

Fock-Rosly description of the moduli space of flat connections on a Riemann surface

Yiannis Loizides, Winter 2013

1 Introduction

In a well-known paper [1], Atiyah and Bott showed how the moduli space of flat connections on a Riemann surface could be equipped with a Poisson structure, via an infinite dimensional reduction from the space of all connections. In this essay, we will focus on a later construction of the same Poisson structure due to Fock and Rosly [3]. Their method avoids the infinite dimensional reduction, as well as providing a natural treatment of a class of functions on the moduli space (roughly: Wilson loops), which had been considered by others (c.f. [2] for an exposition and references). The main focus of this essay will be to describe in detail the proof that the construction of Fock and Rosly gives the same Poisson structure on the moduli space as that of Atiyah and Bott.

Let's begin by reviewing the method of Atiyah and Bott. Let S be a compact Riemann surface, possibly with boundary. Let G be a Lie group with Lie algebra \mathfrak{g} . We assume that \mathfrak{g} is equipped with a non-degenerate, symmetric bilinear form $B : \mathfrak{g} \otimes \mathfrak{g} \rightarrow \mathbb{R}$ which is invariant with respect to the adjoint action. We consider the trivial bundle $S \times G$ and fix a trivialization. A connection is then simply a Lie algebra-valued 1-form $A \in \Omega^1(S, \mathfrak{g})$. The infinite dimensional vector space $\mathcal{A}^S = \Omega^1(S, \mathfrak{g})$ carries a symplectic structure

$$\omega(a, b) = \int_S B(a \wedge b).$$

The group $G^S = C^\infty(S, G)$ acts on \mathcal{A}^S by gauge transformations

$$A \mapsto Ad_g A - dg g^{-1}$$

preserving the symplectic structure. Let $\Gamma = \partial S$, and G_Γ^S the subgroup of based gauge transformations, i.e. gauge transformations restricting to the identity on Γ . The smooth dual of the Lie algebra $\mathfrak{g}_\Gamma^S = \{x \in C^\infty(S, \mathfrak{g}) \mid x|_\Gamma = 0\}$ can be identified with $\Omega^2(S, \mathfrak{g})$ via the pairing

$$\langle \alpha, x \rangle = \int_S B(\alpha \wedge x) \quad \alpha \in \Omega^2(S, \mathfrak{g}), \quad x \in \mathfrak{g}_\Gamma^S.$$

Atiyah-Bott showed that the action of G_Γ^S on \mathcal{A}^S is Hamiltonian, where the moment map $F : \mathcal{A}^S \rightarrow (\mathfrak{g}_\Gamma^S)^*$ is the curvature $F(A) = dA + A \wedge A$. The symplectic reduction

$$\mathcal{M}_\Gamma = F^{-1}(0)/G_\Gamma^S$$

is the space of flat connections \mathcal{A}_{fl}^S modulo based gauge transformations.

We would like to perform a further reduction by G^Γ in order to obtain the space of flat connections modulo all gauge transformations (not only those based on Γ). Roughly, the moment map should send a connection A to its pull-back to Γ . But this doesn't quite work, as the resulting map is not Poisson but instead an extra 2-cocycle appears. To get a true moment map it's necessary to pass to a central extension of G^S . Let $\Gamma_1, \dots, \Gamma_k$ denote the components of Γ . Each Γ_i is a circle, and so we can consider the affine Kacs-Moody algebra $\widehat{\mathfrak{g}^{\Gamma_i}}$. This is the 1-dimensional central extension of \mathfrak{g}^{Γ_i} where the 2-cocycle is

$$c(x, y) = \int_{\Gamma_i} B(x \wedge dy).$$

The dual $(\widehat{\mathfrak{g}^{\Gamma_i}})^*$ can be identified with $\Omega^1(\Gamma_i, \mathfrak{g}) \oplus \mathbb{R}$ via the pairing

$$\langle (\alpha, c), (x, d) \rangle = \int_{\Gamma_i} B(\alpha \wedge x) + cd.$$

The map $\Phi : \mathcal{M}_\Gamma \rightarrow \prod_i (\widehat{\mathfrak{g}^{\Gamma_i}})^*$ given by

$$\Phi(A) = \{(A|_{\Gamma_1}, 1), \dots, (A|_{\Gamma_k}, 1)\}$$

is a moment map for the action of $\mathcal{H} = \prod_i \widehat{G^{\Gamma_i}}$ on \mathcal{M}_Γ . Notice that this map is well-defined on equivalence classes, since above we took a quotient by *based* gauge transformations, which do not affect the boundary values.

