Spectral Networks and Harmonic Maps to Buildings

3^{rd} Itzykson Colloquium
Fondation Mathématique Jacques Hadamard
IHES, Thursday 7 November 2013

C. Simpson, **joint work** with
Ludmil Katzarkov, Alexander Noll, and
Pranav Pandit
(in progress)

We wanted to understand the spectral networks of Gaiotto, Moore and Neitzke from the perspective of euclidean buildings. This should generalize the trees which show up in the SL_2 case. We hope that this can shed some light on the relationship between this picture and moduli spaces of stability conditions as in Kontsevich-Soibelman, Bridgeland-Smith, . . .

We thank Maxim and also Fabian Haiden for important conversations.

Consider X a Riemann surface, $x_0 \in X$, $E \to X$ a vector bundle of rank r with $\bigwedge^r E \cong \mathcal{O}_X$, and

$$\varphi : E \to E \otimes \Omega^1_X$$

a Higgs field with $\mathrm{Tr}(\varphi) = 0$. Let

$$\Sigma \subset T^*X \xrightarrow{p} X$$

be the spectral curve, which we assume to be reduced.

We have a tautological form

$$\phi \in H^0(\Sigma, p^*\Omega^1_X)$$

which is thought of as a multivalued differential form. Locally we write

$$\phi = (\phi_1, \ldots, \phi_r), \quad \sum \phi_i = 0.$$

The assumption that Σ is reduced amounts to saying that ϕ_i are distinct.

Let $D = p_1 + \ldots + p_m$ be the locus over which Σ is branched, and $X^* := X - D$. The ϕ_i are locally well defined on X^*.

There are 2 kinds of WKB problems associated to this set of data.

(1) The **Riemann-Hilbert** or **complex WKB** problem:

Choose a connection ∇_0 on E and set

$$\nabla_t := \nabla_0 + t\varphi$$

for $t \in \mathbb{R}_{\geq 0}$. Let

$$\rho_t : \pi_1(X, x_0) \to SL_r(\mathbb{C})$$

be the monodromy representation. We also choose a fixed metric h on E.

From the flat structure which depends on t we get a family of maps

$$h_t : \widetilde{X} \to SL_r(\mathbb{C})/SU_r$$

which are ρ_t-equivariant. We would like to understand the asymptotic behavior of ρ_t and h_t as $t \to \infty$.

Definition: For $P, Q \in \widetilde{X}$, let $T_{PQ}(t) : E_P \to E_Q$ be the transport matrix of ρ_t. Define the *WKB exponent*

$$\nu_{PQ} := \limsup_{t \to \infty} \frac{1}{t} \log \|T_{PQ}(t)\|$$

where $\|T_{PQ}(t)\|$ is the operator norm with respect to h_P on E_P and h_Q on E_Q.

(2) The **Hitchin WKB problem**:

Assume X is compact, or that we have some other control over the behavior at infinity. Suppose (E, φ) is a stable Higgs bundle. Let h_t be the Hitchin Hermitian-Yang-Mills metric on $(E, t\varphi)$ and let ∇_t be the associated flat connection. Let $\rho_t : \pi_1(X, x_0) \to SL_r(\mathbb{C})$ be the monodromy representation.

Our family of metrics gives a family of *harmonic maps*

$$h_t : \widetilde{X} \to SL_r(\mathbb{C})/SU_r$$

which are again ρ_t-equivariant.

We can define $T_{PQ}(t)$ and ν_{PQ} as before, here using $h_{t,P}$ and $h_{t,Q}$ to measure $\|T_{PQ}(t)\|$.

Gaiotto-Moore-Neitzke explain that ν_{PQ} should vary as a function of $P, Q \in X$, in a way dictated by the *spectral networks*. We would like to give a geometric framework.

Remark: In the complex WKB case, one can view $T_{PQ}(t)$ in terms of Ecalle's *resurgent functions*. The Laplace transform

$$\mathcal{L}T_{PQ}(\zeta) := \int_0^\infty T_{PQ}(t)e^{-\zeta t}dt$$

is a holomorphic function defined for $|\zeta| \geq C$. It admits an analytic continuation having infinite, but locally finite, branching.

One can describe the possible locations of the branch points, and this description is compatible with the discussion of G.M.N., however today we look in a different direction.

How are **buildings** involved?

