Notes on Moduli Spaces of Flat Connections with Boundary Conditions

Yiannis Loizides, Spring 2012

Contents

1 Introduction

In a recent paper [3] P. Ševera applies the method of Atiyah and Bott [2] to construct a symplectic form on the moduli space of flat connections on a surface with boundary satisfying boundary conditions given by Lagrangian Lie subalgebras. In these notes we give an exposition of the early parts of Ševera's paper (those giving the construction of the symplectic form), and also discuss briefly the possibility of obtaining Ševera's moduli spaces using the theory of Lie group-valued moment maps [1].

2 Notation and setup

If f is a smooth map between manifolds, Tf will denote the induced tangent map. Let G be a connected Lie group with Lie algebra $\mathfrak{g}$ possessing an Ad-invariant inner product $\langle , \rangle$. Given $g \in G$, l_g and r_g will denote the maps $G \to G$ given by left and right multiplication by g respectively. We use θ^L, θ^R to denote the left and right Maurer-Cartan forms respectively. We will be most interested in the case where $\langle , \rangle$ has split signature. If M is a manifold then $G(M)$, $\mathfrak{g}(M)$ will denote the space of smooth maps $M \to G$ and $M \to \mathfrak{g}$ respectively.

The situation that Ševera considers is as follows. Let Σ be a compact oriented surface with boundary, and we allow the boundary to have corners (recall for manifolds with boundary we allow some charts to have half-spaces $\mathbb{R}_{\geq 0} \times \mathbb{R}^{n-1}$ for domains; in this case we also allow some charts to have $(\mathbb{R}_{\geq 0})^2$ for domains). For simplicity Ševera also assumes that none of the components of Σ is closed, and that on each component of $\partial\Sigma$ there is at least one corner (he makes a further assumption (see below) that ensures there are at least two corners on each component of the boundary). We thus get a picture of the boundary of Σ as consisting of a finite number of components, each of which looks like an n-gon for some $n = 2, 3, 4, \ldots$ (for $n = 2$ imagine curved sides). We refer to segments of the boundary between two adjacent corners as "arcs" or "edges".

Ševera then defines a *colored surface* as a surface Σ as described above whose arcs a have been decorated with Lagrangian Lie subalgebras $\mathfrak{h}_a^{\perp} = \mathfrak{h}_a \subset \mathfrak{g}$, and such that whenever a and b are adjacent arcs then $\mathfrak{h}_a \cap \mathfrak{h}_b = \{0\}$ (this in turn means that for a colored surface, each component of the boundary must have at least two corners). By the Lie correspondence, for each $\mathfrak{h}_a$ we get a corresponding connected Lie subgroup $H_a \subset G$ (i.e. the subgroup generated by $\exp(\mathfrak{h}_a)$, or equivalently the leaf through the identity $e \in G$ of the integrable distribution $g \in G \mapsto g \cdot \mathfrak{h}_a \subset T_gG$; we recall that this is a Lie group, but in general it need not be closed in G, it is an immersed submanifold of G).

3 Moduli space of colored flat connections

Given G as above and a colored surface Σ, Ševera defines a symplectic manifold $\mathbf{M}_\Sigma$, which is a moduli space of flat connections over Σ satisfying certain boundary conditions, i.e. the connection must be a *colored* connection, to be defined below.

3.1 Connections

We begin by briefly recalling some basic facts about connections–we'll be quite selective and only describe facts that are crucial to what follows (see one or more standard texts for more information). Let $\pi : P \to \Sigma$ be a principal G bundle, with left G-action. For each $\xi \in \mathfrak{g}$ we define a vector field ξ_P on P by

$$\xi_P(p) = \frac{d}{dt}\bigg|_{t=0} \exp(-t\xi) \cdot p.$$

Recall that a connection on a principal G-bundle $\pi : P \to \Sigma$ is an equivariant $\mathfrak{g}$-valued 1-form A on P satisfying

$$A(\xi_P) = \xi.$$

Equivariance means that for $g \in G$ we have $L_g^* A = Ad_g \circ A$, where L_g denotes left action by g. We denote the space of all connections on a G-bundle P by $\mathcal{A}(P)$.

Given an open covering U_i of Σ together with smooth local sections $\phi_i : U_i \to \pi^{-1}(U_i)$, we may think of a connection as a collection of $\mathfrak{g}$-valued 1-forms $A_i = \phi_i^{\star} A$ on U_i which obey a compatibility condition on overlaps:

$$A_i = Ad_{t_{ij}} \circ A_j - t_{ij}^* \theta^R, \tag{1}$$

where $t_{ij}\phi_j = \phi_i$ on $U_i \cap U_j$ is the transition function. The proof follows essentially from the Leibniz rule. Let $X = x'(0)$ be a vector tangent to the curve $x(\lambda)$ in Σ and put $x_0 = x(0)$. Then

$$\begin{aligned}
T\phi_i \cdot X &= T(t_{ij}\phi_j) \cdot X \\
&= \frac{d}{d\lambda}\bigg|_0 t_{ij}(x(\lambda))\phi_j(x(\lambda)) \\
&= \frac{d}{d\lambda}\bigg|_0 t_{ij}(x(\lambda))\phi_j(x_0) + \frac{d}{d\lambda}\bigg|_0 t_{ij}(x_0)\phi_j(x(\lambda)) \qquad \text{(Leibniz)} \\
&= \frac{d}{d\lambda}\bigg|_0 t_{ij}(x(\lambda))t_{ij}(x_0)^{-1}\phi_i(x_0) + L_{t_{ij}}T\phi_j \cdot X \\
&= \frac{d}{d\lambda}\bigg|_0 \exp(\lambda t_{ij}^*\theta^R(X)) \cdot \phi_i(x_0) + L_{t_{ij}}T\phi_j \cdot X \\
&= -\xi_P + L_{t_{ij}}T\phi_j \cdot X,
\end{aligned}$$

where ξ_P is the vertical vector on P generated by $t_{ij}^*\theta^R(X)$. Applying A to both sides and using the defining properties of A (i.e. $L_g^*A = Ad_g A$ and $A(\xi_P) = \xi$) yields

$$A_i(X) = -t_{ij}^*\theta^R(X) + Ad_{t_{ij}}A_j(X).$$

Conversely a collection of $\mathfrak{g}$-valued 1-forms A_i defined on opens sets U_i obeying this compatibility condition give rise to a unique connection on P. Sometimes we will take the collection of local sections (U_i, ϕ_i) as understood, and (informally) write $A \in \Omega^1(\Sigma) \otimes \mathfrak{g}$ (understanding that in fact this refers to a collection of local expressions for A, compatible on overlaps). We will switch between the local and global descriptions of a connection depending on convenience.

A connection on a principal bundle P gives rise to a covariant derivative on associated vector bundles; let ζ be a section of an associated vector bundle $E = P \times_\rho V$ (so $\zeta : P \to V$ and is equivariant with respect to the action of G on P and on the representation (V, ρ)), then using a local section $\phi : U \subset \Sigma \to P$ to pull back the connection and the section ζ we have

$$d_A\zeta = d\zeta + A \cdot \zeta,$$

where $A \cdot \zeta$ indicates the Lie algebra action (and we really mean ϕ^*A, $\phi^*\zeta$ but have left this out to make the formula less cluttered). In particular recall that $ad\ P = P \times_{Ad\ G} \mathfrak{g}$ (i.e. $\mathfrak{g}$ carries the adjoint representation). For ξ a section of $ad\ P$ we have

$$d_A\xi = d\xi + [A, \xi].$$

The *horizontal distribution* of A is the distribution $p \in P \mapsto ker(A_p)$. A vector $X \in T_pP$ is called *horizontal* if it lies in the horizontal distribution. If we have a smooth curve

$\gamma : [0,1] \to \Sigma$ together with a point $p \in \pi^{-1}(\gamma(0))$, then there is a unique curve $\tilde{\gamma}$ in P called the *horizontal lift* of γ through p satisfying (1) $\pi \circ \tilde{\gamma} = \gamma$, (2) $\tilde{\gamma}$ has horizontal tangent vector field, and (3) $\tilde{\gamma}(0) = p$. If γ is closed then $\tilde{\gamma}(1)$ lies in the same fibre as $\tilde{\gamma}(0)$, and so there is some unique group element $g \in G$ such that $\tilde{\gamma}(1) = g \cdot \tilde{\gamma}(0)$. This prescription defines a mapping from the collection of smooth closed curves at $\pi(p)$ to G, namely $\gamma \mapsto g$. We'll refer to this map as the *holonomy map*; to define it we need both a connection and a choice of base point $p \in P$. The image of this mapping is denoted $Hol_p(A)$ and is called the *holonomy group* of A at p.

On the other hand if $h \cdot p =: q \in \pi^{-1}(x)$ is some other point in the same fibre as p, and if $\tilde{\gamma}$ is the horizontal lift of a closed curve based at x having $\tilde{\gamma}(0) = p$, $\tilde{\gamma}(1) = g \cdot p$ then $h \cdot \tilde{\gamma}$ is a horizontal lift of the same closed curve based at x but now having $(h \cdot \tilde{\gamma})(0) = h \cdot p = q$ and $(h \cdot \tilde{\gamma})(1) = h \cdot g \cdot p = (Ad_h g) \cdot q$. This shows that we have a canonical isomorphism

$$Hol_{p \cdot h}(A) = Ad_h Hol_p(A). \tag{2}$$

If p and q are in different fibres, and assuming Σ is path-connected, we can choose any curve from $\pi(p)$ to $\pi(q)$, lift to a horizontal curve α from q to some point p' in the same fibre as p. Then the curves $\alpha_g := g \cdot \alpha$ as g ranges over G are all horizontal. If γ is a closed curve at $\pi(p')$ with $\tilde{\gamma}$ the horizontal lift starting at p' and ending at $h \cdot p'$ then we get a horizontal curve starting at q by following α then $\tilde{\gamma}$ and then following α_h backwards. This gives an isomorphism from $Hol_p(A)$ to $Hol_q(A)$, although notice that in this case (as for fundamental groups), the isomorphism is not canonical since it depended on the choice of a curve from $\pi(q)$ to $\pi(p)$.

Given a connection A, its curvature F is a G-equivariant, *horizontal* $\mathfrak{g}$-valued 2-form on P given by

$$F = dA + \tfrac{1}{2}[A, A].$$

Horizontal means that F vanishes on directions tangent to the fibres, i.e. if $v \in ker(T\pi)$ then $\iota(v)F = 0$. Since F is both equivariant and horizontal, it may equivalently be thought of as an *ad* P-valued 2-form on Σ. There are various interpretations of the curvature (see one or more standard texts, or e.g. the discussion in [2]). One that is particularly important for us (and so we give a proof) is the following.

Proposition 3.1. *Let X, Y be horizontal vector fields, then*

$$A([X, Y]) = -F(X, Y).$$

Proof. Using the formula for the exterior derivative of a 1-form we have

$$dA(X, Y) = XA(Y) - YA(X) - A([X, Y]).$$

The first two terms vanish since X, Y are horizontal vector fields. Also

$$\tfrac{1}{2}[A, A](X, Y) = [A(X), A(Y)],$$

which also vanishes, again because X, Y are horizontal. Plugging these two results into the definition of F yields

$$F(X,Y) = -A([X,Y]).$$

□

By Frobenius' theorem, the curvature is exactly the obstruction to the integrability of the horizontal distribution. We will be interested in flat connections, so $F = 0$ and the horizontal distribution gives us a foliation of P. This has several useful consequences. Fix $p \in P$. The leaf of the horizontal foliation through p, together with the map π gives a covering space of Σ. So we can carry over standard results about covering spaces. In particular, the endpoint of the horizontal lift to p of a curve starting at $x = \pi(p)$ depends only on the homotopy class of the curve. The holonomy mapping thus induces a homomorphism from the fundamental group $\pi_1(\Sigma, x)$ to G.