By standard results on moment maps, this means that the quotient

$$\mathcal{M} = \mathcal{M}_\Gamma / \mathcal{H}$$

is Poisson with symplectic leaves given by the inverse images of coadjoint orbits in \mathcal{H} . Recall that the coadjoint orbits of the affine Kacs-Moody group are classified according to the conjugacy class of the holonomy around the loop. Since \mathcal{H} is a product of affine Kacs-Moody groups, a choice of coadjoint orbit corresponds to a choice of conjugacy class for each component Γ_i . In summary, the space \mathcal{M} of flat connections modulo gauge transformations is a finite dimensional Poisson manifold (possibly with singularities); the symplectic leaves are obtained by fixing the conjugacy class of the holonomy around each of the boundary components.

2 Fock-Rosly description

The basic idea of the Fock-Rosly approach is to work with a graph l embedded in the surface rather than the surface itself. The graph gives a kind of discretization of the surface (in the same spirit as lattice approximations for gauge theories), and so should completely encode its topology. Fock-Rosly begin with a finite dimensional Poisson manifold \mathcal{A}^l (in general it is a product of copies of G) the points of which represent holonomies along the edges of the graph, and then perform a finite reduction to obtain \mathcal{M} . The space \mathcal{A}^l may be thought of

as resulting from the space of flat connections on taking a quotient by the group of gauge transformations based at the vertices of the graph. The reduction is then just by a finite product of copies of G , equal to the number of vertices.

Throughout we will assume that S has at least one boundary component; say S has genus g and b boundary components, with $b \geq 1$. In this case, the fundamental group of the surface is a free group with $2g + b - 1$ generators. For simplicity, suppose one of the boundary components Γ_0 has been marked. We would like to choose a graph l and an embedding of l in S that “encodes the topology” of S . The minimal requirement is that by following various sequences of edges in the (embedded) graph, we should be able to obtain a set of curves which generate the fundamental group of the surface obtained from S by closing one of the holes (say Γ_0) to obtain a surface with genus g and $b - 1$ boundary components. We will in fact require that the graph l be a free group with $2g + b - 1$ generators. The simplest possibility is to have a single vertex and $2g + b - 1$ loops (take the standard generators for the fundamental group of S except for the generator corresponding to the marked boundary component Γ_0 as this is the hole that gets closed). More complicated graphs can arise by (e.g.) adding more vertices. We can even recover the surface from such a graph together with a cyclic ordering of the edges at each vertex which determines the embedding (when the graph is embedded in S , this is induced from the orientation)—Fock and Rosly call this a *fat graph*.

Let $E(l)$ denote the set of ends of edges (meaning each edge corresponds to two different elements of $E(l)$), and let $N(l)$ denote the set of vertices. The map $\alpha \mapsto \alpha^\vee$ sending an end to the other end of the same edge, defines an involution of $E(l)$. Let $[\alpha]$ denote the vertex associated with the end α , and let $[\alpha, \alpha^\vee]$ denote the edge associated to α . Following Fock and Rosly, we define a *graph connection* to be a map $A : E(l) \rightarrow G$ satisfying $A(\alpha^\vee) = A(\alpha)^{-1}$. Let \mathcal{A}^l denote the set of all graph connections. If m is the number of edges of l , then

$$\mathcal{A}^l \simeq \left\{ (g_\alpha) \in \prod_{\alpha \in E(l)} G^{(\alpha)} \mid g_\alpha = g_{\alpha^\vee}^{-1} \right\} \simeq \left\{ (g_1, g_1^{-1}, \dots, g_m, g_m^{-1}) \in G^{2m} \right\}. \quad (1)$$

We also have an analog of the gauge group, $\mathcal{G}^l \simeq G^n$ (n =number of vertices), which we define to be the set of maps $g : N(l) \rightarrow G$. The action is

$$(g \cdot A)(\alpha) = g(\alpha^\vee)^{-1} A(\alpha) g(\alpha).$$

We'll also write this action as $g \cdot (g_\alpha)$ when using the description above. Notice that this is exactly the usual effect of gauge transformations on holonomy.

Now let e^i denote an orthonormal basis on \mathfrak{g} , and let L^i, R^i denote the corresponding left and right vector fields on G . Let L_α^i, R_α^i denote the corresponding vector fields on the α -copy of G (i.e. $G^{(\alpha)}$) in the product appearing in (1). Define vector fields on the product $\prod_\alpha G^{(\alpha)}$ by

$$X_\alpha^i := L_\alpha^i - R_{\alpha^\vee}^i.$$

Notice that at points in \mathcal{A}^l these are tangent to \mathcal{A}^l and so define vector fields on \mathcal{A}^l . We'll see why these vector fields are used later on (for now, notice that if $\iota : G \rightarrow G \times G$ is the

map $\iota(g) = (g, g^{-1})$ then $\iota_*L^i = L_1^i - R_2^i$.