Basic idea: let $\mathcal{K}$ be a "field" of germs of functions on $\mathbb{R}_{\gg 0}$, with valuation given by "exponential growth rate". Then

$$\{\rho_t\} : \pi_1(X, x_0) \to SL_r(\mathcal{K}).$$

So, π_1 acts on the Bruhat-Tits building $\mathbf{B}(SL_r(\mathcal{K}))$, and we could try to choose an equivariant harmonic map

$$\widetilde{X} \to \mathbf{B}(SL_r(\mathcal{K}))$$

following Gromov-Schoen.

However, it doesn't seem clear how to make this precise.

Luckily, **Anne Parreau** has developed just such a theory, based on work of **Kleiner-Leeb:**

Look at our maps h_t as being maps into a symmetric space with distance rescaled:

$$h_t : \widetilde{X} \to \left(SL_r(\mathbb{C})/SU_r, \frac{1}{t} d \right).$$

Then we can take a “Gromov limit” of the symmetric spaces with their rescaled distances, and it will be a building modelled on the same affine space $\mathbb{A}$ as the SL_r Bruhat-Tits buildings.

The limit construction depends on the choice of *ultrafilter* ω, and the limit is denoted Cone_ω. We get a map

$$h_\omega : \widetilde{X} \to \mathsf{Cone}_\omega,$$

equivariant for the limiting action ρ_ω of π_1 on Cone_ω which was the subject of Parreau's paper.

The main point for us is that we can write

$$d_{\mathsf{Cone}_\omega}(h_\omega(P), h_\omega(Q)) =$$

$$\lim_\omega \frac{1}{t} d_{SL_r\mathbb{C}/SU_r}(h_t(P), h_t(Q))\,.$$

There are several distances on the building, and these are all related by the above formula to the corresponding distances on $SL_r\mathbb{C}/SU_r$.

- The Euclidean distance $\leftrightarrow$ Usual distance on $SL_r\mathbb{C}/SU_r$

- Finsler distance $\leftrightarrow$ log of operator norm

- Vector distance $\leftrightarrow$ dilation exponents

We are most interested in the *vector distance*. In the affine space

$$\mathbb{A} = \{(x_1, \ldots, x_r) \in \mathbb{R}^r, \ \sum x_i = 0\} \cong \mathbb{R}^{r-1}$$

the vector distance is translation invariant, defined by

$$\overrightarrow{d}(0, x) := (x_{i_1}, \ldots, x_{i_r})$$

where we use a Weyl group element to reorder so that $x_{i_1} \geq x_{i_2} \geq \cdots \geq x_{i_r}$.

In Cone_ω, any two points are contained in a common apartment, and use the vector distance defined as above in that apartment.

In $SL_r\mathbb{C}/SU_r$, put

$$\overrightarrow{d}(H,K) := (\lambda_1, \ldots, \lambda_k)$$

where

$$\|e_i\|_K = e^{\lambda_i}\|e_i\|_H$$

with $\{e_i\}$ a simultaneously H and K orthonormal basis.

In terms of transport matrices,

$$\lambda_1 = \log \|T_{PQ}(t)\|,$$

and one can get

$$\lambda_1 + \ldots + \lambda_k = \log \| \bigwedge^k T_{PQ}(t)\|,$$

using the transport matrix for the induced connection on $\bigwedge^k E$. Intuitively we can restrict to mainly thinking about λ_1. That, by the way, is the "Finsler metric".

Remark: For $SL_r\mathbb{C}/SU_r$ we are only interested in these metrics "in the large" as they pass to the limit after rescaling.

Our rescaled distance becomes

$$\frac{1}{t}\log\|T_{PQ}(t)\|.$$

Define the *ultrafilter exponent*

$$\nu^{\omega}_{PQ} := \lim_{\omega}\frac{1}{t}\log\|T_{PQ}(t)\|.$$

Notice that $\nu^{\omega}_{PQ} \leq \nu_{PQ}$. Indeed, the ultrafilter limit means the limit over some "cleverly chosen" subsequence, which will in any case be less than the $\limsup$.

Furthermore, we can say that these two exponents are equal in some cases, namely:

(a) for any fixed choice of P, Q, there exists a choice of ultrafilter ω such that $\nu_{PQ}^{\omega} = \nu_{PQ}$. Indeed, we can subordinate the ultrafilter to the condition of having a sequence calculating the $\limsup$ for that pair P, Q. It isn't *a priori* clear whether we can do this for all pairs P, Q at once, though. In our example, it will follow *a posteriori*!

