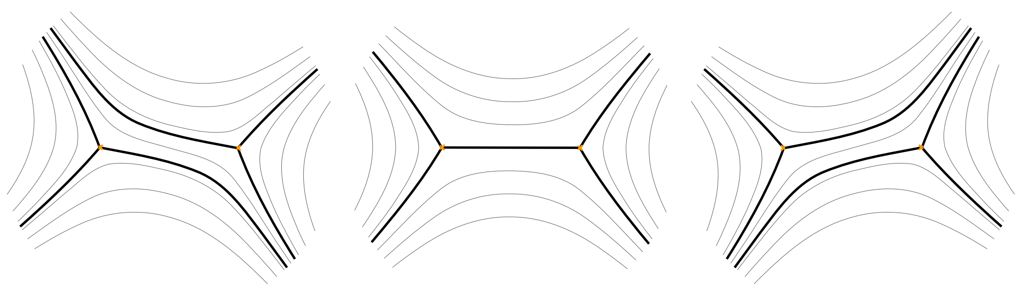

$\vartheta < \vartheta_c$ $\qquad$ $\vartheta = \vartheta_c$ $\qquad$ $\vartheta > \vartheta_c$

The segment appearing in the middle is called a *saddle connection*.

Example 2.17 (BPS phases for $N = 3$). Here is an example of a BPS phase for $N = 3$. In this case the finite object which appears is a three-stringed junction.

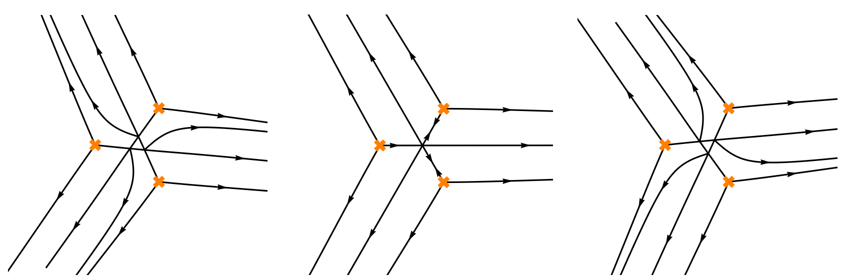

$\vartheta < \vartheta_c$ $\qquad$ $\vartheta = \vartheta_c$ $\qquad$ $\vartheta > \vartheta_c$

3. Exact WKB for opers

Now we consider one place where the WKB spectral network occurs "in nature." [Voros, Ecalle, . . . , Gaiotto-Moore-N, Iwaki-Nakanishi, Hollands-N] Many of the needed analytical statements are still conjectures, but there is considerable numerical evidence that things work as claimed, at least in examples simple enough to investigate.

3.1. **Data.** We take the same data as before:

- A compact Riemann surface C.
- An integer $N \geq 1$.
- A tuple $\vec{\phi} = (\phi_1, \ldots, \phi_N)$ where ϕ_i is a meromorphic section of $K_C^{\otimes i}$, with $\phi_1 = 0$.

In addition we take two extras:

- A spin structure on C (only necessary for N even).
- A complex projective structure on C.

3.2. **Opers.** Given these data, there is a canonical construction of an order N differential operator (oper) acting on meromorphic sections,

$$\mathcal{D}_{\vec{\phi}} : K_C^{\frac{1-N}{2}} \to K_C^{\frac{1+N}{2}}. \tag{3.1}$$

Example 3.1 (Opers for $N = 2$). Suppose $N = 2$. Choose a local coordinate patch on C, lying in the atlas given by the complex projective structure. Let z be the coordinate. Trivialize $K_C^{\frac{1}{2}}$ by one of the two choices of $\sqrt{\mathrm{d}z}$, similarly trivialize all powers $K_C^{\frac{k}{2}}$ by $\sqrt{\mathrm{d}z}^k$, and write $\phi_2 = P_2(z)\mathrm{d}z^2$. Relative to these trivializations, $\mathcal{D}_{\vec{\phi}}$ is given by

$$\mathcal{D}_{\vec{\phi}}\psi = (\partial_z^2 + P_2)\psi. \tag{3.2}$$

(So it looks like a meromorphic version of a Schrödinger operator.)

Example 3.2 (Opers for $N = 3$). Suppose $N = 3$ and we choose a patch and trivializations as above. Relative to these trivializations, $\mathcal{D}_{\vec{\phi}}$ is given by[1]

$$\mathcal{D}_{\vec{\phi}}\psi = \left(\partial_z^3 + P_2\partial_z + \frac{1}{2}\partial_z P_2 + P_3\right)\psi. \tag{3.3}$$

More generally we introduce an extra parameter $\hbar \in \mathbb{C}^\times$, and consider the rescaling $\hbar^{-1}\vec{\phi} = (\hbar^{-2}\phi_2, \hbar^{-3}\phi_3, \dots, \hbar^{-N}\phi_N)$. Then we let $\mathcal{D}^\hbar_{\vec{\phi}} = \mathcal{D}_{\hbar^{-1}\vec{\phi}}$. So e.g. (3.2) becomes

$$\mathcal{D}^\hbar_{\vec{\phi}}\psi = (\partial_z^2 + \hbar^{-2}P_2)\psi. \tag{3.4}$$

Now the kernel of $\mathcal{D}^\hbar_{\vec{\phi}}$ acting on holomorphic sections is a rank N local system $L^\hbar_{\vec{\phi}}$ over the locus $C' \subset C$ where $\vec{\phi}$ is holomorphic. It comes with some additional structure: namely, along any ray going into a singularity of $\vec{\phi}$, there is a filtration by the rates of growth of the sections. The filtration depends on the ray in a piecewise-constant fashion: said otherwise, there is a circle of asymptotic directions near the singularity, and this circle is divided into arcs, with a complete filtration along each arc. We call $L^\hbar_{\vec{\phi}}$, with this filtration data, a *decorated* local system. We won't develop this systematically here but let us discuss two key examples.