Given a connection A on S and a graph l embedded in S , we obtain a graph connection (also denoted by A) by putting $A(\alpha) = hol(A, \alpha)$ for each $\alpha \in E(l)$ (as above, we are assuming a trivial bundle $S \times G$ with fixed trivialization). Let $G_{N(l)}^S$ denote the set of gauge transformations based at the vertices of l . Since elements of $G_{N(l)}^S$ don't change the holonomy of a connection along the edges of the graph, it's clear that we have a map $hol_l : \mathcal{A}_{fl}^S / G_{N(l)}^S \rightarrow \mathcal{A}^l$. A flat connection is completely determined by its holonomy along representatives for a set of generators of the fundamental group; in fact such a connection gives a representation $\pi_1 \rightarrow G$ of the fundamental group. It follows that the map hol_l is injective. In fact hol_l is also surjective. This is because if we're given a graph connection, we can always find a flat connection on S with the prescribed holonomies (to show this formally, the basic idea is to pass to the space \tilde{S} of (fixed endpoint-) homotopy classes of paths having initial point in the set of vertices $N(l)$ (for example, this would be the universal covering space if there is only one vertex), and then use the trivial connection on $\tilde{S} \times G$ to obtain an appropriate flat connection on $S \times G \simeq (\tilde{S} \times G) / \pi_1(S, N(l))$ where $\pi_1(S, N(l))$, the groupoid of homotopy classes of paths between the vertices $N(l)$, acts on G according to the prescribed holonomies along a sequence of edges homotopic to the path; c.f. [4] for an exposition).

One remark is that it is possible to consider more general graphs (for example, a triangulation of S), but in this case the map hol_l has image equal to the submanifold of *flat* graph connections, i.e. those for which if we form a closed curve from a sequence of edges which is homotopic (in S) to the constant curve, then the product of the group elements assigned to the edges going around the curve is the identity.

This shows that $hol_l : \mathcal{A}_{fl}^l / \mathcal{G}^l \simeq \mathcal{M}$ at least as (possibly singular) manifolds. Fock and Rosly now proceed to define a Poisson structure on \mathcal{A}^l and \mathcal{G}^l , and then prove that hol_l is a Poisson map for this structure. In order to do this, they introduce a linear order $<$ on the ends of edges at each vertex, calling the resulting data a *ciliated fat graph*—they visualize the linear order being specified by adding a small “cilium” at each vertex indicating where we should start counting.

For each vertex $n \in N(l)$, we choose an r-matrix $r(n) \in \mathfrak{g} \otimes \mathfrak{g}$, i.e. a solution of the Yang-Baxter equation, and require that the symmetric part of $r(n)$ equal the quadratic form B . With respect to the orthonormal basis e^i we can write $r_{ij} = B_{ij} + a_{ij}$ where $a_{ij} = -a_{ji}$. We will not go into any detail on r-matrices, but for our purposes it suffices to know that these conditions (specifically, the Yang-Baxter equation) imply that $r(n)$ induces a Poisson bracket on \mathfrak{g} , and hence G becomes a Poisson Lie group. Following Fock and Rosly, we define a bivector field on \mathcal{A}^l by

$$\pi = \sum_{n \in N(l)} \left(\sum_{\alpha < \beta \in n} r_{ij}(n) X_\alpha^i \wedge X_\beta^j + \frac{1}{2} \sum_{\alpha \in n} r_{ij}(n) X_\alpha^i \wedge X_\alpha^j \right). \quad (2)$$

We can simplify this expression using the decomposition $r_{ij} = B_{ij} + a_{ij}$. We have

$$\sum_{\alpha < \beta} r_{ij} X_\alpha^i \wedge X_\beta^j = \sum_{\alpha < \beta} B_{ij} X_\alpha^i \wedge X_\beta^j + \sum_{\alpha \neq \beta} a_{ij} X_\alpha^i \otimes X_\beta^j,$$

and similarly, using that $B_{ij} = \pm \delta_{ij}$ (the e^i are orthonormal) the other term is

$$\begin{aligned} \frac{1}{2} \sum_{\alpha} r_{ij} X_\alpha^i \wedge X_\alpha^j &= \frac{1}{2} \sum_{\alpha} B_{ij} X_\alpha^i \wedge X_\alpha^j + a_{ij} X_\alpha^i \wedge X_\alpha^j \\ &= \sum_{\alpha} 0 + a_{ij} X_\alpha^i \otimes X_\alpha^j. \end{aligned}$$