We can also ask the reverse question: given a homomorphism ϕ from $\pi_1(\Sigma, x)$ to G, is there a principal G bundle $\pi : P \to \Sigma$ and a point $p \in \pi^{-1}(x)$ together with a flat connection A on P such that the holonomy mapping is the homomorphism ϕ? The answer is yes and we'll see this in detail below when we discuss colored connections. We'll show further that holonomy completely characterizes a connection once we allow for "gauge equivalence", which we discuss next.

3.2 Gauge group

Let $\pi : P \to \Sigma$ be a principal G bundle as above. The gauge group $Gau(P)$ consists of smooth automorphisms f of P (that is, fibre-preserving G-equivariant diffeomorphisms) which induce the identity map on Σ: $\pi \circ f = \pi$. Such a gauge transformation may be described by a smooth function $\tilde{f} : P \to G$

$$f(p) = \tilde{f}(p) \cdot p,$$

which is equivariant

$$\tilde{f}(g \cdot p) = Ad_g \tilde{f}(p).$$

We check that

$$\begin{aligned} f(g \cdot p) &= \tilde{f}(g \cdot p) g \cdot p \\ &= g \cdot \tilde{f}(p) \cdot p. \end{aligned}$$

For some purposes it is simpler to deal with gauge transformations by working locally, i.e. assume we have local trivializations (U_α, ϕ_α) and work with the local expressions $A_\alpha := \phi_\alpha^* A$, which are $\mathfrak{g}$-valued 1-forms on open subsets U_α of Σ. On a local patch U_α a gauge transformation $\tilde{f}$ is represented by the map $g_\alpha := \tilde{f} \circ \phi_\alpha : U_\alpha \to G$ (the full gauge transformation over U_α can be recovered from g_α using equivariance). Thus a gauge transformation is equivalently given by a collection of maps $g_\alpha : U_\alpha \to G$ which are compatible on overlaps (the point being that this way we can think of a gauge transformation as something defined on

Σ, rather than as something defined on P but which is constrained by equivariance). For simplicity of notation we will usually just write $A \in \Omega^1(\Sigma) \otimes \mathfrak{g}$ and $g \in G(\Sigma)$, and by this notation we implicitly mean that we have a collection of such things defined locally and compatible on overlaps.

In this notation, let's now write down the action of the gauge group on the space of connections. If $g \in G(\Sigma)$ and $A \in \Omega^1(\Sigma) \otimes \mathfrak{g}$ then

$$g \cdot A := Ad_g A - g^* \theta^R. \tag{3}$$

In the global picture, the action $f \in Gau(P)$ can be thought of as pushing-forward the horizontal subspaces by Tf (so that in particular, horizontal curves are mapped to horizontal curves), or equivalently as pulling the globally defined 1-form on P back by the inverse $(f^{-1})^*$. Working out a local expression gives (3)–the proof is the same as that given above for equation (1). Some intuition as to its form: the first term does not represent a change in the connection, but rather arises because we've moved by g on P–it comes from the equivariance property $L_g^* A = Ad_g \circ A$. The second term represents a "tilting" of the horizontal subspaces caused by the fact that g varies on the base, the pull-back of the Maurer-Cartan form giving exactly this correction.

Under a gauge transformation the curvature $F \in \Omega^2(\Sigma) \otimes \mathfrak{g}$ transforms more simply

$$F \mapsto Ad_g F. \tag{4}$$

Unlike (3), there is no affine term. This means that the subset of flat connections $\mathcal{A}^{flat}(P)$ is preserved by the gauge group.

Given a $\mathfrak{g}$-valued function ξ on Σ, we get a 1-parameter group of gauge transformations by putting $g(\lambda) = \exp(\lambda\xi)$. If we put this into the formula (3) and take the derivative with respect to λ at 0 we get the vector field $\xi_{\mathcal{A}}$ generated by ξ on $\mathcal{A}(P)$

$$\xi_{\mathcal{A}}(A) = -[A, \xi] - d\xi = -d_A \xi. \tag{5}$$

We see that tangents to the gauge orbits are d_A-exact forms.

3.3 Colored connections

The moduli space of flat connections $\mathbf{M}_\Sigma$ which Ševera considers are connections (up to gauge equivalence) on Σ satisfying certain boundary conditions. Recall that each edge a is decorated with a Lagrangian Lie subalgebra $\mathfrak{h}_a$, and the corresponding connected Lie subgroup is $H_a \subset G$. Ševera defines a *colored G-bundle* over Σ to be a G-bundle over Σ along with additional structure: for each edge a a submanifold $P_a \subset \pi^{-1}(a)$ which is a principal H_a-bundle over a (in other words, we have a reduction of the structure group to H_a over a) together with a choice of point $p_x \in P_a \cap P_b$ for each pair of edges a, b meeting at a corner x. So we can denote a colored G-bundle by (P, P_a, p_x) where a is understood to run over all the edges, and x over all the vertices. In particular notice that we have a canonical trivialization of the fibres over the corners (map p_x to $e \in G$). We will call these special chosen points in the fibres above the corners "marked points" or "special points".

Two colored G-bundles $(P^{(i)}, P_a^{(i)}, p_x^{(i)})$, $i = 1, 2$ having the same $\mathfrak{h}_a$ decorating the edges a, are isomophic if there is a fibre-preserving, G-equivariant diffeomorphism $f : P^{(1)} \to P^{(2)}$ which induces the identity on Σ and such that $f(P_a^{(1)}) = P_a^{(2)}$, $f(p_x^{(1)}) = f(p_x^{(2)})$. We will identify colored G-bundles up to such isomorphisms. So for example if we have connections A_i on isomorphic colored G-bundles $(P^{(i)}, P_a^{(i)}, p_x^{(i)})$, $i = 1, 2$, we will identify the bundles and assume we have two connections A_i, $i = 1, 2$ on a fixed colored G-bundle (P, P_a, p_x). If we like we can assume we've chosen a unique fixed representative (P, P_a, p_x) of each isomorphism class.

Given a colored G-bundle (P, P_a, p_x), Ševera defines a *colored connection* to be a connection which restricts to an $\mathfrak{h}_a$-valued form on each P_a (this makes sense because P_a is a principal H_a-bundle, and $\mathfrak{h}_a$ is invariant under the adjoint action of H_a). The space of colored connections on P is denoted $\mathcal{A}_{col}(P)$, and collecting these spaces together for non-isomorphic P we get the space of colored connections over Σ, denoted $\mathcal{A}_{col}(\Sigma)$. We remark again that when we say "non-isomorphic" we are using the (stronger) notion of isomorphism of *colored* G-bundles. If we restrict to flat connections, the corresponding spaces are denoted $\mathcal{A}_{col}^{flat}(P)$ and $\mathcal{A}_{col}^{flat}(\Sigma)$.

Similarly, we will be interested in a subgroup of the usual gauge group which preserves the colored G-bundle structure, i.e. the submanifolds P_a and the marked points p_x. We will refer to this subgroup as the *colored gauge transformations* $Gau_{col}(P)$. Since these gauge transformations preserve the submanifolds P_a, if we describe a gauge transformation locally in terms of maps $g_i : V_i \subset \Sigma \to G$, then we can assume the g_i to be H_a-valued over a (they restrict to a gauge transformation of the H_a-bundle P_a over a). Colored gauge transformations restrict to the identity over the corners so that the marked points $p_x \in P_a \cap P_b$ are preserved. In particular, colored gauge transformations are automatically *based* (this means in particular that the action of the colored gauge group on the space of colored connections is free since for a non-trivial transformation the affine term is always present, and (neglecting analytical issues) the resulting quotient is nonsingular).

Given a connection A on a colored bundle, we can make sense of the "holonomy" of A along a curve γ between two corners $\gamma(0) = x, \gamma(1) = y$. This is because we have canonical trivializations of the fibres above the corners. In detail, let p_x, p_y be the marked points in the fibres above x, y respectively. Let $\tilde{\gamma}$ be the lift of γ with initial point $\tilde{\gamma}(0) = p_x$. Then the holonomy of A along γ is the element $g \in G$ such that $\tilde{\gamma}(1) = g \cdot p_y$. Let us introduce some notation: $hol(A, \gamma)$ will denote this unique element $g \in G$. We can think of $hol(\gamma)$ as a function $A \in \mathcal{A}_{col}(\Sigma) \mapsto hol(A, \gamma) \in G$. Because the group of *colored* gauge transformations act as the identity on fibres above corners, it preserves $hol(A, \gamma)$, and so $hol(\gamma)$ can be thought of as a function on the (colored) gauge orbits.

If furthermore A is flat, then $hol(A, \gamma)$ depends only on the homotopy class of γ. Let X

denote the set of corners of Σ, and $\Pi(\Sigma, X)$ the set of homotopy classes of paths between points of X (note in particular that each edge $a \in \partial\Sigma$ may be thought of as an element in $\Pi(\Sigma, X)$). Then we've shown that we have a (gauge invariant) function

$$hol : \Pi(\Sigma, X) \times \mathcal{A}_{col}^{flat}(P) \to G.$$

We also note that $\Pi(\Sigma, X)$ is a *groupoid* over X with source $s : \Pi_1(\Sigma, X) \to X$ the map taking a homotopy class of curves to its common *end point*, and target $t : \Pi(\Sigma, X) \to X$ the map taking a homotopy class of curves to its common initial point. The multiplication is just concatenation of (homotopy classes of) paths $(\alpha, \beta) \mapsto \alpha \star \beta$ (where this means first follow α, then β). If we fix a flat connection A then the map

$$\begin{aligned} hol(A) : \Pi(\Sigma, X) &\to G \\ \gamma &\mapsto hol(A, \gamma) \end{aligned}$$

is a homomorphism of groupoids (i.e. a covariant functor). Let's show this.

Let $x, y, z \in X$ with corresponding marked points p_x, p_y, p_z, let $\alpha, \beta \in \Pi(\Sigma, X)$ be curves with $t(\alpha) = x, s(\alpha) = t(\beta) = y, s(\beta) = z$, and let $\tilde{\alpha}, \tilde{\beta}$ be lifts of α, β with initial points p_x, p_y respectively. A fixed connection A is understood throughout. By definition, $\tilde{\alpha}(1) = hol(\alpha) \cdot p_y$ and $\tilde{\beta}(1) = hol(\beta) \cdot p_z$. Also $\tilde{\alpha} \star (hol(\alpha) \cdot \tilde{\beta})$ is a lift of $\alpha \star \beta$ and so

$$\begin{aligned} hol(\alpha \star \beta) \cdot p_z &= (\tilde{\alpha} \star (hol(\alpha) \cdot \tilde{\beta}))(1) \\ &= hol(\alpha) \cdot hol(\beta) \cdot p_z \\ \Rightarrow hol(\alpha \star \beta) &= hol(\alpha) \cdot hol(\beta). \end{aligned}$$

Using the same method we get a similar result for inverses: $hol(\bar{\alpha}) = hol(\alpha)^{-1}$ where $\bar{\alpha}$ denotes α traversed in the opposite direction.

3.4 The moduli space

The moduli space of flat colored connections on a bundle P is

$$\mathbf{M}_\Sigma(P) := \mathcal{A}_{col}^{flat}(P)/Gau_{col}(P).$$

We take the (disjoint) union of these spaces for non-isomorphic colored G-bundles P to get the moduli space $\mathbf{M}_\Sigma$. We'll write this as

$$\mathbf{M}_\Sigma = \mathcal{A}_{col}^{flat}(\Sigma)/Gau_{col}$$

to indicate that $\mathbf{M}_\Sigma$ is the space of colored flat connections on Σ up to (colored) gauge equivalence.

Ševera gives two other equivalent descriptions of $\mathbf{M}_\Sigma$ in terms of holonomies. We'll repeat these other descriptions here.