Example 3.3 (Polynomial differentials for $N = 2$). Set $N = 2$, $C = \mathbb{CP}^1$, and $\phi_2 = (z^3 - 1)\mathrm{d}z^2$. Then $C' = \mathbb{C}$, and the space of global sections of the local system $L^\hbar_{\vec{\phi}}$ over C' is a 2-dimensional vector space V.

In this case the circle around $z = \infty$ is divided into 5 arcs, as shown below when $\arg \hbar = 0$. (The blue lines mark the boundaries between arcs; the middle of each arc is marked with a red dot.)

[1] The extra term $\frac{1}{2}\partial_z\phi_2$ might be a little surprising; one quick way to see that it needs to be there is to consider the covariance of this expression under Möbius transformations. It also arises from naive quantization of the curve Σ using the symmetric ordering prescription.

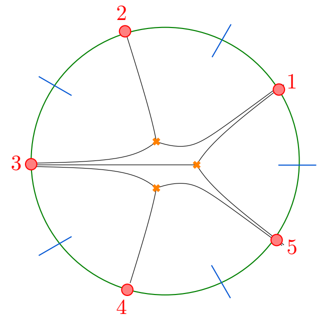

Thus V gets 5 complete flags, which just means 5 distinct lines $\ell_1, \ldots, \ell_5 \subset V$. ℓ_i consists of those sections of $L^\hbar_{\vec{\phi}}$ which decay exponentially as $z \to \infty$ in the i-th arc. Thus the isomorphism class of the decorated local system $L^\hbar_{\vec{\phi}}$ is a point of $\mathrm{Conf}_5(\mathbb{CP}^1)$ (note, this $\mathbb{CP}^1 \simeq \mathbb{P}(V)$ has nothing to do with C!)

Example 3.4 (Polynomial differentials more generally). In a similar way, if we take $\phi_2 = P_2(z)\mathrm{d}z^2$ where P_2 is a polynomial of degree d, we get $d+2$ asymptotic arcs and a point of $\mathrm{Conf}_{d+2}(\mathbb{CP}^1)$.

More generally, when $C = \mathbb{CP}^1$ and $\vec{\phi} = (\phi_2, \ldots, \phi_N)$ consists of sufficiently generic polynomial differentials, the pattern of filtrations that one gets around $z = \infty$ is more complicated. If $\vec{\phi}$ is a sufficiently low-degree perturbation of $(0, 0, \ldots, 0, z^K \mathrm{d}z^N)$ then I hope that $L^\hbar_{\vec{\phi}}$ is equivalent to a point of $\mathrm{Gr}(N, K+N)/(\mathbb{C}^\times)^{K+N}$. The case above was the case $N = 2$, $K = d$, where $\mathrm{Gr}(2, d+2)/(\mathbb{C}^\times)^{d+2} \simeq \mathrm{Conf}_{d+2}(\mathbb{CP}^1)$.

3.3. **The Riemann-Hilbert map.** The map $\vec{\phi} \mapsto L_{\vec{\phi}}$ is called the *Riemann-Hilbert map*. We are interested in computing this map practically.

As a practical matter this could mean computing the trace of holonomy around some loop, or some invariants of the filtrations around irregular singularities. Here is one concrete example:

Example 3.5 (Cross-ratio question for polynomial differentials). Say $N = 2$, $C = \mathbb{CP}^1$, $\phi_2 = (z^3 - 1)\mathrm{d}z^2$ as in Example 3.3. Then we ask: what is the corresponding point $L^\hbar_{\vec{\phi}} \in \mathrm{Conf}_5(\mathbb{CP}^1)$? Concretely, we choose a coordinate system on $\mathrm{Conf}_5(\mathbb{CP}^1)$ and ask, what are the invariant cross-ratios

$$r(\ell_1, \ell_2, \ell_3, \ell_5), \qquad r(\ell_5, \ell_1, \ell_3, \ell_4)\,? \tag{3.5}$$

3.4. **Formal WKB solutions.** How do we study the Riemann-Hilbert map? We need a practical understanding of the solutions of $\mathcal{D}_{\vec{\phi}}\psi = 0$. The basic idea: consider

$$\mathcal{D}^\hbar_{\vec{\phi}}\psi = 0 \tag{3.6}$$

instead, and expand in series around $\hbar = 0$. We make the following ansatz:

$$\psi^{(y)}_{\mathrm{formal}}(z, \hbar) = \exp\left(\hbar^{-1} \int_{y_0}^{y} \tilde{\lambda}_{\mathrm{formal}}(\hbar)\right) \tag{3.7}$$

where

- $y \in \pi^{-1}(z) \subset \Sigma_{\vec{\phi}}$,

- $\tilde{\lambda}_{\text{formal}}(\hbar)$ is a 1-form on $\Sigma_{\vec{\phi}}$, written as a formal power series in $\hbar$,

$$\tilde{\lambda}_{\text{formal}}(\hbar) = \sum_{n=0}^{\infty} \hbar^n \lambda_n \tag{3.8}$$

  with each λ_n a 1-form on $\Sigma_{\vec{\phi}}$, and $\lambda_0 = \lambda$ the original Liouville form.

The equation (3.6), expanded in powers of $\hbar$, determines the λ_n recursively.

Example 3.6 (WKB ansatz when $N = 2$). In the case $N = 2$, the equation we are solving is (3.4). Then the leading-order term in our ansatz is

$$\psi(z, \hbar) \sim \exp\left(\hbar^{-1} \int^z \sqrt{-P_2}\, \mathrm{d}z\right) \tag{3.9}$$

which indeed solves Equation 3.4 to leading order. Then plugging in the full ansatz (3.7) gives[2]

$$\tilde{\lambda}^2 + \pi^* \phi_2 + \hbar \partial \tilde{\lambda} = 0, \tag{3.10}$$

which determines

$$\tilde{\lambda}_{\text{formal}}(\hbar) = \left(\sqrt{-P_2} - \hbar \frac{P_2'}{4P_2} + \hbar^2 \sqrt{-P_2} \frac{5P_2'^2 - 4P_2 P_2''}{32 P_2^3} + \cdots\right) \pi^* \mathrm{d}z. \tag{3.11}$$

(Note $\sqrt{-P_2}$ is single-valued on Σ.)