(b) If $\limsup_t \ldots = \lim_t \ldots$ then it is the same as $\lim_\omega \ldots$. This applies in particular for the local WKB case. It would also apply in the complex WKB case, for generic angles, if we knew that $\mathcal{L}T_{PQ}(\zeta)$ didn't have essential singularities.

Theorem ("Classical WKB"):
Suppose $\xi : [0,1] \to \widetilde{X}^*$ is a short path, which is *noncritical* i.e. $\xi^*\mathrm{Re}\phi_i$ are distinct for all $t \in [0,1]$. Reordering we may assume

$$\xi^*\mathrm{Re}\phi_1 > \xi^*\mathrm{Re}\phi_2 > \ldots > \xi^*\mathrm{Re}\phi_r.$$

Then, for the complex WKB problem we have

$$\frac{1}{t}\overrightarrow{d}\,(h_t(\xi(0)), h_t(\xi(1)) \sim (\lambda_1, \ldots, \lambda_r)$$

where

$$\lambda_i = \int_0^1 \xi^*\mathrm{Re}\phi_i.$$

Corollary: At the limit, we have

$$\overrightarrow{d}_\omega (h_t(\xi(0)), h_t(\xi(1)) = (\lambda_1, \ldots, \lambda_r).$$

Conjecture: The same should be true for the Hitchin WKB problem.

Corollary: If $\xi : [0,1] \to \widetilde{X}^*$ is any noncritical path, then $h_\omega \circ \xi$ maps $[0,1]$ into a single apartment, and the vector distance which determines the location in this apartment is given by the integrals:

$$\overrightarrow{d}_\omega (h_t(\xi(0)), h_t(\xi(1)) = (\lambda_1, \ldots, \lambda_r).$$

This just follows from a fact about buildings: if A, B, C are three points with

$$\overrightarrow{d}(A,B) + \overrightarrow{d}(B,C) = \overrightarrow{d}(A,C)$$

then A, B, C are in a common apartment, with A and C in opposite chambers centered at B or equivalently, B in the Finsler convex hull of $\{A, C\}$.

Corollary: Our map

$$h_\omega : \widetilde{X} \to \mathsf{Cone}_\omega$$

is a harmonic ϕ-map in the sense of Gromov and Schoen. In other words, any point in the complement of a discrete set of points in $\widetilde{X}$ has a neighborhood which maps into a single apartment, and the map has differential Reϕ (no “folding”).

This finishes what we can currently say about the general situation: we get a harmonic ϕ-map

$$h_\omega : \widetilde{X} \to \mathsf{Cone}_\omega$$

depending on choice of ultrafilter ω, with

$$\nu^\omega_{PQ} \leq \nu_{PQ},$$

and we can assume that equality holds for one pair P, Q. Also equality holds in the local case. We expect that one should be able to choose a single ω which works for all P, Q.

Now, we would like to analyse harmonic ϕ-maps in terms of spectral networks.

The main observation is just to note that the reflection hyperplanes in the building, pull back to curves on $\tilde{X}$ which are imaginary foliation curves, including therefore the spectral network curves.

Indeed, the reflection hyperplanes in an apartment have equations $x_{ij} = \text{const.}$ where $x_{ij} := x_i - x_j$, and these pull back to curves in $\widetilde{X}$ with equation $\text{Re}\phi_{ij} = 0$. This is the equation for the spectral network curves.

The Berk-Nevins-Roberts (BNR) example

In order to see how the the *collision* spectral network curves play a role in the harmonic map to a building, we decided to look closely at a classical example: it was the original example of Berk-Nevins-Roberts.

They seem to be setting $\hbar = 1$, a standard physicist's move. If we undo that, we can say that they consider a family of differential equations with large parameter t, of the form

$$(\frac{1}{t^3}\frac{d^3}{dx^3} - \frac{3}{t}\frac{d}{dx} + x)f = 0.$$

When we use the companion matrix we obtain a Higgs field φ with spectral curve given by the equation

$$\Sigma \quad : \quad y^3 - 3y + x = 0$$

where $X = \mathbb{C}$ with variable x, and y is the variable in the cotangent direction.

The differentials ϕ_1, ϕ_2 and ϕ_3 are of the form $y_i dx$ for y_1, y_2, y_3 the three solutions.