Using cuts

Cut Σ along curves connecting corners until we get a polygon. Notice that the two vertices bounding any edge in the resulting polygon came from corners in the original surface Σ. Let s label the sides of the polygon. The moduli space $\mathbf{M}_\Sigma$ is the space of assignments $s \mapsto \gamma_s \in G$ satisfying

1. if s is part of the boundary of Σ then $\gamma_s \in H_s$,
2. if s, s' are the two sides which arise from a single cut then $\gamma_{s'} = \gamma_s^{-1}$,
3. the product of the γ_s going cyclically around the boundary of the polygon is e.

This picture arises from the description of $\mathbf{M}_\Sigma$ in terms of connections by sending each colored flat connection A to the assignment $s \mapsto hol(A, s)$. The first condition is satisfied by virtue of the fact that A is colored, the second is automatic, and the third is satisfied because A is flat and the curve in Σ corresponding to the product of the γ_s around the boundary of the polygon is homotopic to a point. It doesn't matter which connection A we choose in any given Gau_{col} equivalence class because colored gauge transformations leave fixed the fibres over the corners and so don't change the holonomies (this is why unlike the pure Atiyah-Bott case, we don't need to quotient out by conjugations).

Let n be the number of cuts. This description of $\mathbf{M}_\Sigma$ shows us that

$$\mathbf{M}_\Sigma = \Phi^{-1}(e),$$

where

$$\begin{aligned} \Phi : \left(\prod_a H_a\right) \times G^n &\to G \\ (h_1, ..., h_k, g_1, ..., g_n) &\mapsto \prod_s \gamma_s. \end{aligned}$$

Here the h_i label segments of $\partial\Sigma$ while the g_i label cuts (there is one g_i per cut, and it appears in the product on the right once as g_i and once as g_i^{-1}).

As groupoid morphisms

The moduli space can also be described as a collection of functors F:

$$\mathbf{M}_\Sigma = \{F : \Pi(\Sigma, X) \to G | F(a) \in H_a \text{ for each edge } a \text{ in } \partial\Sigma\}.$$

This description is nice because it makes it clear that in the above description in terms of cuts, the resulting space does not depend on the particular cuts chosen. The cuts made in the polygon description correspond to a choice of generators for the groupoid $\Pi(\Sigma, X)$.

Given a Gau_{col}-equivalence class of connections $[A]$, we've seen already that $F := hol(A)$ gives a functor satisfying the above. We still need to explain the reverse, namely given a holonomy description (either as an assignment $s \mapsto \gamma_s$ or a functor $F : \Pi(\Sigma, X) \to G$), how to construct a colored flat connection with the required holonomies.

Let $\tilde{\Sigma}$ denote the space of homotopy classes of paths in Σ having initial point in X (recall that the universal covering space of Σ is the space of homotopy classes of paths having some fixed initial point, so $\tilde{\Sigma}$ is just a disjoint union of n copies of the universal covering space, where n is the number of corners). We have natural maps $f : \tilde{\Sigma} \to X$, which maps a homotopy class of curves α to their common initial point, and $q : \tilde{\Sigma} \to \Sigma$ the covering projection which maps a class of curves α to their common end point. There is a groupoid action of $\Pi(\Sigma, X)$ on $\tilde{\Sigma}$: let $\gamma \in \Pi(\Sigma, X)$ and $\alpha \in \tilde{\Sigma}$, whenever $s(\gamma) = f(\alpha)$ (i.e. γ has end point equal to the initial point of α), we can form $\gamma \star \alpha \in \tilde{\Sigma}$. When we go to the orbit space, we erase any information of about the homotopy class, and so the orbit space is isomorphic to Σ with isomorphism induced by the map q.

Now suppose we have a functor F satisfying the above. Consider the trivial bundle $\tilde{P} = \tilde{\Sigma} \times G$ with the product connection $\tilde{A}$, i.e. if $(\alpha, g) \in \tilde{\Sigma} \times G$ then $T_{(\alpha,g)}(\tilde{\Sigma} \times G) = T_\alpha \tilde{\Sigma} \times T_g G$ and if $(\tilde{Y}, -\xi \cdot g)$ is a tangent vector then

$$\tilde{A}((\tilde{Y}, -\xi \cdot g)) = \xi.$$

We can lift the action of the groupoid $\Pi(\Sigma, X)$ to $\tilde{P}$. When $s(\gamma) = f(\alpha)$ we define

$$\gamma \cdot (\alpha, g) = (\gamma \star \alpha, gF(\gamma)^{-1}).$$

Let $P = \tilde{P}/\Pi(\Sigma, X)$; we denote elements of the quotient using square brackets, so $[\alpha, g]$ denotes the orbit of (α, g). In particular when $s(\gamma) = f(\alpha)$ we have $[\alpha, g] = [\gamma \star \alpha, gF(\gamma)^{-1}]$. Because the groupoid action is equivariant for the left G-action on $\tilde{P}$, the quotient carries a left G-action. The projection is induced by q: $\pi([\alpha, g]) = q(\alpha)$. Thus P is a principal G-bundle over Σ.

Next let's show that $\tilde{A}$ induces a connection on P. Let $\gamma \in \Pi(\Sigma, X)$, $\alpha \in \tilde{\Sigma}$ with $s(\gamma) = f(\alpha)$, and $(\tilde{Y}, -\xi \cdot g) \in T_{(\alpha,g)}\tilde{P}$. The tangent map $T\gamma$ acts by

$$\begin{aligned} T\gamma\ (\tilde{Y}, -\xi \cdot g) &= (T\gamma \cdot \tilde{Y}, -\xi \cdot (g \cdot F(\gamma)^{-1})) \\ \Rightarrow (\gamma^* \tilde{A})(\tilde{Y}, -\xi \cdot g) &= \xi. \end{aligned}$$

This shows that $\tilde{A}$ is invariant under the groupoid action and so descends to a connection A on P. Since $\tilde{A}$ is flat, A is also flat. We can horizontally lift a curve γ in Σ to P by first lifting to $\tilde{\Sigma}$ (via q), then to $\tilde{P}$ using $\tilde{A}$, and then applying the quotient map. In detail, let γ be a curve in Σ with initial point x. Choose a curve (or homotopy class of curves) $\alpha \in \tilde{\Sigma}$ with initial point in X and end point x. Then the lift of γ to $\alpha \in \tilde{\Sigma}$ is just the curve $t \in [0, 1] \mapsto \alpha \star \gamma_t$ where

$$\gamma_t(\lambda) = \begin{cases} \gamma(\lambda) & \text{if } \lambda < t \\ \gamma(t) & \text{if } \lambda \geq t \end{cases}$$

(this is a curve in $\tilde{\Sigma}$, i.e. a 1-parameter family of (homotopy classes of) curves). In particular its end point is $\alpha \star \gamma$. Since $\tilde{A}$ is the product connection, the horizontal lift of $\alpha \star \gamma_t$ to a point $(\alpha, g) \in \tilde{P}$ is the curve $t \in [0,1] \mapsto (\alpha \star \gamma_t, g)$. Applying the quotient, we get a curve $t \mapsto [\alpha \star \gamma_t, g]$ in P, the horizontal lift of γ to the point $[\alpha, g] \in P$.

For each $x \in X$, let c_x denote the (homotopy class of the) constant curve at x. We take the marked point in the fibre above x to be $p_x = [c_x, e]$. As explained earlier, with marked points and a flat connection A, we get a map $hol(A) : \Pi(\Sigma, X) \to G$. Now suppose we have a curve $\gamma \in \Pi(\Sigma, X)$ with initial point x_0 and end point x_1 in X. As explained above, the lift to $c_{x_0} \in \tilde{\Sigma}$ is $c_{x_0} \star \gamma_t$. Then lift horizontally to $(c_{x_0} \star \gamma_t, e)$. Applying the quotient map, we find that the horizontal lift of γ to $p_{x_0} \in P$ is $t \in [0,1] \mapsto [c_{x_0} \star \gamma_t, e]$. The end point, which lies in the fibre above x_1 is $[c_{x_0} \star \gamma, e]$. Now $c_{x_0} \star \gamma$ is homotopic to $\gamma \star c_{x_1}$, hence

$$\begin{aligned} [c_{x_0} \star \gamma, e] &= [\gamma \star c_{x_1}, e] \\ &= [\gamma \star c_{x_1}, F(\gamma)F(\gamma)^{-1}] \\ &= [c_{x_1}, F(\gamma)] \\ &= F(\gamma) \cdot [c_{x_1}, e] \\ &= F(\gamma) \cdot p_{x_1}. \end{aligned}$$

This shows that $hol(A, \gamma) = F(\gamma)$ for every $\gamma \in \Pi(\Sigma, X)$. We've thus shown that A has the desired holonomies between pairs of corners.

The marked points also give us the rest of the colored structure. Consider an edge a bounded by corners $x_0, x_1 \in X$. We have submanifolds $H_a \cdot p_{x_0} \subset \pi^{-1}(x_0)$ and $H_a \cdot p_{x_1} \subset \pi^{-1}(x_1)$ above the corners. Now choose a point $p := h \cdot p_{x_0} \in H_a \cdot p_{x_0}$ and let $\tilde{a}$ denote the horizontal lift of a through p. By the above, the endpoint of $\tilde{a}$ is just $h \cdot F(a) \cdot p_{x_1}$. Since by assumption $F(a) \in H_a$, we have $h \cdot F(a) \cdot p_{x_1} \in H_a \cdot p_{x_1}$. This shows that if we lift a to a horizontal curve through each point in $H_a \cdot p_{x_0}$, the set of end points is exactly $H_a \cdot p_{x_1}$. It follows that the submanifold P_a obtained by parallel transporting $H_a \cdot p_{x_0}$ along a is the same as that got by parallel transporting $H_a \cdot p_{x_1}$ along a (in the opposite direction). The submanifold P_a is clearly an H_a sub-bundle, because left multiplication by an element $h \in H_a$ just permutes the curves linking $H_a \cdot p_{x_0}$ and $H_a \cdot p_{x_1}$ as described above. By construction it's clear that A restricts to P_a (i.e. the lift of a to any point in P_a doesn't leave P_a), and since H_a acts transitively on the fibres, the restriction takes values in $\mathfrak{h}_a$.

We've now explained how to obtain a colored flat connection on a colored G-bundle having holonomies prescribed by a functor F. We will refer to the specific colored G-bundle and connection arising from this construction (i.e. from $\tilde{\Sigma} \times G$ and the product connection by quotienting by an F-lift of the $\Pi(\Sigma, X)$-action) as the "standard example for the functor F". The last thing we would like to show is that the correspondence between equivalence classes of colored flat connections on colored G-bundles and holonomy descriptions F is one-one. We will do this by considering an arbitrary colored flat connection A on a colored G-bundle (P, P_a, p_x) with $F = hol(A)$, and showing that it is isomorphic to the standard example for the functor F.