Exercise 3.1. Derive the series (3.11) up to the order shown. Write it out more explicitly in the case where $P_2(z) = z$. (In this case the equation (3.4) is the Airy equation, and the desired solutions are the Airy and Biry functions.)

Formally, this works for each $y \in \pi^{-1}(z)$, and thus gives N solutions in any simply connected domain.

3.5. **Borel summation.** The crucial fact about (3.8) is that (whenever there exists at least one branch point) it necessarily has zero radius of convergence. Nevertheless, one can try to make it into an honest 1-form, by the procedure of Borel summation. The key claim is that this can actually be done, *but only away from the spectral network*:

Conjecture 3.7 (Borel summability of WKB series). The series $\tilde{\lambda}_{\text{formal}}(y, \hbar)$ is Borel summable to an actual 1-form $\tilde{\lambda}(y, \hbar)$, except when $y \in p_1(\mathcal{W}(\vartheta = \arg \hbar))$.

Proposition 3.8 (Borel summability when $N = 2$ with enough poles). Suppose $N = 2$, ϕ_2 has at least one double pole, and no residue m^2 of ϕ_2 has $\arg m = \vartheta$. Then Conjecture 3.7 is true. [Koike-Schafke, Takei, Allegretti, Nikolaev, ...]

Question 3.1. Can we establish Conjecture 3.7 more generally?

(From now on let's reduce the notation by writing $\mathcal{W} = \mathcal{W}(\vartheta = \arg \hbar)$, $L = L^{\hbar}_{\vec{\phi}}$, $\Sigma = \Sigma_{\vec{\phi}}$.)

[2] In the equation below $\partial \tilde{\lambda}$ is a quadratic differential on Σ, whose definition involves the complex projective structure on C; in a coordinate on C in the atlas determined by the projective structure, we would write it as $\partial \tilde{\lambda} = \pi^*(\partial_z(\pi_* \tilde{\lambda}/\mathrm{d}z)\mathrm{d}z^2)$.

For $z \notin \pi(\mathcal{W})$ this works for each $y \in \pi^{-1}(z)$, and thus the ansatz

$$\psi^{(y)}(z,\hbar) = \exp\left(\hbar^{-1}\int_{y_0}^{y}\tilde{\lambda}(\hbar)\right) \tag{3.12}$$

gives N distinct solutions in any simply connected domain away from $\pi(\mathcal{W})$. These solutions depend on the basepoints y_0 we choose, but changing y_0 just changes them by a scalar multiple.

Conjecture 3.9 (Linear independence of solutions given by WKB ansatz). The solutions $\psi^{(y)}(z,\hbar)$ for $y \in \pi^{-1}(z)$ form a linearly independent set.

Thus they give a decomposition of the N-dimensional space of solutions into N distinguished lines. Said more geometrically: we have a local system A over $\Sigma \setminus p_1(\mathcal{W})$, whose sections are the solutions of the equation

$$(\mathrm{d} - \hbar^{-1}\tilde{\lambda})\psi = 0. \tag{3.13}$$

Then what we found is that

$$L\,|_{C\setminus\pi(\mathcal{W})} = \pi_* A|_{C\setminus\pi(\mathcal{W})}. \tag{3.14}$$

This is like saying that one can canonically *diagonalize* the local system L away from $\pi(\mathcal{W})$.

Moreover, suppose $\vartheta = \arg\hbar$ is not a BPS phase; then there is a canonical way of patching A across $p_1(\mathcal{W})$ to get a rank 1 local system on $\Sigma \setminus \pi^{-1}(\Delta_C)$, with holonomy -1 around each point of $\pi^{-1}(\Delta_C)$. We call this an almost-local system over Σ. Then, for a loop $\gamma \in H_1(\Sigma \setminus \pi^{-1}(\Delta_C), \mathbb{Z})$ we define

$$X_\gamma = \mathrm{Hol}_\gamma A. \tag{3.15}$$

The holonomy X_γ is a function we can say something about:

Conjecture 3.10 (Asymptotic expansion of abelian holonomies). If $\vartheta = \arg\hbar$ is not a BPS phase, then the series

$$\tilde{Z}_\gamma^{\mathrm{formal}}(\hbar) = \sum_{n=0}^{\infty}\hbar^n \oint_\gamma \tilde{\lambda}_n \tag{3.16}$$

is Borel summable to a function $\tilde{Z}_\gamma(\hbar)$, and

$$X_\gamma = \exp(\tilde{Z}_\gamma/\hbar). \tag{3.17}$$

Moreover, $\oint_\gamma \tilde{\lambda}_1 \in \pi\mathrm{i}\mathbb{Z}$. In particular,

$$X_\gamma \sim \pm\exp(Z_\gamma/\hbar) \tag{3.18}$$

as $\hbar \to 0$.

Moreover the X_γ are useful in practice:

Example 3.11 (Abelian holonomies as cross-ratios). Return to Example 3.5. The two cross-ratios can be expressed as

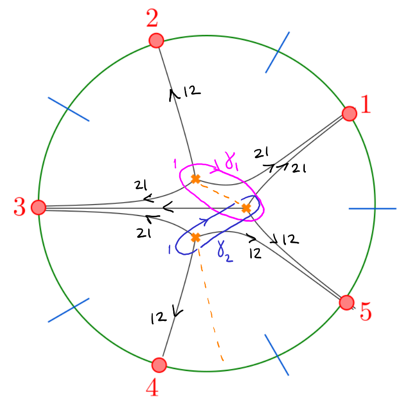

$$r(\ell_1, \ell_2, \ell_3, \ell_5) = X_{\gamma_1}, \qquad r(\ell_5, \ell_1, \ell_3, \ell_4) = X_{\gamma_2}. \tag{3.19}$$