Adding the two parts we find

$$\begin{aligned} \pi &= \sum_{n \in N(l)} \left(\sum_{\alpha, \beta \in n} a_{ij}(n) X_\alpha^i \otimes X_\beta^j + \sum_{\alpha < \beta \in n} B_{ij} X_\alpha^i \wedge X_\beta^j \right) \\ &= \sum_{n \in N(l)} \left(a_{ij}(n) X_\Delta^i(n) \otimes X_\Delta^j(n) + \sum_{\alpha < \beta \in n} B_{ij} X_\alpha^i \wedge X_\beta^j \right) \end{aligned} \quad (3)$$

where we've introduced the vectors $X_\Delta^i(n) := \sum_{\alpha \in n} X_\alpha^i$, which are tangent to the orbit of \mathcal{G}^l . This means that the part of π involving the anti-symmetric part a_{ij} gives trivial contribution in the quotient $\mathcal{A}^l/\mathcal{G}^l$ (since for a \mathcal{G}^l -invariant function f on \mathcal{A}^l , clearly $X_\Delta^i(n)f = 0$). In particular, this shows that the Poisson structure π induced on the quotient is independent of the choices of r-matrices (the symmetric part is already constrained to be B_{ij}).

For now we will take for granted (1) that this formula defines a Poisson structure on \mathcal{A}^l , (2) that if \mathcal{G}^l is equipped with the direct product Poisson structure (using the r-matrices $r(n)$ on the n th factor), then the action on \mathcal{A}^l is Poisson, (3) consequently the quotient $\mathcal{A}^l/\mathcal{G}^l$ acquires a Poisson structure. Instead we will focus on showing that the bracket induced by π on the quotient agrees with the Atiyah-Bott Poisson structure on \mathcal{M} (thus giving an alternative proof that π induces a Poisson structure on the quotient).

3 Poisson brackets of certain functions

As preparation for the proof, we will compute the Poisson bracket of certain functions on \mathcal{M} in the Atiyah-Bott description. Let $f : G \rightarrow \mathbb{R}$ be a smooth function on G . Consider a simple curve $\alpha : [0, 1] \rightarrow S$ in the surface. Given a connection A , let $h_\alpha(A) = \text{hol}(A, \alpha)$ denote the holonomy of A along α . As above, for simplicity we are taking the bundle to be $S \times G$ with fixed trivialization. So $A \in \Omega^1(S, \mathfrak{g})$ and $h_\alpha(A) = h(1)$ where $h(s)$ is the unique solution to the linear ODE

$$\dot{h} = A(\dot{\alpha}(s))h(s), \quad h(0) = Id.$$

We define a function on the space of connections by $f^\alpha = f \circ h_\alpha$. Fix $u \in \Omega^1(S, \mathfrak{g})$. And let $g = h_\alpha(A)$ (constant). By definition

$$\begin{aligned} df_A^\alpha(u) &= \left. \frac{d}{dt} \right|_0 f_\alpha(A + tu) \\ &= \left\langle (df)_g, \left. \frac{d}{dt} \right|_0 h_\alpha(A + tu) \right\rangle \\ &= \left\langle (df)_g, \int_\alpha g \cdot v_\alpha(u) \right\rangle \\ &= \left\langle (df)_g, \int_S \delta_\alpha \wedge g \cdot v_\alpha(u) \right\rangle \end{aligned}$$

where the second to last line defines v_α as an $End(\mathfrak{g})$ -valued function on α , and δ_α denotes an appropriate delta-type 1-form supported on α (altogether $\delta_\alpha g \cdot v_\alpha$ is the variational derivative of the functional h_α at the point $A \in \mathcal{A}^S$). Since B is non-degenerate, we can define a function $F : G \rightarrow \mathfrak{g}$ by

$$\langle (df)_g, X_g \rangle = B(F(g), g^{-1}X_g).$$

Fix an orthonormal basis e^i for \mathfrak{g} (so $B(e^i, e^j) = \pm\delta^{ij}$). Let L^i, R^i denote the corresponding left and right invariant vector fields on G respectively. Writing $F(g) = F_i(g)e^i$, $B(e^i, e^j) = B^{ij}$ we find

$$F_i = B_{ij}L^j f.$$

Using F the differential can be written

$$\begin{aligned} df_A^\alpha(u) &= \int_S B(F(g)\delta_\alpha \wedge v_\alpha(u)) \\ &= \int_S B(v_\alpha^T F(g)\delta_\alpha \wedge u), \end{aligned}$$

here v^T denotes the adjoint with respect to B . Comparing with the equation for the symplectic form, this shows that the Hamiltonian vector field for the function f^α is

$$Hf^\alpha(A) = v_\alpha^T F(h_\alpha(A))\delta_\alpha.$$

Now let α, β be two curves in S which we assume intersect transversely. Choose functions f_1, f_2 on G . Using what we've just found, their Poisson bracket is

$$\begin{aligned} \{f_1^\alpha, f_2^\beta\} &= \omega(Hf_1^\alpha, Hf_2^\beta) \\ &= \int_S B(v_\alpha^T F_1 \circ h_\alpha, v_\beta^T F_2 \circ h_\beta)\delta_\alpha \wedge \delta_\beta \\ &= \sum_{x \in \alpha \cap \beta} \epsilon(x) B(v_\alpha^T(x)F_1 \circ h_\alpha, v_\beta^T(x)F_2 \circ h_\beta). \end{aligned}$$