Suppose we have a colored flat connection A on a colored G-bundle (P, P_a, p_x), and let $F = hol(A)$ be the corresponding holonomy functor. These induce a colored G-bundle $\tilde{P}$

and colored flat connection $\tilde{A}$ on $\tilde{\Sigma}$ by pull-back, i.e. $\tilde{P} = q^*P$, $\tilde{A} = q^*A$. As $\tilde{\Sigma}$ is simply connected, $\tilde{P}$ must be trivial–indeed we can find an explicit global section as follows. For each $x \in X$, we consider the point $c_x \in \tilde{\Sigma}$ (recall c_x denotes the homotopy class of the constant curve at x). Notice that the points c_x for different x are in different connected components of $\tilde{\Sigma}$ (recall that $\tilde{\Sigma}$ has n connected components (where n is the size of X), each isomorphic to the universal covering space of Σ; the connected component containing c_x with $x \in X$ consists of homotopy classes of curves having initial point x). Since $\tilde{A}$ is flat, its horizontal distribution gives rise to a foliation. Let x_i run through the points of X, and for each i let $\tilde{\Sigma}_i$ denote the connected component of $\tilde{\Sigma}$ containing c_{x_i}. For each connected component $\tilde{\Sigma}_i$ of $\tilde{\Sigma}$, we choose the leaf L_i of the foliation that passes through the point (c_{x_i}, p_{x_i}). The leaf L_i over $\tilde{\Sigma}_i$ is a covering space of $\tilde{\Sigma}_i$. Since $\tilde{\Sigma}_i$ is simply connected, it follows that L_i is isomorphic to $\tilde{\Sigma}_i$ via the fibre projection, and therefore corresponds to a section ϕ_i of $\tilde{P}$ defined over $\tilde{\Sigma}_i$. Putting the ϕ_i defined over separate connected components of $\tilde{\Sigma}$ together we get a global section ϕ of P which passes through all of the points (c_{x_i}, p_{x_i}). We use this section to trivialize $\tilde{P} = \tilde{\Sigma} \times G$. In particular the constant section e of $\tilde{\Sigma} \times P$ is a leaf of the horizontal foliation of the connection $\tilde{A}$ on $\tilde{\Sigma} \times G$ and so $\tilde{A}$ is the trivial (product) connection.

So now in principal we have two projection mappings $\tilde{P} \to P$, the first induced from the fact that $\tilde{P} = q^*P$, and the second from quotienting $\tilde{P} = \tilde{\Sigma} \times G$ by the F-lift of the $\Pi(\Sigma, X)$ groupoid action (as described previously). These maps agree on the points $(c_{x_i}, p_{x_i}) = (c_{x_i}, e)$ (one in each connected component of $\tilde{P}$) and are otherwise completely determined by parallel transport, and so must agree everywhere. Since the projection mappings map leaves of the horizontal foliation of $\tilde{A}$ into leaves of the horizontal foliation of A, it follows that the connection A is induced by $\tilde{A}$ in the manner of the standard example for the functor F. We've thus shown that our colored G-bundle and connection is isomorphic to the standard example for the functor F.

3.5 Symplectic structure

The space of connections $\mathcal{A}(P)$ on a G-bundle P form an affine space, and so in particular its tangent space at a point A is the vector space that $\mathcal{A}(P)$ is modelled on, and so can be thought of as the set of differences $a := A_1 - A_2$ where $A_i \in \mathcal{A}(P)$. Then a is an Ad-equivariant *horizontal* 1-form on P, or equivalently a 1-form with values in the vector bundle $ad\ P$. Given a local section ϕ defined on V, the pull-back ϕ^*a will be a $\mathfrak{g}$-valued 1-form on V. As for connections, we will sometimes write $a \in \Omega^1(\Sigma) \otimes \mathfrak{g}$ with the understanding that we mean a collection of forms defined locally and compatible on overlaps. Note that under a change of local section, the pull-back of a transforms according to the adjoint representation (see equation (1), the affine term cancels). Similarly if P is colored, then the sub-affine space $\mathcal{A}_{col}(P)$ consists of Ad-equivariant horizontal 1-forms on P which take values in $\mathfrak{h}_b$ when restricted to the sub-bundle P_b over each edge b. When local sections (V_i, ϕ_i) are chosen such that they are P_b-valued on each edge b, the pull-backs ϕ_i^*a are $\mathfrak{g}$-valued 1-forms on V_i taking values in $\mathfrak{h}_b$ over each edge b. We will sometimes write this as $a \in \Omega^1_{col}(\Sigma) \otimes \mathfrak{g}$, where as before it is understood that we refer to a collection of forms defined locally and compatible on overlaps.

We consider a colored flat connection $A \in \mathcal{A}_{col}^{flat}(P)$. We would like to determine the tangent space $T_A\mathcal{A}_{col}^{flat}(P)$. Let $a \in \Omega^1_{col}(\Sigma) \otimes \mathfrak{g}$, then writing $F(B)$ for the curvature of a connection B, we have

$$\begin{aligned} F(A+\lambda a) &= d(A+\lambda a) + \tfrac{1}{2}[A+\lambda a, A+\lambda a] \\ &= F(A) + \lambda d_A a + O(\lambda^2) \\ &= \lambda d_A a + O(\lambda^2). \end{aligned}$$

If a is tangent to the space of flat connections, then the curvature of $F(A+\lambda a)$ should vanish to first order in λ. From the above we see that this happens iff $d_A a = 0$. Hence

$$T_A\mathcal{A}_{col}^{flat}(P) = \{a \in \Omega^1_{col}(\Sigma) \otimes \mathfrak{g} | d_A a = 0\}.$$

We also note that if $b \in \Omega^k(\Sigma) \otimes \mathfrak{g}$ then by the graded Jacobi identity we have

$$[A,[A,b]] = \tfrac{1}{2}[[A,A],b].$$

And therefore

$$\begin{aligned} d_A^2 b &= (d + [A,\cdot]) \cdot (db + [A,b]) \\ &= d([A,b]) + [A,db] + [A,[A,b]] \\ &= [dA,b] - [A,db] + [A,db] + \tfrac{1}{2}[[A,A],b] \\ &= [F(A),b]. \end{aligned}$$

In particular if A is flat then $F(A) = 0$ and d_A is a differential. Recall that we also showed above that if $\xi \in \mathfrak{g}(\Sigma)$ generates a gauge transformation then the tangent to the gauge orbit in $\mathcal{A}_{col}^{flat}(P)$ is $-d_A\xi$. Putting this together we see that

$$T_{[P,A]}\mathbf{M}_\Sigma \simeq T_A\mathcal{A}_{col}^{flat}(P)/T_A(Gau_{col} \cdot A) = H^1(\Omega_{col}(\Sigma) \otimes \mathfrak{g}, d_A). \tag{6}$$

We define a 2-form on $\mathcal{A}_{col}(\Sigma)$. Let (V_i, ϕ_i) be a collection of local sections with the V_i covering Σ, and let ρ_i be a partition of unity subordinate to V_i. Then we define

$$\omega_{\mathcal{A}}(a,b) = \sum_i \rho_i \int_{V_i} \langle \phi_i^* a, \phi_i^* b\rangle. \tag{7}$$

That this is independent of the choice of local sections ϕ_i is immediate from the fact that under a change of local section the pull-backs of a, b transform according to the adjoint representation, while the inner product is Ad-invariant. In fact this means that the inner product $\langle \phi_i^* a, \phi_i^* b\rangle$ is a well-defined 2-form on Σ, and so abusing notation slightly we will write simply:

$$\omega_{\mathcal{A}}(a,b) = \int_\Sigma \langle a, b\rangle. \tag{8}$$

(The same holds for manipulations below involving $ad\ P$-valued forms on Σ.)

In fact as we'll show, when restricted to the flat connections $\mathcal{A}_{col}^{flat}(\Sigma)$ this vanishes on gauge orbit directions (and is constant on orbits), so it gives us a 2-form on $\mathbf{M}_\Sigma$:

$$\omega([a],[b]) = \int_\Sigma \langle a, b\rangle. \tag{9}$$

Let $b \in T_A\mathcal{A}_{col}^{flat}(\Sigma)$, so $d_A b = 0$. Let $\xi \in \mathfrak{g}(\Sigma)$ generate a 1-parameter group of colored gauge transformations, so ξ gives rise to a vector field $\xi_\mathcal{A}(A) = -d_A\xi$ on $\mathcal{A}_{col}^{flat}(\Sigma)$ tangent to the gauge orbit. Note that using compatibility of the inner product we have

$$\begin{aligned}\langle d\xi + [A,\xi], b\rangle &= d\langle \xi, b\rangle - \langle \xi, db\rangle - \langle \xi, [A,b]\rangle \\ &= d\langle \xi, b\rangle - \langle \xi, d_A b\rangle,\end{aligned}$$

and hence

$$\begin{aligned}\int_\Sigma \langle d_A\xi, b\rangle &= \int_\Sigma d\langle \xi, b\rangle - \int_\Sigma \langle \xi, d_A b\rangle \\ &= \int_{\partial\Sigma} \langle \xi, b\rangle.\end{aligned} \tag{10}$$

(The other integral vanishes because b is d_A-closed.) Now on each edge of $\partial\Sigma$, both ξ and b take values in the same Lagrangian Lie subalgebra (when ξ generates a *colored* gauge transformation and $b \in \Omega^1_{col}(\Sigma)\otimes\mathfrak{g}$). Hence $\langle \xi, b\rangle = 0$ on $\partial\Sigma$, and the expression above vanishes. This shows that the 2-form (9) is well-defined on $\mathbf{M}_\Sigma$. Furthermore, it is non-degenerate by Poincare duality, and in addition it is "constant" on $\mathbf{M}_\Sigma$ e.g. in the sense that the expression does not depend on A. So (9) warrants being called a symplectic form on $\mathbf{M}_\Sigma$; in the case without boundary this is the Atiyah-Bott symplectic structure on the moduli space of flat connections.

As noticed originally by Atiyah-Bott, the moduli space of flat connections modulo gauge transformations may be thought of as a symplectic reduction of the space of connections. The group is the gauge group and, in the case of a closed surface, the moment map is the curvature. In Ševera's case something similar works; notice from (10) we had

$$\begin{aligned}\omega_\mathcal{A}(d_A\xi, b) &= \int_{\partial\Sigma}\langle \xi, b\rangle - \int_\Sigma \langle \xi, d_A b\rangle \\ &= \frac{d}{d\lambda}\bigg|_0 \left(\int_{\partial\Sigma}\langle \xi, \lambda b\rangle - \int_\Sigma \langle \xi, F(A+\lambda b)\rangle\right),\end{aligned}$$

(this holds for arbitrary $\xi \in \mathfrak{g}(\Sigma)$, $A \in \mathcal{A}(P)$ and $b \in T_A\mathcal{A}(P)$). This expression can in turn be re-written in a suggestive form. Let $\xi_\mathcal{A} = -d_A\xi$ be the vector field generated by the gauge action. Then the equation above becomes

$$(\iota_{\xi_\mathcal{A}}\omega_\mathcal{A})(b) = \frac{d}{d\lambda}\bigg|_0 \left(-\int_{\partial\Sigma}\langle \xi, \lambda b\rangle + \int_\Sigma \langle \xi, F(A+\lambda b)\rangle\right), \tag{11}$$

the expression on the right hand side is the derivative of $\langle \mu, \xi\rangle$ in the direction b, where $\mu : \Omega^1(\Sigma)\otimes\mathfrak{g} \to Lie(Gau)^* = \mathfrak{g}^*(\Sigma)$ is defined by

$$\langle \mu(A), \xi\rangle = -\int_{\partial\Sigma}\langle \xi, A\rangle + \int_\Sigma \langle \xi, F(A)\rangle.$$

We see that (11) defines a moment map. It's zero level set $\mu^{-1}(0)$ is $\mathcal{A}_{col}^{flat}(\Sigma)$, and so we see that $\mathbf{M}_\Sigma = \mu^{-1}(0)/Gau_{col}$ is a symplectic quotient (the symplectic form is defined by the same formula on both spaces).

We remark that in the case of a surface with boundary (and no boundary conditions), in general this moment map is not equivariant. However in the case with Lagrangian Lie subalgebra boundary conditions it is; it's straightforward to compute the Poisson bracket

$$\{\langle\mu,\xi_1\rangle,\langle\mu,\xi_2\rangle\} = \langle\mu,[\xi_1,\xi_2]\rangle + \int_{\partial\Sigma}\langle\xi_1, d\xi_2\rangle,$$

and the boundary term vanishes when ξ_1, ξ_2 take values in the same Lagrangian Lie subalgebras along the edges of the boundary, so we see that μ is equivariant.