Cluster aficionados will recognize this picture: according to [Fock-Goncharov] one should associate a "cluster" coordinate system on $\mathrm{Conf}_5(\mathbb{CP}^1)$ to every triangulation of the pentagon. The cross-ratios $(X_{\gamma_1}, X_{\gamma_2})$ are the cluster coordinates associated with the triangulation determined by $\mathcal{W}$:

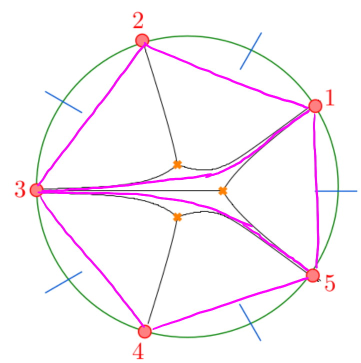

So what we have seen is that these coordinates are in some sense the natural coordinates for analysis of opers. Indeed this turns out to be true more generally:

Proposition 3.12 (Abelian holonomies as cluster coordinates). Suppose $N = 2$, ϕ_2 has at least one pole of order ≥ 2, ϕ_2 has only simple zeroes, and $\vartheta = \arg \hbar$ is not a BPS phase. Then the spectral network $\mathcal{W}$ determines an ideal triangulation of a blown-up version of C, as we discussed above. For an internal edge E of the triangulation, there is a corresponding cycle γ_E shown in the figure. Then the abelian holonomy X_{γ_E} is the Fock-Goncharov coordinate $X_E(L)$. (As above, it can be expressed as a cross-ratio between distinguished sections associated to the four vertices of the quadrilateral.)

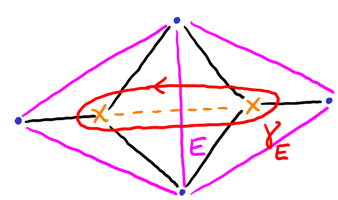

We expect a similar relation between the X_γ and cluster coordinates for $N > 2$ at least in sufficiently nice (finite) cases. This kind of relation is known in some examples. For instance one can recover cluster coordinates on $\mathrm{Gr}(3,6)$ this way, or on moduli spaces of flat $SL(N)$-connections with complete flags on the boundary, a la Fock-Goncharov. [refs]

4. Path lifting

Where are we now? We started with:

- A compact Riemann surface C,
- An integer $N \geq 1$,
- A tuple $\vec{\phi} = (\phi_2, \ldots, \phi_N)$ where ϕ_i is a meromorphic section of $K_C^{\otimes i}$,

and

- A spin structure on C (only necessary for N even),
- A complex projective structure on C.

From these data we built:

- A rank N local system $L = L^\hbar_{\vec{\phi}}$ over C,
- A spectral network $\mathcal{W} = \mathcal{W}(\vartheta, \vec{\phi})$ on C,
- (Depending on conjectures in some cases), a rank 1 local system $A = A^\hbar_{\vec{\phi}}$ over the spectral cover $\Sigma = \Sigma_{\vec{\phi}}$.

The idea was that L is something difficult we want to study (eg a Schrödinger equation), and we simplify by replacing it with A which is easier to study.

How can we describe the relation between them directly, without referring to WKB, Borel summation etc?

At least when $\mathcal{W}$ is finite, the full local system L can be expressed in a precise way in terms of A, as we now describe. What is involved is a way of lifting paths from C to Σ. Naively lifting paths would not be homotopy invariant: when we perturb a path across a branch point, we change the homotopy classes of its lifts.

But there is a way of repairing this problem using the extra data of the network $\mathcal{W}$: this is the path-lifting rule.

To formulate it we need a little notation:

Definition 4.1 (Path categories). For $Y \subset X$:

- Let $\mathrm{Path}(X, Y)$ be the category of paths in X with endpoints in Y, enriched over abelian groups (so $\mathrm{Hom}(y, y')$ is the abelian group of formal $\mathbb{Z}$-linear combinations of paths from y to y' in X, and composition is extended linearly).
- Given a map $\pi : Z \to X$, let $\mathrm{Path}^Z(X, Y)$ be the category of paths in Z with endpoints in $\pi^{-1}(Y)$, where $\mathrm{Hom}(y, y')$ is the abelian group of formal $\mathbb{Z}$-linear combinations of paths from any $z \in \pi^{-1}(y)$ to any $z' \in \pi^{-1}(y')$ in Z, with composition taken to be zero for paths which don't concatenate.

Proposition 4.2 (WKB path-lifting rule). If $\mathcal{W}$ is finite, then there exists a "path-lifting" functor

$$F : \mathrm{Path}(C, C \setminus \pi(\mathcal{W})) \to \mathrm{Path}^{\Sigma}(C, C \setminus \pi(\mathcal{W})) \tag{4.1}$$

such that:

- F is almost-homotopy-invariant, i.e. if $\mathcal{P} \sim \mathcal{P}'$ then $F(\mathcal{P}) \sim^{al} F(\mathcal{P}')$, where $\sim^{al}$ means ordinary homotopy except that if we move a path across a point of $\pi^{-1}(\Delta_C)$ we get a sign -1.
- For any path $\mathcal{P}$, all terms in $F(\mathcal{P})$ are obtained by splicing paths $p(c)$, where c is a WKB soliton with phase ϑ, into lifts of $\mathcal{P}$.

We don't define "splicing" carefully here, but we do indicate the most fundamental and important example:

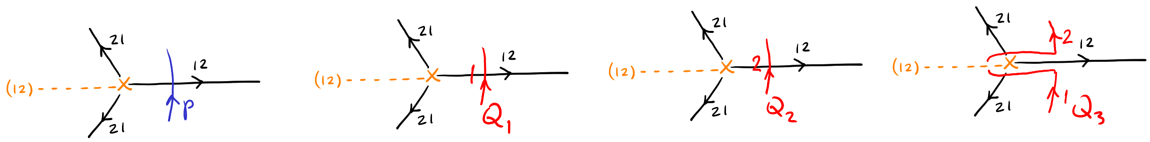

In this example $F(\mathcal{P}) = Q_1 + Q_2 + Q_3$.