In this equation, $\epsilon(x) = \pm 1$ is the oriented intersection number at x , which arises from integrating over the wedge of the delta-type 1-forms.

To proceed further, we need to determine $v_\alpha(x)$ at the points x of intersection. Below, we will only need to know the result in the case where A vanishes on a neighbourhood which includes each of the endpoints of α as well as all of the intersection points—say A vanishes on all of α except possibly $\alpha([s_0, s_1])$ where $0 < s_0 < s_1 < 1$. Now define $h_t(s)$ to be the family of solutions to the linear ODEs

$$\partial_s h_t = (A + tu)(\dot{\alpha}(s))h_t(s), \quad h_t(0) = Id, \quad (4)$$

as above, so that $h_t(1) = h_\alpha(A + tu)$. Integrating both sides from 0 to 1 we obtain

$$h_\alpha(A + tu) = h_t(1) = h_t(0) + \int_0^1 (A + tu)(s)h_t(s)ds.$$

In the above formula (and also below) we are abusing notation and writing $A(\dot{\alpha}(s))$ as $A(s)$, and likewise for u . Now take $\partial_t|_0$ of both sides, which gives

$$\partial_t|_0 h_\alpha(A + tu) = \int_0^1 \left[A(s)\partial_t|_0 h_t(s) + u(s)h_0(s) \right] ds$$

We only need to determine v_α at points in $\alpha([0, s_0])$ and $\alpha([s_1, 1])$ (since we are assuming all the intersection points fall in these intervals), so it suffices to consider two types of variations u . First take $u(s) = 0$ for $s \in [0, s_1]$. This implies that $h_t(s)$ is independent of t for $s \in [0, s_1]$, and consequently $\partial_t|_0 h_t(s) = 0$ for $s \in [0, s_1]$. Also $A(s) = 0$ for $s \in [s_1, 1]$ by our assumption on A . Together this implies that the first term in the integral vanishes. Moreover, for the second term, u vanishes on $[0, s_1]$, and $h_0(s) = g = h_\alpha(A)$ is constant for $s \in [s_1, 1]$ (this is because A is identically zero on that interval, so the holonomy doesn't change after reaching s_1). And so the integral simplifies to

$$\int_0^1 g \cdot v_\alpha(u)(s)ds = \partial_t|_0 h_\alpha(A + tu) = \int_0^1 u(s) \cdot g ds.$$

Since u was allowed to vary freely away from $[0, s_1]$, this suffices to show that $v_\alpha(x) = Ad_{g^{-1}}$ for $x \in \alpha([s_1, 1])$. This has a simple interpretation if we imagine the holonomy as a parallel transport operator acting on the left. We multiply by g first (that is the holonomy along $\alpha([0, s_1])$) and then by the integral of u , which is the holonomy along $\alpha([s_1, 1])$ to first order. We get something similar if we take $u(s) = 0$ for $s \in [s_0, 1]$, except the contribution of u to the holonomy comes *before* the segment where A is nonzero (instead of after), so the two are multiplied in the reverse order and we get

$$\int_0^1 g \cdot v_\alpha(u)(s)ds = \partial_t|_0 h_\alpha(A + tu) = \int_0^1 g \cdot u(s)ds,$$

which shows that $v_\alpha(x) = Id$ for $x \in \alpha([0, s_0])$.

Borrowing from the previous section, we can think of α, α^\vee as labelling the two *ends* of the curve (where α labels the end where the parametrisation starts, etc.). If an intersection point x falls in the initial third $\alpha([0, s_0])$ we will say that it is associated with α , while if it

falls in the final third $\alpha([s_1, 1])$ we will say that it is associated with α^\vee . Write αx for the end (either α or α^\vee) that a given intersection point x is associated with. We can summarize the previous paragraph by saying that if $\alpha x = \alpha^\vee$ then $v_\alpha(x) = Ad_{g^{-1}}$, otherwise $\alpha x = \alpha$ and $v_\alpha(x) = Id$.