We've seen a definition of the symplectic form on $\mathbf{M}_\Sigma$ using the connection picture. Next we consider the symplectic structure in the holonomy picture.

4 Ševera's formula for the symplectic form

In his paper, Ševera describes a clever way of computing the symplectic form in the holonomy picture. The formula is similar to the one used to construct fusion products of quasi-Hamiltonian G spaces [1] (also see below). Let U be a manifold and let $G(U)$ be the collection of smooth maps $U \to G$. Pointwise multiplication turns $G(U)$ into a group.

4.1 Central extension by closed 2-forms

We now define a certain central extension of this group by $\Omega^2_{cl}(U)$, the closed 2-forms on U. The product looks strange at first but ends up doing exactly what we want. Let

$$\tilde{G}(U) := \{(g,t) \in G(U) \times \Omega^2(U) | dt = g^*\eta\},$$

where here η denotes the Cartan 3-form on G; our convention is

$$\eta = -\tfrac{1}{12}\langle[\theta^L,\theta^L],\theta^L\rangle.$$

This becomes a group when we define multiplication and inverses by

$$(g_1,t_1)(g_2,t_2) := (g_1g_2, t_1+t_2+\tfrac{1}{2}\langle g_1^*\theta^L, g_2^*\theta^R\rangle), \tag{12}$$

$$(g,t)^{-1} = (g^{-1},-t). \tag{13}$$

The proof that $\tilde{G}(U)$ is indeed closed under (12) follows from the key formula

$$Mult^*\eta = pr_1^*\eta + pr_2^*\eta + \tfrac{1}{2}d\langle pr_1^*\theta^L, pr_2^*\theta^R\rangle \tag{14}$$

where $Mult : G \times G \to G$ is the group multiplication, and pr_1, pr_2 are the projections onto the first and second factors respectively of the product $G \times G$. Indeed if we compute

$$\begin{aligned}(g_1 g_2)^*\eta &= (Mult \circ (g_1, g_2))^*\eta \\ &= (g_1^*, g_2^*) \circ Mult^*\eta \\ &= (g_1^*, g_2^*)(pr_1^*\eta + pr_2^*\eta + \tfrac{1}{2}d\langle pr_1^*\theta^L, pr_2^*\theta^R\rangle) \\ &= (pr_1 \circ (g_1, g_2))^*\eta + (pr_2 \circ (g_1, g_2))^*\eta + \tfrac{1}{2}d\langle (pr_1 \circ (g_1, g_2))^*\theta^L, (pr_2 \circ (g_1, g_2))^*\theta^R\rangle \\ &= g_1^*\eta + g_2^*\eta + \tfrac{1}{2}d\langle g_1^*\theta^L, g_2^*\theta^R\rangle \\ &= d\big(t_1 + t_2 + \tfrac{1}{2}\langle g_1^*\theta^L, g_2^*\theta^R\rangle\big)\end{aligned}$$

we get exactly the multiplication formula above. The equation (14) itself can be proved by a somewhat tedious but direct calculation.

Next we want to find the Lie algebra of $\tilde{G}(U)$. We claim that it is

$$\tilde{\mathfrak{g}}(U) := \mathfrak{g}(U) \oplus \Omega^2_{cl}(U),$$

with Lie bracket

$$[(x_1, s_1), (x_2, s_2)] := ([x_1, x_2], \langle dx_1, dx_2\rangle).$$

The Jacobi identity follows from the Jacobi identity for $\mathfrak{g}$ and the Ad-invariance of the inner product on $\mathfrak{g}$. Clearly the bracket vanishes on $\Omega^2_{cl}(U)$, so this is a central extension of the Lie algebra $\mathfrak{g}(U)$ (with pointwise bracket). We now want to verify that this is indeed the Lie algebra of $\tilde{G}(U)$.

Proposition 4.1. *The Lie algebra $\tilde{\mathfrak{g}}(U)$ is the Lie algebra of $\tilde{G}(U)$.*

Proof. Consider a curve $(e^{\lambda x}, S(\lambda))$ in $\tilde{G}(U)$ with $S(0) = 0$ (so that the curve passes through the identity), where here $x \in \mathfrak{g}(U)$, and for each λ, $S(\lambda)$ is a 2-form on U satisfying $dS(\lambda) = (e^{\lambda x})^*\eta$. Notice that

$$(e^{\lambda x})^*\theta^L\Big|_{\lambda=0} = id^*\theta^L = 0, \tag{15}$$

where $id : U \to G$ denotes the constant map $id(u) = e$ for all $u \in U$. Hence by the Leibniz rule

$$\frac{d}{d\lambda}\Big|_0 (e^{\lambda x})^*\eta = 0.$$

This shows that $d(\partial_\lambda S(\lambda)|_0) = 0$, and so the tangent space to $\tilde{G}(U)$ at the identity is indeed $\tilde{\mathfrak{g}}(U)$.

Recall that if G is any Lie group and $X \in \mathfrak{g} = Lie(G)$, the exponential map is defined to be the unique solution of the ODE

$$\begin{aligned}\frac{d}{d\lambda}\Big|_{\lambda_0} \exp(\lambda X) &= Tl_{\exp(\lambda_0 X)} \cdot X \\ &= \frac{d}{d\lambda}\Big|_0 \exp(\lambda_0 X) \cdot \exp(\lambda X)\end{aligned} \tag{16}$$

having initial condition $\exp(0) = 1$. We'll use $\exp(X)$ to denote the exponential map in $\tilde{G}(U)$ and reserve e^X for the exponential map in $G(U)$. Write

$$\exp(\lambda(x, s)) = (e^{\lambda x}, S(\lambda)) \tag{17}$$

where $S(0) = 0, S'(0) = s$. Putting (17) into the definition (16) and noting that

$$\frac{d}{d\lambda}\bigg|_0 (e^{\lambda x})^* \theta^L = \frac{d}{d\lambda}\bigg|_0 (e^{\lambda x})^* \theta^R = dx. \tag{18}$$

yields

$$S'(\lambda) = s + \tfrac{1}{2}\langle (e^{\lambda x})^* \theta^L, dx \rangle. \tag{19}$$

If desired we can integrate this

$$S(\lambda) = \lambda s + \tfrac{1}{2} \int_0^\lambda \langle (e^{\tau x})^* \theta^L, dx \rangle d\tau. \tag{20}$$

Now let's compute the bracket. Recall that for any Lie group G and $X, Y \in Lie(G)$ we have

$$[X, Y] = ad_X Y = \frac{d}{d\lambda}\bigg|_0 Ad_{\exp(\lambda X)} Y = \frac{d}{d\lambda}\bigg|_0 \frac{d}{d\lambda'}\bigg|_0 \big(\exp(\lambda X) \exp(\lambda' Y) \exp(-\lambda X)\big). \tag{21}$$

We apply this to the case at hand using the formula (12) for the product in $\tilde{G}(U)$. This gives

$$\begin{aligned}
[(x_1, s_1), (x_2, s_2)] &= \frac{d}{d\lambda}\bigg|_0 \frac{d}{d\lambda'}\bigg|_0 \big(e^{\lambda x_1}, S_1(\lambda)\big)\big(e^{\lambda' x_2}, S_2(\lambda')\big)\big(e^{-\lambda x_1}, -S_1(\lambda)\big) \\
&= \frac{d}{d\lambda}\bigg|_0 \frac{d}{d\lambda'}\bigg|_0 \Bigg(e^{\lambda x_1} e^{\lambda' x_2} e^{-\lambda x_1}, S_1(\lambda) + S_2(\lambda') - S_1(\lambda) \\
&\quad + \tfrac{1}{2}\langle (e^{\lambda x_1})^* \theta^L, (e^{\lambda' x_2})^* \theta^R \rangle + \tfrac{1}{2}\langle (e^{\lambda x_1} e^{\lambda' x_2})^* \theta^L, (e^{-\lambda x_1})^* \theta^R \rangle\Bigg) \\
&= \Bigg([x_1, x_2], \tfrac{1}{2} \frac{d}{d\lambda}\bigg|_0 \frac{d}{d\lambda'}\bigg|_0 \Big[\langle (e^{\lambda x_1})^* \theta^L, (e^{\lambda' x_2})^* \theta^R \rangle + \langle (e^{\lambda x_1} e^{\lambda' x_2})^* \theta^L, (e^{-\lambda x_1})^* \theta^R \rangle\Big]\Bigg).
\end{aligned}$$

Let l_g denote the map $G \to G$ given by left-multiplication by g. Recalling that for a map $g : U \to G$, $g^* \theta^L = T l_{g^{-1}} \circ Tg$ (which can be abbreviated $g^{-1} dg$), we have

$$(e^{\lambda x_1} e^{\lambda' x_2})^* \theta^L = Ad_{e^{-\lambda' x_2}} (e^{\lambda x_1})^* \theta^L + (e^{\lambda' x_2})^* \theta^L.$$

We apply this to the expression above:

$$\begin{aligned}
[(x_1, s_1), (x_2, s_2)] = \Bigg([x_1, x_2], \tfrac{1}{2} \frac{d}{d\lambda}\bigg|_0 \frac{d}{d\lambda'}\bigg|_0 \Big[&\langle (e^{\lambda x_1})^* \theta^L, (e^{\lambda' x_2})^* \theta^R \rangle \\
&+ \langle Ad_{e^{-\lambda' x_2}} (e^{\lambda x_1})^* \theta^L + (e^{\lambda' x_2})^* \theta^L, (e^{-\lambda x_1})^* \theta^R \rangle\Big]\Bigg).
\end{aligned}$$

To compute the derivatives, we make use of (15) and (18). Most of the terms drop out and we're left with

$$\begin{aligned}[(x_1, s_1), (x_2, s_2)] &= \big([x_1, x_2], \tfrac{1}{2}\langle dx_1, dx_2\rangle - \tfrac{1}{2}\langle dx_2, dx_1\rangle\big) \\ &= \big([x_1, x_2], \langle dx_1, dx_2\rangle\big)\end{aligned}$$

as required. □

Ševera's formula is simplest when expressed as a product in a bigger group which Ševera calls $\tilde{G}_{big}(U)$ defined as

$$\tilde{G}_{big}(U) := G(U) \times \Omega^2(U),$$

with the same formula for the product and inverse as $\tilde{G}(U)$.

Recall that by cutting Σ we get a polygon with labelled edges, γ_s is the label on edge s (s runs cyclically around the polygon). Since each edge connects two vertices in Σ, as explained before we can think of γ_s as a function $\gamma_s : u \in \mathbf{M}_\Sigma \mapsto hol(A_u, s) \in G$ where A_u is a connection in the equivalence class $u \in \mathbf{M}_\Sigma$. Let's now state Ševera's formula for the symplectic form on $\mathbf{M}_\Sigma$:

$$(e, \omega) = \prod_s (\gamma_s, 0), \tag{22}$$

where the product is taken in the group $\tilde{G}_{big}(\mathbf{M}_\Sigma)$.

4.2 Proof of the formula

Ševera gives an interesting proof that this formula indeed gives the Atiyah-Bott symplectic form on the moduli space. Although it relies on formal arguments involving connections on infinite dimensional principal bundles (and the analytic details are not filled in), it gives valuable intuition as to why the formula works. The essential idea is that the formula should be viewed as an instance of "Stokes theorem"–the symplectic form being the integral of a "curl" (in this case, the curvature) and the product of the holonomies along the edges of the boundary being the "circulation".