Exercise 4.1. Check that for the path $\mathcal{P}'$ shown below, which has $\mathcal{P} \sim \mathcal{P}'$, indeed $F(\mathcal{P}) \sim^{al} F(\mathcal{P}')$. (Note this is a purely topological exercise.)

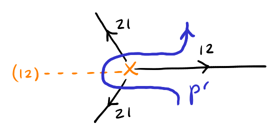

Here is another example of splicing:

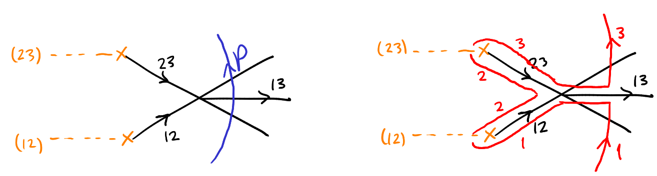

(This one helps to explain the need for the birth of extra half-leaves at intersection points; indeed, the homotopy invariance could not have worked without them.)

Question 4.1. Can we formulate a version of the path-lifting rule which works without the finiteness assumption?

The path-lifting functor is unique when ϑ is not a BPS phase. Otherwise some more care is needed with the definition of splicing, and there can be multiple choices; in particular one gets two different ones by taking the limits $\vartheta \to \vartheta^{\pm}$.

4.1. **Path lifting vs WKB.**

Definition 4.3 (Nonabelianization map). Given an almost-local system A over Σ, we define a local system $L = Nab_{\mathcal{W}}(A)$ over $C \setminus \pi(\mathcal{W})$ as follows: for an open set U, $L(U) = A(\pi^{-1}(U))$; for a path $\mathcal{P}$ from U to U', the map $L(\mathcal{P}) : L(U) \to L(U')$ is $A(F(\mathcal{P}))$.

This is a purely topological operation, and for a fixed $\mathcal{W}$ it is often easy to define and study. But it agrees with what we get from WKB analysis:

Proposition 4.4 (WKB and nonabelianization). The $A = A^{\hbar}_{\vec{\phi}}$ and $L = L^{\hbar}_{\vec{\phi}}$ which we considered above are related by $L = Nab_{\mathcal{W}}(A)$ where $\mathcal{W} = \mathcal{W}(\vartheta = \arg \hbar, \vec{\phi})$.

Corollary 4.5 (Expansions of traces of holonomies). Suppose $\mathcal{P}$ is a closed loop on C. Then

$$\operatorname{Tr} L(\mathcal{P}) = \sum_{\gamma \in H_1(\Sigma,\mathbb{Z})} \overline{\underline{\Omega}}(\mathcal{P},\gamma) X_\gamma \tag{4.2}$$

for some (computable) $\overline{\underline{\Omega}}(\mathcal{P},\gamma) \in \mathbb{Z}$.

In particular, since we know the asymptotic expansion of X_γ as $\hbar \to 0$, this allows us to determine the asymptotic expansion of $\operatorname{Tr} L(\mathcal{P})$ as $\hbar \to 0$.

5. Donaldson-Thomas-type invariants

5.1. **Chamber structure.** We have been discussing a decorated $GL(N)$-local system $L^{\hbar}_{\vec{\phi}}$ over C and a $GL(1)$-almost-local system $A^{\hbar}_{\vec{\phi}}$ over Σ, related by $L^{\hbar}_{\vec{\phi}} = Nab^{\hbar}_{\vec{\phi}}(A^{\hbar}_{\vec{\phi}})$. Working modulo equivalence we get the corresponding map on moduli spaces of decorated local systems, of the form

$$Nab^{\hbar}_{\vec{\phi}} : \mathcal{M}^{tw}(\Sigma_{\vec{\phi}}, GL(1)) \to \mathcal{M}(C, GL(N)).$$

$Nab^{\hbar}_{\vec{\phi}}$ depends only on the topology of the spectral network $\mathcal{W}(\vartheta = \arg \hbar, \vec{\phi})$.

Conjecture 5.1. The map $Nab^{\hbar}_{\vec{\phi}}$ is a local symplectomorphism, with respect to the natural (Atiyah-Bott) structures on the two moduli spaces.

This is true for $N = 2$, just because we already know the Fock-Goncharov coordinates are Darboux coordinates. There is an idea of proof in the spectral networks paper; I believe [Morrissey] working on real proof.

It follows that $F^{\hbar}_{\vec{\phi}}$ is at least locally invertible, to give a map

$$Ab^{\hbar}_{\vec{\phi}} : \mathcal{M}(C, GL(N)) \to \mathcal{M}^{tw}(\Sigma_{\vec{\phi}}, GL(1)) \simeq (\mathbb{C}^\times)^n \tag{5.1}$$

We think of $Ab^{\hbar}_{\vec{\phi}}$ as giving a local coordinate system on $\mathcal{M}(C, GL(N))$. We discussed some of these coordinate systems earlier (cluster cooordinates).

So we have a decomposition of the parameter-space of $(\vartheta, \vec{\phi})$ into chambers. Each chamber is labeled by a topology of the spectral network $\mathcal{W}(\vartheta, \vec{\phi})$. The walls are the locus where ϑ is a BPS phase. Each wall w is labeled by some birational automorphism a_w of the torus $\mathcal{M}^{tw}(\Sigma_{\vec{\phi}}, GL(1))$, determined by the relation

$$G^+ = a_w \circ G^- \tag{5.2}$$

where $G^\pm$ are the coordinate systems on the two sides of the wall.

The walls may be dense in some part of the parameter space, but as usual there are some nice enough examples where they are not.

Example 5.2 (Cubic $N = 2$ example). If we take $C = \mathbb{CP}^1$, $\phi_2 = z^3 + z + u$ and let u vary with fixed small imaginary part, we get a picture with the topology below.

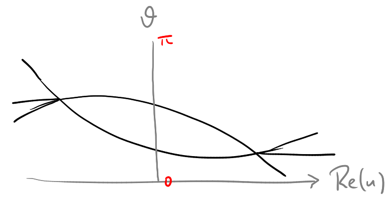

5.2. **Automorphisms.** What are the automorphisms of $\mathcal{M}^{tw}(\Sigma_{\vec{\phi}}, GL(1))$ attached to the walls?