Applying this we can simplify the formula above for the Poisson bracket, at least for the special case of interest. Let $g_\alpha = hol_\alpha(A)$ and similarly for β . Suppose that A vanishes in neighbourhoods around the two endpoints of both curves (as above) and suppose further that all the intersection points $x \in \alpha \cap \beta$ lie in these portions of the two curves (each x has to be in both). Notice that if $v_\alpha(x) = Id$ then $v_\alpha^T(x) = Id$, while if $v_\alpha(x) = Ad_{g_\alpha^{-1}}$ then $v_\alpha^T(x) = Ad_{g_\alpha}$ by Ad-invariance of B . Using Ad-invariance of B again, we have that $\langle (df)_g, R_g^i \rangle = B(Ad_g F(g), e^i)$, and therefore $Ad_g F(g) = B_{ij}(R_g^i f) e^j$. Let's use $Y_{\alpha,x}^i$ to denote either L^i or R^i depending on whether $\alpha x = \alpha$ or $\alpha x = \alpha^\vee$ and similarly for β . Applying these considerations, the formula becomes

$$\{f_1^\alpha, f_2^\beta\}(A) = \sum_{x \in \alpha \cap \beta} B_{ij}(Y_{\alpha,x}^i f_1)(Y_{\beta,x}^j f_2) \epsilon(x) \quad (5)$$

where $Y_{\alpha,x}^i f_1$ is evaluated at $g_\alpha = h_\alpha(A) \in G$, and $Y_{\beta,x}^j f_2$ is evaluated at $g_\beta = h_\beta(A) \in G$.

Let's now write this equation in a slightly different form (which will begin to indicate the similarity with the Fock-Rosly Poisson structure—to be upgraded to a proof in the next section). Let $N = \{(g_1, g_2) \in G \times G | g_1 = g_2^{-1}\}$ and let $\iota : G \rightarrow N$ denote the isomorphism $\iota(g) = (g, g^{-1})$. Given a function f on G , we can use ι to define a function ϕ on N by pulling back (by ι^{-1}), i.e. $\phi(g, g^{-1}) := f(g)$. Then $L^i f(g) = (L_1^i - R_2^i) \phi(g, g^{-1})$ (because $\iota_* L^i = L_1^i - R_2^i$), and similarly $R^i f(g) = -(R_1^i - L_2^i) \phi(g, g^{-1})$. Let α be a curve as above, and let α^\vee denote α traversed in the opposite direction (to connect with the previous section, think of α and α^\vee as the two *ends* of the curve, with the α direction being that which starts at the end labelled by α , etc.). Then N is the space of graph connections on the curve α , where $g_1 = A(\alpha)$, $g_2 = g_1^{-1} = A(\alpha^\vee)$. The vector field $L_1^i - R_2^i$ is X_α^i while $R_1^i - L_2^i = X_{\alpha^\vee}^i$, therefore

$$\begin{aligned} L^i f(g) &= X_\alpha^i \phi(g, g^{-1}) \\ R^i f(g) &= -X_{\alpha^\vee}^i \phi(g, g^{-1}). \end{aligned}$$

Notice the crucial minus sign! It suggests doing the following. We want to re-write (5) in terms of functions on N and the vector fields X_α^i . The terms involving $(R^i f) \epsilon(x)$ (i.e. for x associated with α^\vee) will become $(-X_{\alpha^\vee}^i \phi) \epsilon(x)$, where here $\epsilon(x)$ is the intersection number relative to α . Instead we can absorb the extra minus sign by taking $\epsilon(x)$ to be the intersection number relative to α^\vee (which has the opposite orientation, and so the intersection number switches sign). The result is that the orientation number at an intersection x is taken relative to the two ends that x is associated with.

Let's now do this. Let a and b be unoriented simple curves in S , with ends labelled α, α^\vee and β, β^\vee respectively. Let ϕ_1, ϕ_2 be functions on N which we lift to functions on \mathcal{A}^S by

$\phi_1^a = \phi_1 \circ \text{hol}_a$, $\phi_2^b = \phi_2 \circ \text{hol}_b$ where $\text{hol}_a(A) = (\text{hol}_\alpha(A), \text{hol}_{\alpha^\vee}(A)) = (g_\alpha, g_\alpha^{-1})$, etc. Let A be a connection on S where as above we assume that A vanishes on the initial and final thirds of both curves, and that all the intersection points x fall in these portions. Then equation (5) becomes

$$\{\phi_1^a, \phi_2^b\}(A) = \sum_{x \in a \cap b} B_{ij}(X_{\alpha x}^i \phi_1)(X_{\beta x}^j \phi_2) \epsilon(x), \quad (6)$$

where in this equation $\epsilon(x) = \pm 1$ denotes the intersection number for the point x relative to the αx , βx orientations on a and b .