In general since the curvature $F = dA + \frac{1}{2}[A, A]$ has a non-linear term, we can't think of it as the "curl" of a "vector field" A whose "circulation" around the boundary is equal to its holonomy. But Ševera begins by discussing a situation in which we can. We consider a central extension of Lie groups

$$C \to \tilde{K} \to K,$$

a disc D and a principal $\tilde{K}$-bundle $\tilde{P} \to D$. Then $P := \tilde{P}/C$ is a principal K-bundle over D. Let A be a flat connection on P and $\tilde{A}$ any connection on $\tilde{P}$ lifting A. As D is simply connected, without loss of generality we can take $\tilde{P} = D \times \tilde{K}$, $P = D \times K$, we have global coordinates, and we can take A to be the trivial product connection on P. Now $\tilde{A}$ need not be flat, but in order to induce the flat connection A on the quotient, it's curvature must take values in $\mathfrak{c}$ the Lie algebra of C (this makes sense because C is in the centre of $\tilde{K}$, and in particular $\mathfrak{c}$ is fixed under the adjoint action of $\tilde{K}$). But $\mathfrak{c}$ is abelian, and so in this case the non-linear term vanishes and we have

$$\tilde{F} = d\tilde{A}$$

and hence by Stokes' theorem, together with the fact that we may think of $\tilde{P}$ as a product $D \times \tilde{K}$ (so can use global coordinates regarding $\tilde{A}$ as a $\mathfrak{c}$-valued 1-form on D) and the definition of holonomy we have

$$\begin{aligned} \exp\left(\int_D \tilde{F}\right) &= \exp\left(\int_{\partial D} \tilde{A}\right) \\ &= hol_{\partial D}\tilde{A}. \end{aligned}$$

Ševera applies this general result (in an infinite-dimensional context) to the central extension

$$\Omega^2_{cl}(U) \to \tilde{G}(U) \to G(U),$$

where U will denote open subsets (coordinate charts) of the moduli space $\mathbf{M}_\Sigma$. The disc D will be the result of cutting Σ to get a polygon (so the holonomy around its boundary will be the product of the holonomies along the edges).

The space $\mathbf{M}_\Sigma$ is covered by "coordinate charts" U, ψ with $\psi : U \to \mathcal{A}^{flat}_{col}(P)$ such that $\psi(u)$ is a connection in the equivalence class u for each $u \in U$, and here P is a colored principal G-bundle over Σ (so U smoothly parametrizes a collection of flat colored connections living on the same bundle $P \to \Sigma$). Recall that the 2-form $\omega_\mathcal{A}$ on $\mathcal{A}^{flat}_{col}(\Sigma)$ gave rise to a symplectic form ω on $\mathbf{M}_\Sigma$. Since $\omega_\mathcal{A}$ is constant on gauge orbits and vanishes on gauge directions, we have

$$\omega = \psi^*\omega_\mathcal{A}. \tag{23}$$

Note that we have inclusions $G \hookrightarrow G(U)$ and $G \hookrightarrow \tilde{G}(U)$, which just means that we can think of an element $g \in G$ as a constant G-valued function on U (with zero 2-form part). This means that P induces $G(U)$ and $\tilde{G}(U)$ principal bundles over Σ (for example, we use the same transition functions). We denote these bundles P_U and $\tilde{P}_U$ respectively, and we have that $P_U \simeq \tilde{P}_U/\Omega^2_{cl}(U)$. $G(U)$ and $\tilde{G}(U)$ are infinite dimensional. Following Ševera's paper we treat these formally: we think of a section of a $G(U)$ bundle over Σ as a being given locally by a smooth function $\sigma : V \subset \Sigma \to G(U)$, or equivalently as a family of smooth functions $\sigma(u) : V \to G$ smoothly parametrized by $u \in U$. Likewise we think of a connection on P_U as being given locally (on $V \subset \Sigma$) by 1-forms $A \in \Omega^1(V) \otimes \mathfrak{g}(U)$, or equivalently as a family of smooth 1-forms $A(u) \in \Omega^1(V) \otimes \mathfrak{g}$ smoothly parametrized by $u \in U$.

We can use our chart $\psi : U \to \mathcal{A}^{flat}_{col}(P)$ to define a family of 1-forms smoothly parametrized by U. Let $\phi : V \subset \Sigma \to \pi^{-1}(V) \subset P$ be a local section, and define

$$A(u) = \phi^*\psi(u) \tag{24}$$

Each $A(u)$ is a $\mathfrak{g}$-valued 1-form on V, so we have a family $A(u) \in \Omega^1(V) \otimes \mathfrak{g}$ smoothly parametrized by $u \in U$–by definition a connection on P_U at least defined locally over $V \subset \Sigma$. Because the $A(u)$ are just pull-backs of connections, if we have two coordinate patches V_1, ϕ_1 and V_2, ϕ_2 the corresponding 1-forms $A_1(u)$ and $A_2(u)$ will obey a compatibility condition on overlaps

$$A_2(u) = Ad_{t_{12}} A_1(u) - t_{12}^* \theta^R,$$

where $t_{12} : V_1 \cap V_2 \to G$ is the transition function. Since the bundle P_U has the same transition functions as P, this condition (taken for all $u \in U$ simultaneously) is exactly the compatibility condition for 1-forms defining a connection on P_U. Thus the equation (24) defines a connection on P_U, which we denote by A. On local patches V in Σ it is just the family of $\mathfrak{g}$-valued 1-forms $A(u)$ smoothly parametrized by $u \in U$.

Furthermore A is flat! This is just because it is "composed" of a collection of flat connections. Explicitly, computing locally in $V \subset \Sigma$ we have

$$\begin{aligned} F = F(u) &= dA(u) + \tfrac{1}{2}[A(u), A(u)] \\ &= 0 \end{aligned}$$

because for fixed u, $A(u)$ is flat (notice that the d which appears is with respect to Σ, so that the u-dependence plays no role).

We can explain this in a slightly different way which sheds additional light on what is going on. Locally we think of P over $V \subset \Sigma$ as the product $V \times G$, and we think of P_U as the product $V \times G(U)$. So (locally) we can think of a point p in P_U as a smooth function $p : U \to G$; its evaluation at a point $u \in U$ yields an element $p(u) \in G$, which we may think of as a point in P (all over the same base point $x \in V \subset \Sigma$). Hence with local charts, we are free to think of a point p of P_U as representing a collection of points $p(u) \in P$ smoothly parametrized by U. Now if we have a curve $f(\lambda), \lambda \in [0, 1]$ in V and a point $p = p(u) \in P_U$, what is the horizontal lift to p? If we let $\tilde{f}_u(\lambda)$ denote the $A(u)$ lift of f to $p(u) \in P$, then the lift is the curve $\tilde{f}(\lambda) := \tilde{f}_u(\lambda)$, where (separately for each λ) we are using the identification of points of P_U with collections of points in P smoothly parametrized by U. The flatness of a connection may be characterized in terms of its horizontal lifts: a connection is flat iff whenever f is a smooth closed curve homotopic to a point, it's horizontal lift closes. For fixed u, $A(u)$ is flat and therefore if f is closed and homotopic to a point then the curve $\tilde{f}_u$ is closed. That is, $p(u) = \tilde{f}_u(0) = \tilde{f}_u(1)$ and since this holds for all u we get $p = \tilde{f}(0) = \tilde{f}(1)$, so the lift closes. Thus we see that A is indeed flat.

A connection on $\tilde{P}_U$ can be thought of locally as a $\tilde{\mathfrak{g}}(U) = \mathfrak{g}(U) \oplus \Omega^2_{cl}(U)$-valued 1-form. Hence $\tilde{A} = (A, 0)$, or more explicitly $\tilde{A}(u) = (A(u), 0)$, defines a connection on $\tilde{P}_U$ which is a lift of the connection A on P_U. However as discussed before, this connection need not be flat. Let's compute the curvature. By definition

$$\tilde{F} = d\tilde{A} + \tfrac{1}{2}[\tilde{A}, \tilde{A}],$$

where here the exterior derivative is with respect to Σ. Now recall that $[\tilde{A}, \tilde{A}]$ means that we take the bracket of the Lie algebra parts and wedge the 1-form parts, the result being independent of the decomposition since both operations are bilinear. So write A locally as

$$A = \sum_{\mu} A_\mu dx^\mu$$

where x^μ are coordinates on Σ. Then

$$\begin{aligned}
\tilde{F} &= (dA, 0) + \tfrac{1}{2}[(A,0),(A,0)] \\
&= (dA, 0) + \tfrac{1}{2}\sum [(A_\mu, 0), (A_\nu, 0)] dx^\mu \wedge dx^\nu \\
&= (dA, 0) + \tfrac{1}{2}\sum ([A_\mu, A_\nu], \langle d^U A_\mu, d^U A_\nu\rangle) dx^\mu \wedge dx^\nu \\
&= (dA + \tfrac{1}{2}[A,A], \tfrac{1}{2}\langle d^U A, d^U A\rangle) \\
&= (0, \tfrac{1}{2}\langle d^U A, d^U A\rangle).
\end{aligned}$$

We emphasize that the exterior derivatives and wedges happen both on Σ and on U and we have added a superscript U to emphasize this. A wedge on U is implicit in the inner product in lines 3,4,5 (it, as well as the d^U, have come from the formula for the Lie bracket in $\tilde{\mathfrak{g}}(U)$), and in lines 4,5 the inner product implies a wedge on both Σ and U. The resulting form has zero $\mathfrak{g}(U)$ part and so lies in $\Omega^2(\Sigma) \otimes \Omega^2_{cl}(U)$. Hence we can think of it as an $\Omega^2(\Sigma)$-valued 2-form on U. What happens when we evaluate it on two vectors $\alpha, \beta \in T_u U$? Well $d^U A$ is just the tangent map TA for the map $A : u \in U \mapsto A(u) = \phi^* \psi(u) \in \Omega^2(V) \otimes \mathfrak{g}$. Hence $d^U A(\alpha) = TA \cdot \alpha$, so we end up getting

$$\tilde{F}(\alpha, \beta) = \langle TA \cdot \alpha, TA \cdot \beta \rangle.$$

Here to make the formula less cluttered, we've not written the $\mathfrak{g}(U)$ part (which is 0). Once the two $\Omega^2_{cl}(U)$ slots have been used, we're left with a 2-form on $V \subset \Sigma$. We do this in each coordinate patch (V_i, ϕ_i), writing A_i for the map $\phi_i^* \circ \psi$ on U. We can then use a partition of unity ρ_i to put them together and integrate,

$$\begin{aligned}
\int_\Sigma \tilde{F}(\alpha, \beta) &= \sum \rho_i \int_{V_i} \langle TA_i \cdot \alpha, TA_i \cdot \beta \rangle \\
&= \sum \rho_i \int_{V_i} \phi_i^* \langle T\psi \cdot \alpha, T\psi \cdot \beta \rangle \\
&= \omega_{\mathcal{A}}(T\psi \cdot \alpha, T\psi \cdot \beta) \\
&= \psi^* \omega_{\mathcal{A}}(\alpha, \beta) \\
&= \omega(\alpha, \beta),
\end{aligned}$$

where we've used the definitions of ω, $\omega_{\mathcal{A}}$. So we've verified that

$$(0, \omega) = \int_\Sigma \tilde{F}. \tag{25}$$

(Here we've put back in the $\tilde{\mathfrak{g}}(U)$ part, which is zero.)