Example 5.3 (BPS automorphisms from saddle connections when $N = 2$). In case $N = 2$, the simplest kind of wall w arises from a *saddle connection* for the quadratic differential ϕ_2, i.e. a geodesic in the metric $|\phi_2|$ connecting two zeroes. Such a saddle connection appears at the phase where the spectral network jumps, as in the figure:

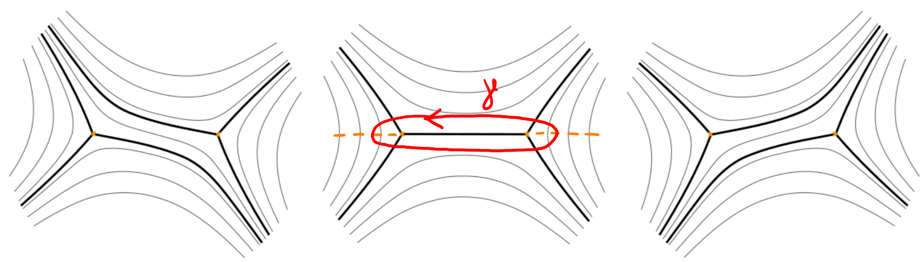

$\vartheta < \vartheta_c$ $\vartheta = \vartheta_c$ $\vartheta > \vartheta_c$

Let the *charge* of the saddle connection be the cycle $\gamma \in H_1(\Sigma, \mathbb{Z})$ in the figure. Note

$$\vartheta_c = \arg Z_\gamma. \tag{5.3}$$

This jump of the spectral network corresponds to the automorphism [Kontsevich-Soibelman]

$$a_w = K_\gamma : X_\mu \to X_\mu (1 - X_\gamma)^{\langle \mu, \gamma \rangle} \tag{5.4}$$

This is also (essentially) an example of a cluster transformation. [Fock-Goncharov]

Example 5.4 (BPS automorphisms from trees). When $N > 2$ we can have saddle connections but also trees with more interesting topologies, e.g.:

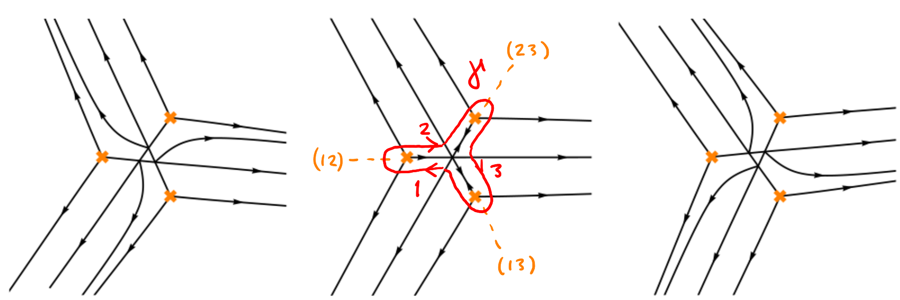

$\vartheta < \vartheta_c$ $\vartheta = \vartheta_c$ $\vartheta > \vartheta_c$

A tree gives the same kind of automorphism as a saddle connection,

$$a_w = K_\gamma : X_\mu \to X_\mu (1 - X_\gamma)^{\langle \mu, \gamma \rangle} \tag{5.5}$$

Example 5.5 (BPS automorphism from ring domain). When $N = 2$, a wall associated to a ring domain gives an automorphism

$$a_w = K_\gamma^{-2} : X_\mu \to X_\mu (1 - X_\gamma)^{-2\langle \mu, \gamma \rangle} \tag{5.6}$$

This is not a cluster transformation, but it plays the same role in the theory as an honest cluster transformation.

When $N > 2$ we can have more interesting objects with a loop attached, which still give the automorphism above. But we can also have more complicated topologies with multiple loops, which can lead to more complicated automorphisms.

5.3. **BPS invariants.** Whatever the automorphisms are, we organize them into a collection of numbers, as follows.

Definition 5.6 (BPS invariants). Fix $\vec{\phi}$ such that $Z_\gamma \parallel Z_{\gamma'} \implies \gamma \parallel \gamma'$. Then for each wall w at $\vec{\phi}$ let a_w be the attached automorphism, and define integers $\Omega(\gamma; \vec{\phi})$ by

$$a_w = \prod_{\gamma \in \Gamma_w} K_\gamma^{\Omega(\gamma; \vec{\phi})} \tag{5.7}$$

For γ which do not appear in a_w for any wall w at $\vec{\phi}$, we set $\Omega(\gamma; \vec{\phi}) = 0$.

There is an algorithm (described in the spectral network paper) for computing $\Omega(\gamma; \vec{\phi})$ from the topology of the spectral network.

Conjecture 5.7 (Integrality of BPS invariants). All $\Omega(\gamma; \vec{\phi}) \in \mathbb{Z}$.

Example 5.8 (BPS invariants when $N = 2$). When $N = 2$,

$$\Omega(\gamma; \phi_2) = \#\{\phi_2 \text{ saddle connections with charge } \gamma\} - 2\#\{\phi_2 \text{ ring domains with charge } \gamma\} \tag{5.8}$$

Example 5.9 (BPS invariants in the cubic $N = 2$ example). In the cubic $N = 2$ example from above, holding $\mathrm{Im}(u)$ fixed and varying $\mathrm{Re}(u)$, we have:

- for small $\mathrm{Re}(u)$, $\Omega(\pm\gamma_1) = 1$, $\Omega(\pm\gamma_2) = 1$, all others zero;
- for large $\mathrm{Re}(u)$, $\Omega(\pm\gamma_1) = 1$, $\Omega(\pm\gamma_2) = 1$, $\Omega(\pm(\gamma_1 + \gamma_2)) = 1$, all others zero.