4 Proof that the Fock-Rosly and Atiyah-Bott Poisson structures agree

Let ϕ_1, ϕ_2 be \mathcal{G}^l -invariant functions on \mathcal{A}^l . We want to compute their bracket using the Fock-Rosly (FR) Poisson structure, and then compare the result with the bracket of their lifts to \mathcal{M} computed using the Atiyah-Bott (AB) Poisson structure. We will choose a fairly specific (embedded) ciliated fat graph l , which in a certain sense is the simplest possible. It will have a single vertex which we assume is on one of the boundary components (we can use a homotopy to move the vertex to the boundary if needed), and with the cilium pointing outside the surface. And it will have one loop for each of the standard generators of the fundamental group, except the generator for the component of the boundary that the vertex lies on. If S has genus g and b boundary components, then there will be $m = 2g + b - 1$ loops. Although we won't go into any detail, this is not such a big restriction. In their paper, Fock-Rosly describe some natural geometric operations on ciliated fat graphs—such as contracting edges between distinct vertices, erasing edges, adding loops—which turn out to be (FR) Poisson maps, and also preserve lifts of functions to \mathcal{M} . If we're given some (connected) graph containing multiple vertices we can contract edges repeatedly until we have a single vertex, coming to a graph that is essentially equivalent to l . Because these maps are Poisson, it is enough to compare the Poisson structures for the graph l .

The computation using the Fock-Rosly structure is immediate. Since there is only one vertex, there is no need to write a sum over vertices. Here and later we will abbreviate expressions of the form $B_{ij} Y_a^i \wedge Y_b^j$ as $Y_a \wedge Y_b$. Using this short form, the Fock-Rosly Poisson bivector is

$$\pi = \sum_{\alpha > \beta} X_\alpha \wedge X_\beta + (X_\Delta \text{ terms})$$

We haven't written the terms involving X_Δ , since they vanish on \mathcal{G}^l -invariant functions:

$$\begin{aligned} \{\phi_1, \phi_2\}_{FR} &= \langle \pi, d\phi_1 \otimes d\phi_2 \rangle \\ &= \sum_{\alpha > \beta} \langle X_\alpha \wedge X_\beta, d\phi_1 \otimes d\phi_2 \rangle. \end{aligned} \quad (7)$$

On the other hand, ϕ_1 and ϕ_2 lift to functions $\phi_1^l = \phi_1 \circ \text{hol}_l$, $\phi_2^l = \phi_2 \circ \text{hol}_l$ on \mathcal{A}^S , where here hol_l maps a connection A to its set of holonomies along the edges of the graph l . Because

of the discussion above, since the graph l captures the topology of the surface S , all gauge-invariant functions on \mathcal{A}_{fl}^S (and hence all functions on the quotient \mathcal{M}) can be obtained in this way. So we will want to compute the Poisson bracket of the lifted functions ϕ_1^l, ϕ_2^l on \mathcal{A}^S at points where the connection A is flat. For this we will use the result of the computation done in a previous section.

We cannot directly apply the formula obtained in the section above since it only applied in the case that two edges intersected transversely; for example it is not directly applicable to computing the bracket of ϕ_1^l and ϕ_2^l . But we can deform the graph l to obtain a slightly different embedding \bar{l} such that the two embeddings have no common vertices and all intersections of edges of l with those of \bar{l} are transverse. Then l and \bar{l} give two different ways of lifting functions to \mathcal{A}^S . But in fact if l, \bar{l} are homotopic, then the two lifts will agree on \mathcal{A}_{fl}^S up to a gauge transformation, and in particular they will agree for gauge invariant functions. So we will compute the Atiyah-Bott bracket for ϕ_1^l and $\phi_2^{\bar{l}}$ using the formula from the section above, and compare the result with the bracket computed for ϕ_1 and ϕ_2 using the Fock-Rosly bivector.

Following Fock and Rosly, we will choose \bar{l} in a specific way so that we know how it relates to l . Imagine two copies of l lying on top of each other. We fix a point in the middle of each edge (a ‘‘pivot point’’) and move the vertex in one of the copies of l (together with its edges but keeping the pivot points in the middle fixed) slightly counter-clockwise as viewed from the outside of the surface (the boundary component is a circle). We make the number of intersections minimal, and ensure that they are transverse. The deformed copy is \bar{l} . See the attached figures for some examples of this.