As Ševera argues, his formula for the symplectic form on $\mathbf{M}_\Sigma$ follows quickly from (25). Cut Σ into a polygon D with side s labelled by the holonomy $\gamma_s : \mathbf{M}_\Sigma \to G$. Then as $\tilde{G}(U)$ is a central extension of $G(U)$ we have

$$\exp\left(\int_\Sigma \tilde{F}\right) = hol_{\partial D}\tilde{A}.$$

By (25), the left hand side is just $\exp((0,\omega)) = (e,\omega)$ (computing the exponential in $\tilde{G}(U)$ using (20)). We would like to write the expression on the right as a product of holonomies along each of the edges. Recall by this we mean that we are using the special points picked out in the fibres above the vertices of the polygon D to give us a trivialization of the fibres above these vertices, and then taking the holonomy relative to these trivializations. Now the holonomy of $\tilde{A}$ along side s is $(\gamma_s, t_s) \in \tilde{G}(U)$ where $\gamma_s(u) = hol(A(u), s)$, and t_s is some 2-form on U satisfying $dt_s = \gamma_s^*\eta$. To compute t_s we use $\tilde{A} = (A, 0)$ together with the formula (20) for the exponential in $\tilde{G}(U)$,

$$\begin{aligned} hol(s, \tilde{A}(u)) &= \exp\left(\int_s \tilde{A}(u)\right) \\ &= \exp\left\{\left(\int_s A(u), 0\right)\right\} \\ &= \left(e^{\int_s A(u)}, \tfrac{1}{2}\int_0^1 d\tau\ \langle \gamma_s(\tau)^*\theta^L, d^U I_s\rangle_u\right) \\ &= (\gamma_s(u), t_s(u)), \end{aligned}$$

where we've put

$$\begin{aligned} I_s(u) &= \int_s A(u) && \text{(a function on } U) \\ \gamma_s(\tau, u) &= e^{\tau I_s(u)} && \text{(note } \gamma_s(1,u) = \gamma_s(u)) \end{aligned}$$

for convenience; $\gamma_s(\tau)$ denotes the function $u \mapsto \gamma_s(\tau, u)$ and the subscript u on the inner product denotes evaluation at $u \in U$.

Now there are two cases to consider. If $s = a$ is in $\partial\Sigma$ (i.e. it comes from an edge a and not a cut), then for each u the connection $A(u)$ is colored and hence takes values in the Lagrangian Lie subalgebra $\mathfrak{h}_a$ along the edge a. Hence for each u, $I_a(u) \in \mathfrak{h}_a$. It follows that $d^U I_a$ and $\gamma_a(\tau)^*\theta^L$ are both $\mathfrak{h}_a$-valued for all τ. The inner product vanishes when restricted to $\mathfrak{h}_a$ and so we see that in this case $t_s(u)$ vanishes for every $u \in U$. The second case is when s arises from a cut. In this case $t_s(u)$ need not be zero. However since s arises from a cut, there is another edge s' which came from the same cut, and so the holonomy of $\tilde{A}(u)$ along it must be exactly the inverse, that is $(\gamma_s'(u), t_s'(u)) = (\gamma_s(u), t_s(u))^{-1} = (\gamma_s(u)^{-1}, -t_s(u))$. In the product, the 2-form parts get added together and cancel. So we may as well drop the $t_s(u)$ part for the purposes of evaluating the product appearing in Ševera's formula (it is for this reason that we want to use $\tilde{G}_{big}(U)$ to compute the product, since when we replace

$(\gamma_s(u), t_s(u))$ with $(\gamma_s(u), 0)$ we leave $\tilde{G}(U)$). The product of the holonomies along the edges, taken cyclically, gives the holonomy around ∂D. We've thus shown that

$$(e, \omega) = \prod_s (\gamma_s, 0),$$

(each side to be evaluated at a point $u \in U \subset \mathbf{M}_\Sigma$). Of course the product of all the γ_s must be e since going all the way around the polygon gives a closed loop homotopic to a point. So the interesting part of the formula is the 2-form part of the resulting product. It is a 2-form on $U \subset M_\Sigma$ and we've shown that it is exactly the Atiyah-Bott 2-form.

5 Group-valued moment maps

Group-valued moment maps were first defined in [1] and provided a method of constructing the Atiyah-Bott symplectic form on the moduli space of flat connections on a compact Riemann surface. One of the advantages gained was to circumvent the infinite-dimensional symplectic reduction. Above we saw that in the case of a surface with boundary conditions given by Lagrangian Lie subalgebras, again we find a finite-dimensional moduli space after an infinite-dimensional reduction. And moreover we have a nice formula (22) for the symplectic form in terms of holonomies. This suggests that we should be able to construct Ševera's moduli spaces using the group-valued moment map approach. This is not solved yet (the remaining difficult part is to prove "minimal degeneracy"), although there is a candidate quasi-Hamiltonian space (see below). We'll review the definition and basic properties of quasi-Hamiltonian G spaces, and briefly discuss how Ševera's moduli spaces might be constructed this way.

As above, let G be a Lie group with Lie algebra $\mathfrak{g}$ having an invariant inner product $\langle , \rangle$. Recall that a *quasi-Hamiltonian G-space* is a G-manifold M carrying an invariant 2-form ω and an equivariant map $\Phi : M \to G$, called the *G-valued moment map* satisfying

1. $\iota(\xi_M)\omega = -\frac{1}{2}\langle \Phi^*(\theta^L + \theta^R), \xi \rangle$
2. $d\omega = \Phi^*\eta$
3. $ker(\omega_m) = \{\xi_M(m) | Ad_{\Phi(m)}\xi = -\xi\}$

where here ξ_M denotes the vector field on M generated by $\xi \in \mathfrak{g}$. The first condition is called the moment map condition, and the third is called the minimal degeneracy condition (minimal because from the moment map condition, ω must vanish on vectors appearing in condition 3). Conditions 1 and 2 may be combined into a single statement:

$$d_G\omega = \Phi^*\eta_G$$

where $d_G = d + \iota(\xi_M)$ is the equivariant differential and $\eta_G(\xi) = -\frac{1}{2}\langle \theta^L + \theta^R, \xi \rangle + \eta$ is an equivariant extension of the Cartan 3-form. These conditions are analogous to those defining an ordinary $\mathfrak{g}^*$-valued moment map, and indeed quasi-Hamiltonian G-spaces have many properties analogous to those of Hamiltonian G-spaces.

We state here two basic properties of quasi-Hamiltonian G-spaces. See [1] for more information and proofs.

Theorem 5.1. *Let* $(M, \omega, \Phi = (\Phi_1, \Phi_2))$ *be a quasi-Hamiltonian* $G \times G$*-space [in particular, if* $(M_1, \omega_1, \Phi_1), (M_2, \omega_2, \Phi_2)$ *are quasi-Hamiltonian, then* $(M_1 \times M_2, \omega_1 + \omega_2, (\Phi_1, \Phi_2))$ *is such a space]. Let*

$$\begin{aligned} \Phi^{fus} &= \Phi_1 \Phi_2 \\ \omega^{fus} &= \omega + \tfrac{1}{2}\langle \Phi_1^* \theta^L, \Phi_2^* \theta^L \rangle. \end{aligned}$$

Then $(M, \omega^{fus}, \Phi^{fus})$ *is a quasi-Hamiltonian* G*-space with the diagonal action.*

The theorem above allows one to take products of quasi-Hamiltonian spaces.

Theorem 5.2. *Let* (M, ω, Φ) *be a quasi-Hamiltonian* G*-space with proper moment map and proper* G*-action. Also suppose the identity* e *is a regular value of* Φ*. Then the reduced space*

$$M//G := \Phi^{-1}(e)/G$$

is a symplectic orbifold (the symplectic structure is induced by ω*).*

A motivating example was a conjugacy class $\mathcal{C} \subset G$ with the adjoint action, momentum map the inclusion (compare to the Kirillov-Kostant-Souriau symplectic structure on coadjoint orbits) and 2-form

$$\omega_g(\xi_1^\#, \xi_2^\#) = \tfrac{1}{2}\langle (Ad_g - Ad_{g^{-1}})\xi_1, \xi_2 \rangle.$$

A second motivating example was the double $D(G) = G \times G$ which is a quasi-Hamiltonian $G \times G$ space (in fact it is a fusion of a certain conjugacy class with itself! see e.g. [4]). If we fuse this example, we find that the fused double $\tilde{D}(G) := G \times G$ is a q-Hamiltonian G-space with conjugation action, and the moment map turns out to be

$$\Phi(a, b) = aba^{-1}b^{-1}.$$

And more generally fusing h copies of $\tilde{D}(G)$ yields G^{2h} as a q-Hamiltonian G-space with moment map

$$\Phi(a_1, b_1, ..., a_h, b_h) = \prod_{k=1}^{h} a_i b_i a_i^{-1} b_i^{-1}.$$

The reduced space $\Phi^{-1}(e)/G$ is exactly the moduli space of flat connections on a compact Riemann surface of genus h, and the symplectic form coming from the q-Hamiltonian reduction agrees with the Atiyah-Bott symplectic form.

5.1 Application to Ševera's moduli spaces?

Here we'll describe a candidate for a quasi-Hamiltonian space that reduces to one of Ševera's moduli spaces. The moduli space of Ševera's that we have in mind is an $n - gon$ labelled by Lagrangian Lie subalgebras $\mathfrak{h}_1, ..., \mathfrak{h}_n$ (note that n must be at least 2; for $n = 2$ we get a single point, more interesting examples start at $n = 3$). Let $H_1, ..., H_n$ be the corresponding connected Lie subgroups of G. For convenience put $N = n + 2$. Define a map

$$\Psi : G^N \to G$$
$$(g_0, ..., g_{n+1}) \mapsto \prod_{k=0}^{n+1} g_k.$$

Let M be the submanifold of G^N defined by

$$\begin{aligned} M &= \{(g_0, ..., g_{n+1}) \in G^N | g_k \in H_k,\ 1 \le k \le n;\ g_0 = g_{n+1}^{-1}\} \\ &= \{(h_1, ..., h_n, g) \in H_1 \times \cdots \times H_n \times G\}. \end{aligned}$$

(In the second description, put $h_k = g_k,\ 1 \le k \le n;\ g = g_{n+1} = g_0^{-1}$.) M is our candidate for a q-Hamiltonian G-space whose reduction will be Ševera's moduli space (for an n-gon). Let $i : M \hookrightarrow G^N$ denote the inclusion and $\Phi := \Psi \circ i$, so

$$\Phi(h_1, ..., h_n, g) = g^{-1} h_1 \cdots h_n g.$$

There is a G-action on M for which Φ is equivariant

$$x \in G : (h_1, ..., h_n, g) \mapsto (h_1, ..., h_n, gx^{-1})$$

which is induced from the action $x \in G : (g_0, ..., g_{n+1}) \mapsto (xg_0, g_1, ..., g_n, g_{n+1}x^{-1})$ on G^N.

The 2-form suggested by Ševera's formula (22) is $\tau^{(N)}$ defined by

$$\prod_{k=0}^{n+1} (g_k, 0) = \left(\prod_{k=0}^{n+1} g_k, \tau^{(N)} \right).$$

We also define $\tau^{(n)}$:

$$\prod_{k=1}^{n} (g_k, 0) = \left(\prod_{k=1}^{n} g_k, \tau^{(n)} \right),$$

which, when pulled back to the submanifold $\{e\} \times H_1 \times \cdots \times H_n \times \{e\} \subset G^N$, is (by Ševera's formula) exactly the symplectic form for Ševera's n-gon moduli space, i.e. if we let $j : \Phi^{-1}(e) \hookrightarrow G^N$ be the inclusion then $j^* \tau^{(n)}$ is the desired symplectic form on $\Phi^{-1}(e)$.

We will show two things,

1. $\tau^{(N)}$ is G-invariant and $(M, i^* \tau^{(N)}, \Phi)$ satisfies the first two properties of a q-Hamiltonian G-space:

(a) $\iota(\xi_M) i^* \tau^{(N)} = -\frac{1}{2}\langle \Phi^*(\theta^L + \theta^R), \xi \rangle$, $\xi \in \mathfrak{g}$

(b) $di^*\tau^{(N)} = \Phi^*\eta$

2. $j^*\tau^{(N)} = j^*\tau^{(n)}$, i.e. the symplectic structure on Ševera's moduli space is what it should be if it's coming from q-Hamiltonian reduction of M.

These are fairly straightforward to show; the missing piece is the minimal degeneracy condition (which we've attempted to show but without success so far).