If we let u vary in the whole plane we have a picture like:

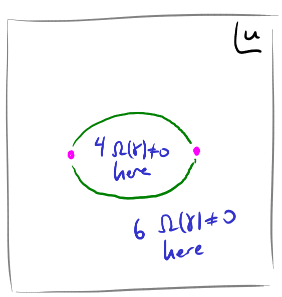

These invariants are counting the saddle connections; so a saddle connection can appear or disappear when we vary u. This is the *wall-crossing* phenomenon.

Note our picture implies some constraints on the wall-crossing phenomenon: given two homotopic paths in the parameter space, the associated automorphisms have to be the same. For example, applying this in the $N = 2$ cubic example to the two paths indicated below

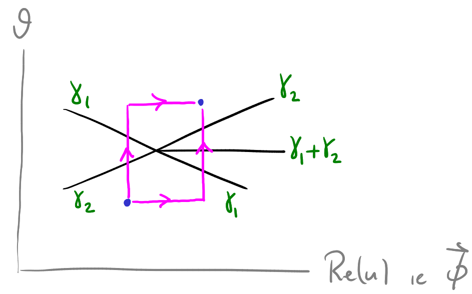

requires the identity

$$K_{\gamma_1} K_{\gamma_2} = K_{\gamma_2} K_{\gamma_1+\gamma_2} K_{\gamma_1} \tag{5.9}$$

which is indeed true (it just uses the fact that $\langle \gamma_1, \gamma_2 \rangle = 1$). In principle, if we know all the $\Omega(\gamma; \vec{\phi})$ for some $\vec{\phi}$, we can use this type of constraint to determine all the $\Omega(\gamma; \vec{\phi}')$. Said otherwise, we expect the $\Omega(\gamma; \vec{\phi})$ to satisfy the Kontsevich-Soibelman wall-crossing formula.

The Kontsevich-Soibelman formula was originally written in the context of DT theory. The following theorem says this is not just a coincidence.

Theorem 5.10 (BPS invariants are DT invariants when $N = 2$). (Roughly) [Bridgeland-Smith] Fix a topological surface S with marked points p_i, labeled by integers n_i, with at least one $n_i \geq 2$. Then there is a triangulated 3CY category $\mathcal{C}$ such that:

- for each (complex structure on S and meromorphic ϕ_2, with a pole of order n_i at p_i for each i), there is a corresponding stability condition on $\mathcal{C}$,
- the BPS invariants $\Omega(\gamma; \phi_2)$ are the generalized Donaldson-Thomas invariants associated to this category and stability condition.

Very recently it was shown that the same theorem is also true without the marked points. [Haiden]

The invariants $\Omega(\gamma; \vec{\phi})$ in case $N > 2$ are much less explored — though some examples computed [N-Galakhov-Longhi-Mainiero-Moore, N-Hollands, N-Hao-Hollands]. In particular, they should be DT invariants as in $N = 2$ but this is not known yet. Unlike $N = 2$ they can have exponential growth as function of $\|\gamma\|$.

5.4. **Riemann-Hilbert problem.** Now let's return to the family of $GL(N)$-local systems $L^{\hbar}_{\vec{\phi}}$ obtained from opers, abelianized by $GL(1)$-local systems $A^{\hbar}_{\vec{\phi}}$, with holonomy functions $X_\gamma(\hbar)$. From what we have said we expect the following properties for $X_\gamma(\hbar)$:

- $X_\gamma(\hbar) \sim \exp(Z_\gamma/\hbar)$ as $\hbar \to 0$.
- $X_\gamma(\hbar)$ is piecewise analytic, with jumps by the automorphisms a_w at the walls w at $\vec{\phi}$.

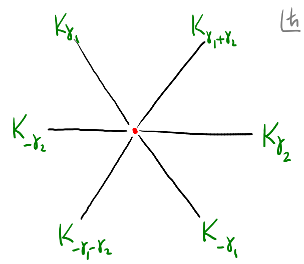

We might also optimistically hope for the following, which is true in some interesting examples:

- $X_\gamma(\hbar)$ has "moderate growth" as $\hbar \to \infty$.

These conditions can be thought of as defining a Riemann-Hilbert problem: given just the Z_γ and the automorphisms a_w, can we find $X_\gamma(\hbar)$? This problem has been studied recently [GMN, Gaiotto, Dumas-Neitzke, Bridgeland, ...] and in simple experiments (mostly with $N = 2$ and $N = 3$) it seems that the answer is yes; so this gives a way of computing the monodromy of $L^\hbar_{\vec{\phi}}$ just in terms of the periods Z_γ and the automorphisms a_w, i.e. the DT invariants.

E.g. the figure below from [Dumas-Neitzke] in an $N = 3$ case, where we took $C = \mathbb{CP}^1$, $\phi_2 = 0$, $\phi_3 = \frac{1}{2}(z^3 - 3z^2 - 2)\mathrm{d}z^2$, $\arg \hbar = 0.1$. We computed by integrating the ODE directly vs. solving the Riemann-Hilbert problem and compared the results, finding good agreement.

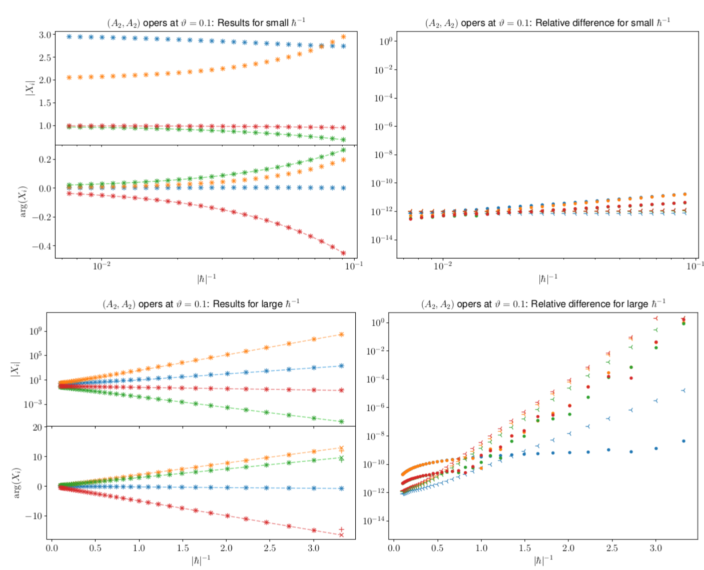

(In this case the relevant moduli space of decorated local systems is $Gr(3,6)/(\mathbb{C}^\times)^6$, and the quantities X_γ plotted here can be thought of as cluster coordinates on the curve $\hbar \mapsto [L^\hbar_{\vec{\phi}}] \in Gr(3,6)/(\mathbb{C}^\times)^6$.)