Notice that $\Sigma \to X$ has branch points

$$p_1 = 2, \quad p_2 = -2.$$

The imaginary spectral network is as in the accompanying picture.

We will continue the discussion using a different pdf with pictures. Let's just sum up here what will be seen.

There are two collision points, which in fact lie on the same vertical collision line.

The spectral network curves divide the plane into 10 regions:
4 regions on the outside to the right of the collision line;
4 regions on the outside to the left of the collision line;
2 regions in the square whose vertices are the singularities and the collisions; the two regions are separated by the interior part of the collision line.

Arguing with the local WKB approximation, we can conclude that each region is mapped into a single Weyl sector in a single apartment of the building Cone_{ω}.

The interior square maps into a single apartment, with a fold line along the "caustic" joining the two singularities. The fact that the whole region goes into one apartment comes from an argument with the axioms of the building. We found the paper of Bennett, Schwer and Struyve about axiom systems for buildings, based on Parreau's paper, to be very useful.

It turns out, in this case, that the two collision points map to the same point in the building. This may be seen by a contour integral using the fact that the interior region goes into a single apartment.

Therefore, the sectors in question all correspond to sectors in the building with a single vertex.

(A small *caveat* is needed: it only seems possible to state that any arbitrarily large compact subset goes into a single apartment; one should perhaps "complete" the system of apartments to include any embedding of $\mathbb{A}$ which is an isometry for the vector distance. We can ignore this point.)

In view of the picture of sectors starting from a single vertex, we can switch over from affine buildings to spherical buildings. An SL_3 spherical building is just a graph, such that any two points have distance maximum 3, any two edges are contained in a hexagon, and there are no loops smaller than a hexagon.

Our situation corresponds to an octagon: the eight exterior sectors. In this case, one can inductively construct a spherical building, by successively completing uniquely each path of length 4 to a hexagon. It corresponds to completing any adjacent sequence of 4 distinct sectors, to their convex hull which is an apartment of 6 sectors.

The two sectors which contain the images of the two interior zones of X, are just the first two new sectors which would be added to the octagon.

The main observation is that opposite edges in the octagon cannot go into a hexagon intersecting the octagon in 5 segments. Rather, they have to go to a twisted hexagon reversing directions. It is here that we see the collision phenomenon.

The inverse image of the apartment corresponding to this twisted hexagon, in X, is disconnected. Thus, if P, Q are points in the opposite sectors, then the distance in the building is not calculated by any integral of a single 1-form from P to Q. The 1-form has to jump when we cross a collision line. This is the collision phenomenon.

We obtain a *universal building* B^ϕ together with a harmonic ϕ-map

$$h^\phi :\to B^\phi$$

such that for any other building $\mathcal{C}$ (in particular, $\mathcal{C} = \mathrm{Cone}_\omega$) and harmonic ϕ-map $X \to \mathcal{C}$ there is a unique factorization

$$X \to B^\phi \xrightarrow{g} \mathcal{C}.$$

In our example, we furthermore have the property that on the Finsler secant subset of the image of X, g is an isometry for any of the distances. It depends on the non-folding property of g.

Therefore, we conclude in our example that *distances in $\mathcal{C}$ between points in X are the same as the distances in B^{ϕ}.*

This allows us to make a stronger conclusion in our example. Remember that for any P, Q it was possible to choose ω such that the distance from P to Q in Cone$_\omega$ was the same as the WKB dilation exponent vector for $T_{PQ}(t)$.

Therefore, again in our example, we conclude that

• the WKB dilation exponent is calculated as the distance in the building B^{ϕ},

$$\overrightarrow{\nu}_{PQ} = \overrightarrow{d}_{B^{\phi}}(h^{\phi}(P), h^{\phi}(Q)).$$

There exist examples (e.g. pullback connections) where we can see that the isometry property cannot be true in general. However, we conjecture that it is true if the spectral curve Σ is smooth and irreducible, and ∇_0 is generic.

The universal property doesn't necessarily require having an isometry, and we make the following

Conjecture: For any spectral curve with multivalued differential ϕ, there is a universal ϕ-map to a "building" or building-like object.

It might be necessary to restrict to some kind of target object which is somewhat smaller than a building.

In the case of SL_2, the universal building B^ϕ is just the space of leaves of the foliation defined by Reϕ.

Hence, we are trying here to obtain a generalization of the "space of leaves" picture, to the higher-rank case.

It is hoped that this will help with stability conditions on categories.