We start by looking at a trick which speeds up the proof. Recall that

$$Mult^*\eta = pr_1^*\eta + pr_2^*\eta + d\sigma, \qquad \sigma := \tfrac{1}{2}\langle pr_1^*\theta^L, pr_2^*\theta^R \rangle$$

Suppose $(\phi_1, \phi_2) : G^N \to G \times G$ are smooth maps. Then we can pull this equation back by (ϕ_1, ϕ_2), giving

$$(\phi_1\phi_2)^*\eta = \phi_1^*\eta + \phi_2^*\eta + \tfrac{1}{2} d\langle \phi_1^*\theta^L, \phi_2^*\theta^R \rangle.$$

Or re-arranging,

$$\tfrac{1}{2} d\langle \phi_1^*\theta^L, \phi_2^*\theta^R \rangle = (\phi_1\phi_2)^*\eta - \phi_1^*\eta - \phi_2^*\eta. \tag{26}$$

Notice in particular that if points of G^N have components labelled $(g_0, ..., g_{N-1})$ then each g_i (and any product thereof) is a function on G^N and so the above equation applies. We will use this formula to compute $d\tau^{(N)}$ inductively.

Proposition 5.1.

$$d\tau^{(N)} = (g_0 \cdots g_{N-1})^*\eta - \sum_{k=0}^{N-1} g_k^*\eta$$

Proof. The base case follows directly from (26),

$$\tau^{(2)} = \tfrac{1}{2}\langle g_0^*\theta^L, g_1^*\theta^R \rangle$$
$$\Rightarrow d\tau^{(2)} = (g_0 g_1)^*\eta - g_0^*\eta - g_1^*\eta.$$

Now for the inductive case: suppose

$$d\tau^{(N-1)} = (g_0 \cdots g_{N-2})^*\eta - \sum_{k=0}^{N-2} g_k^*\eta. \tag{27}$$

From Ševera's formula we have that

$$\tau^{(N)} = \tau^{(N-1)} + \tfrac{1}{2}\langle (g_0 \cdots g_{N-2})^*\theta^L, g_{N-1}^*\theta^R \rangle$$
$$\Rightarrow d\tau^{(N)} = (g_0 \cdots g_{N-2})^*\eta - \sum_{k=0}^{N-2} g_k^*\eta + \tfrac{1}{2} d\langle (g_0 \cdots g_{N-2})^*\theta^L, g_{N-1}^*\theta^R \rangle,$$

where we've used the induction hypothesis. We can calculate the last term using (26) again (for the maps $\phi_1 = g_0 \cdots g_{N-2}$ and $\phi_2 = g_{N-1}$) giving

$$
\begin{aligned}
d\tau^{(N)} &= (g_0 \cdots g_{N-2})^* \eta - \sum_{k=0}^{N-2} g_k^* \eta + (g_0 \cdots g_{N-1})^* \eta - (g_0 \cdots g_{N-2})^* \eta - g_{N-1}^* \eta \\
&= (g_0 \cdots g_{N-1})^* \eta - \sum_{k=0}^{N-1} g_k^* \eta.
\end{aligned}
$$

□

Using the map Ψ we can write this equation as

$$
d\tau^{(N)} = \Psi^* \eta - \sum_{k=0}^{N-1} g_k^* \eta. \tag{28}
$$

Now we pull this equation back to M using i. Notice that on M, $g_0 = g_{N-1}^{-1}$ so that restricted to M, $g_0^* \eta = (g_{N-1}^{-1})^* \eta = -g_{N-1}^* \eta$ and so these two terms cancel in the sum in (28) when that equation is restricted to M. On the other hand on M, $g_k \in H_k$ for $k = 1, ..., n = N-2$, so that $i^* g_k^* \theta^L$ is $\mathfrak{h}_k$-valued. Since $\eta = -\frac{1}{12}\langle [\theta^L, \theta^L], \theta^L \rangle$ and $\mathfrak{h}_k$ is Lagrangian, this means that $i^* g_k^* \eta = 0$ for $k = 1, ..., n$. So pulling back, (28) becomes

$$
d i^* \tau^{(N)} = i^* \Psi^* \eta = \Phi^* \eta,
$$

proving one of the q-Hamiltonian conditions.

Now let's show that the moment map condition holds. For convenience we define

$$
\varphi(g_0, ..., g_{n+1}) = \prod_{k=1}^{n} g_k.
$$

Now compute $\tau^{(N)}$,

$$
\begin{aligned}
\prod_{k=0}^{n+1} (g_k, 0) &= (g_0, 0) \left(\prod_{k=1}^{n} (g_k, 0) \right) (g_{n+1}, 0) \\
&= (g_0, 0)(\varphi, \tau^{(n)})(g_{n+1}, 0) \\
&= (\Psi, \tau^{(n)} + \tfrac{1}{2}\langle g_0^* \theta^L, \varphi^* \theta^R \rangle + \tfrac{1}{2}\langle (g_0 \varphi)^* \theta^L, g_{n+1}^* \theta^R \rangle) \\
\Rightarrow \tau^{(N)} &= \tau^{(n)} + \tfrac{1}{2}\langle g_0^* \theta^L, \varphi^* \theta^R \rangle + \tfrac{1}{2}\langle (g_0 \varphi)^* \theta^L, g_{n+1}^* \theta^R \rangle.
\end{aligned}
$$

Now $(g_0 \varphi)^* \theta^L = Ad_{\varphi^{-1}} g_0^* \theta^L + \varphi^* \theta^L$, which we use in the expression above to get the useful formula

$$
\tau^{(N)} = \tau^{(n)} + \tfrac{1}{2}\langle g_0^* \theta^L, \varphi^* \theta^R \rangle + \tfrac{1}{2}\langle Ad_{\varphi^{-1}} g_0^* \theta^L, g_{n+1}^* \theta^R \rangle + \tfrac{1}{2}\langle \varphi^* \theta^L, g_{n+1}^* \theta^R \rangle. \tag{29}
$$

In this form it is especially easy to see that $\tau^{(N)}$ is invariant under the G-action: recall the action left $g_1, ..., g_n$ fixed, while acting on g_0 on the left, and on g_{n+1} on the right. Thus the G-action preserves φ, $\tau^{(n)}$. And since in the above formula g_0 (resp. g_{n+1}) only appears in the form $g_0^*\theta^L$ (resp. $g_{n+1}^*\theta^R$), the left (resp. right) invariance of the left (resp. right) Maurer-Cartan form ensures that $\tau^{(N)}$ is invariant.

Let $\xi \in \mathfrak{g}$. The vector field generated by the action on G^N is

$$\xi^{\#}(g_0, ..., g_{n+1}) = (-\xi g_0, 0, ..., 0, g_{n+1}\xi),$$

(we write ξg_0 as an abbreviation for $Tr_{g_0} \cdot \xi$ where r_{g_0} denotes the right action, etc.). In particular since $T\varphi$ vanishes on the 0th and (n+1)st factors (φ does not depend on those variables) we have $\iota(\xi^{\#})\varphi^*\theta^L = \iota(\xi^{\#})\varphi^*\theta^R = 0$. Similarly $\tau^{(n)}$ also vanishes on the 0th and (n+1)st factors, so $\iota(\xi^{\#})\tau^{(n)} = 0$. Hence

$$-\iota(\xi^{\#})\tau^{(N)} = \tfrac{1}{2}\langle Ad_{g_0^{-1}}\xi, \varphi^*\theta^R\rangle + \tfrac{1}{2}\langle Ad_{\varphi^{-1}}Ad_{g_0^{-1}}\xi, g_{n+1}^*\theta^R\rangle + \tfrac{1}{2}\langle Ad_{\varphi^{-1}}g_0^*\theta^L, Ad_{g_{n+1}}\xi\rangle + \tfrac{1}{2}\langle \varphi^*\theta^L, Ad_{g_{n+1}}\xi\rangle.$$

On M we have $g_0^{-1} = g_{n+1} =: g$ and this expression simplifies to

$$-\iota(\xi_M)i^*\tau^{(N)} = \tfrac{1}{2}\langle Ad_g\xi, \varphi^*(\theta^L + \theta^R) + Ad_\varphi g^*\theta^R + Ad_{\varphi^{-1}}(g^{-1})^*\theta^L\rangle.$$

On the other hand, using $(g^{-1})^*\theta^R = -g^*\theta^L$ we find

$$\begin{aligned}
\tfrac{1}{2}\langle \Phi^*(\theta^L + \theta^R), \xi\rangle &= \tfrac{1}{2}\langle \xi, (g^{-1}\varphi g)^*(\theta^L + \theta^R)\rangle \\
&= \tfrac{1}{2}\langle \xi, Ad_{g^{-1}}Ad_{\varphi^{-1}}(g^{-1})^*\theta^L + Ad_{g^{-1}}\varphi^*\theta^L + g^*\theta^L \\
&\quad - g^*\theta^L + Ad_{g^{-1}}\varphi^*\theta^R + Ad_{g^{-1}}Ad_\varphi g^*\theta^R\rangle \\
&= \tfrac{1}{2}\langle Ad_g\xi, Ad_{\varphi^{-1}}(g^{-1})^*\theta^L + \varphi^*\theta^L + \varphi^*\theta^R + Ad_\varphi g^*\theta^R\rangle,
\end{aligned}$$

in agreement with the previous expression. This shows that the moment map condition is satisfied.

Finally we would like to show that when we pull back to $\Phi^{-1}(e)$ we get Ševera's form. First note that since $\Phi = g^{-1}\varphi g$, on $\Phi^{-1}(e)$ we have $\varphi = e$ is constant. In particular $j^*\varphi^*\theta^L = j^*\varphi^*\theta^R = 0$. So if we pull back (29) by j we get

$$j^*\tau^{(N)} = j^*\tau^{(n)} + \tfrac{1}{2}\langle (g^{-1})^*\theta^L, g^*\theta^R\rangle.$$

Finally notice that $(g^{-1})^*\theta^L = -g^*\theta^R$ and $\langle g^*\theta^R, g^*\theta^R\rangle = 0$ (the wedge is antisymmetric while the inner product is symmetric), so the second term cancels and we're left with

$$j^*\tau^{(N)} = j^*\tau^{(n)}.$$

To close, some brief comments on addressing the minimal degeneracy property. One approach attempted is probably the most direct, i.e. using Ševera's formula to write down an expression for ω and then trying to find it's kernel directly. This seems to lead to complicated formulas and it is not clear how to proceed (even for the smallest n which leads to a non-trivial moduli space). Another approach would be to try to construct M using already-known q-Hamiltonian spaces, for example, a clever choice of conjugacy classes (for example, the double $D(G)$ can be constructed this way), together with the fusion product. Yet another possible approach is through Dirac geometry, using the elegant result of Bursztyn-Crainic that a quasi-Hamiltonian $\mathfrak{g}$-structure on M is equivalent to a strong (forward) Dirac map $\Phi : M \to G$, where G carries the Cartan-Dirac structure (twisted by η)–see [4] for more information and references. One way of proceeding using the set-up above would be to try to work with the fact that $\Psi : G^N \to G$ is a strong Dirac map (again see [4]); at least for initial attempts this runs into trouble though, since it seems that the inclusion $i : M \hookrightarrow G^N$ is not a (forward) Dirac map.

References

[1] A. Alekseev, A. Malkin, and E. Meinrenken. Lie group valued moment maps. *J. Differential Geom.*, 48(3):445–49, 1998.

[2] M. F. Atiyah and R. Bott. The Yang-Mills equations over Riemann surfaces. *Phil. Trans. R. Soc. London*, 308:523–615, 1982.

[3] P. Ševera. Moduli spaces of flat connections and Morita equivalence of quantum tori. *ArXiv e-prints*, June 2011.

[4] E. Meinrenken. Lectures on pure spinors and moment maps. *Poisson Geometry and Mathematical Physics, Contemporary Mathematics*, 450:199–222, 2008.