6. Hyperkähler metrics

Finally we very briefly consider the moduli space of Higgs bundles. So, fix a compact Riemann surface C. Then there is an associated moduli space

$$\mathcal{M}^H(C, GL(N)) = \{(E,\varphi)\,\text{stable} : E \text{ holomorphic } GL(N)\text{-bundle over } C, \varphi \in H^0(\operatorname{End} E \otimes K_C)\}/\sim \tag{6.1}$$

There is also a "ramified" version where C has marked points $z_1, \dots, z_k$, some asymptotic data around the punctures (we won't specify this precisely), E is a parabolic bundle and φ is meromorphic. In either case $\mathcal{M}^H(C, GL(N))$ is known to be hyperkähler[Hitchin, Simpson, Biquard-Boalch] — this means it has a family of holomorphic symplectic structures (I^ζ, ϖ^ζ) labeled by $\zeta \in \mathbb{CP}^1$, obeying some relations.

One description of the holomorphic symplectic structure (I^ζ, ϖ^ζ): given a Higgs bundle (E, φ) there is an associated Hermitian metric h in E such that the family of connections

$$\nabla^\zeta_{(E,\varphi)} = \zeta^{-1}\varphi + D_h + \zeta\varphi^{\dagger_h} \tag{6.2}$$

is flat (here D_h denotes the Chern connection). To find this metric (and thus to find the connection $\nabla^\zeta_{(E,\varphi)}$) directly, one has have to solve an elliptic PDE on C. At any rate, $(E,\varphi) \mapsto \nabla^\zeta_{(E,\varphi)}$ gives a map

$$NAH^\zeta : \mathcal{M}^H(C, GL(N)) \to \mathcal{M}(C, GL(N)) \tag{6.3}$$

and then

$$\varpi^\zeta = (NAH^\zeta)^* \varpi_{AB} \tag{6.4}$$

with ϖ_{AB} the standard Atiyah-Bott form.

So to determine the hyperkähler structure it is enough to determine the connections $\nabla^\zeta_{(E,\varphi)}$ up to equivalence, or their corresponding local systems $L^\zeta_{(E,\varphi)}$. The idea: we can study this family *without* actually solving Hitchin's PDE, using WKB methods much like what we discussed for the family $L^\hbar_{\vec{\phi}}$. Indeed the behavior here as $\zeta \to 0$ is very similar to the behavior as $\hbar \to 0$ we had before.

Indeed let $\vec{\phi}$ be the coefficients in the characteristic polynomial of φ; then we hope for the same kind of structure as before, i.e. $L^\zeta_{(E,\varphi)}$ can be abelianized by a rank 1 almost-local system $A^\zeta_{(E,\varphi)}$ over the spectral curve $\Sigma_{\vec{\phi}}$. (We could also say this in the language of ζ-connections, and then as $\zeta \to 0$ this abelianization of ζ-connections should go over to the usual abelianization of Higgs bundles.) Then the $X_\gamma(\zeta)$ give Darboux coordinates for ϖ^ζ.

In particular, for fixed (E,φ) lying on the Hitchin section the functions $X_\gamma(\zeta)$ should obey a Riemann-Hilbert problem like the one we had above, but now with a symmetry between $\zeta = 0$ and $\zeta = \infty$:

- $X_\gamma(\zeta) \sim \exp(Z_\gamma/\zeta)$ as $\zeta \to 0$.
- $X_\gamma(\zeta)$ is piecewise analytic, with jumps by the automorphisms a_w at the walls w at $\vec{\phi}$.
- $X_\gamma(\zeta) \sim \exp(\overline{Z}_\gamma\zeta)$ as $\zeta \to \infty$.

Then solving this Riemann-Hilbert problem gives a computation of ϖ^ζ. This gives some interesting information about it (e.g. asymptotics).

This can also be tested numerically, eg the figure below from [Dumas-Neitzke] shows the case of Higgs bundles $\varphi = \begin{pmatrix} 0 & 1 \\ z^3 - c & 0 \end{pmatrix} \mathrm{d}z$, E trivialized, in the plane; restricted to this family one has Kähler metric $g(c)|\mathrm{d}c|^2$.

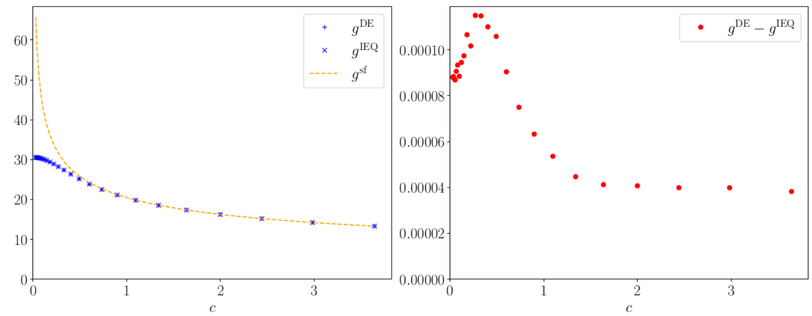

Figure 18. Left: The metric coefficient $g(c)$ for $\Lambda = 0$ and $c \in \mathbb{R}_+$. The blue marks show values of $g(c)$ computed using two methods: the direct PDE approach and the integral equations. The dashed line shows the semiflat approximation. Right: The absolute difference $g^{\mathrm{DE}} - g^{\mathrm{IEQ}}$.