Let a_1, \dots, a_m denote the loops of l and $\bar{a}_1, \dots, \bar{a}_m$ the corresponding loops of \bar{l} . Let α_i, α_i^\vee denote the ends of a_i for $i=1, \dots, m$, and likewise $\bar{\alpha}_i, \bar{\alpha}_i^\vee$ denote the corresponding ends in \bar{l} . There is one intersection point at the pivot point $a_i \cap \bar{a}_i$ for each loop, and there is also one intersection point for each pair of ends β, γ in l such that $\beta > \gamma$ (the intersection point being $\beta \cap \bar{\gamma}$). (Roughly speaking, this second set of intersection points arise because the $\bar{\gamma}$ have to cross back over those edges $\beta > \gamma$ which $\bar{\gamma}$ has found itself on the wrong side of due to the fact that we’ve dragged the vertex of \bar{l} counterclockwise.) Probably the best way to see this is by drawing several examples.

We can choose a contractible neighbourhood $U \subset S$ which includes the beginning and ends of all the loops as well as all the intersection points. Also for any intersection point $x \in a \cap \bar{b}$, the set U must contain a segment of a leading to the vertex of l and a segment of \bar{b} leading to the vertex of \bar{l} . Again probably the best way to see that this can be done is by drawing several examples (we give some in the attached figures). Notice that U divides each loop a into three pieces: the part of a outside U , the initial segment of a (relative to some orientation), and the final segment of a . Suppose A is a flat connection at which we would like to evaluate the bracket of $\phi_1^l, \phi_2^{\bar{l}}$. Since U is contractible, we can use a gauge transformation to find another connection A' in the same gauge orbit which vanishes on U . As $\phi_1^l, \phi_2^{\bar{l}}$ are gauge invariant, we can use any connection lying in the same gauge orbit when evaluating

the bracket. So without loss of generality we can assume that we are evaluating the bracket at a flat connection A which vanishes on U . We are now in the situation of the previous section. For each pair of loops a, \bar{b} the connection A vanishes on the initial third and final third of each curve (this was the reason U was chosen the way it was), and so each of the intersection points $x \in a \cap \bar{b}$ becomes associated with one of the ends of each of a and \bar{b} as was described in the previous section. In other words, x is associated to the two ends of edges that can be reached by following loops and staying inside U . For example if, relative to orientations $\alpha, \bar{\beta}$ on a, b respectively, x falls in the initial third of a and the final third of b , then x is associated to α and to $\bar{\beta}^\vee$, and we write this as $lx = \alpha$ and $\bar{l}x = \bar{\beta}^\vee$.

Using (6), the contribution of each point of intersection x to the bracket is

$$B_{ij}(X_{lx}^i \phi_1)(X_{\bar{l}x}^j \phi_2) \epsilon(x)$$

where $\epsilon(x) = \pm 1$ is the intersection number relative to the orientations lx and $\bar{l}x$ on the two loops involved in the intersection. Notice that this expression contains terms like $X_{\bar{\beta}} \phi_2$, etc., which actually just means the same thing as $X_{\beta} \phi_2$, because the functions ϕ_1, ϕ_2 are defined on the same space $\mathcal{A}^l \subset G^{2m}$ (the bar variables only differentiate the two copies of l).

Taking the sum over all the intersection points of l and \bar{l} gives the bracket. We should perhaps make one comment about this. From equation (6), we know that for each pair of intersecting loops a, \bar{b} we should sum over the intersection points of those two loops. But why is it that summing over all the different pairs of loops gives the bracket of the two functions? The reason is a general property of Poisson brackets. Recall that for functions f, g on a Poisson space M and bivector P the Poisson bracket is

$$\{f, g\} = \langle P, df \otimes dg \rangle.$$

Suppose $x_1, \dots, x_n, y_1, \dots, y_k$ are coordinates on M . Then we can write $d = \partial_x + \partial_y$ and expand

$$\{f, g\} = \langle P, \partial_x f \otimes \partial_x g + \partial_x f \otimes \partial_y g + \partial_y f \otimes \partial_x g + \partial_y f \otimes \partial_y g \rangle.$$

Each of the four terms can be given a meaning. Take for example $\langle P, \partial_x f \otimes \partial_y g \rangle$; at a point (c, d) this can be interpreted as the Poisson bracket of the functions $\tilde{f}(x, y) := f(x, d)$ and $\tilde{g}(x, y) := g(c, y)$. A similar thing is happening in our example, except now the “independent variables” correspond to the holonomies of the connection A along the loops a_1, \dots, a_m . This is m “ G -valued variables”, and so the analogous sum has m^2 terms, one for each pair of loops. This explains why summing over all the intersections of the different pairs of loops will give the bracket. It seems to me that for this to work, it must be possible to vary the holonomies along all of the loops *independently* (it is important that the $x_1, \dots, x_n, y_1, \dots, y_m$ form a set of independent variables, otherwise we can’t write $d = \partial_x + \partial_y$). This is possible as long as S has at least one boundary component so that the fundamental group is free. If true, this is another reason for restricting to the case where S has at least one boundary component (both Fock-Rosly and Audin do this early on, though don’t seem to comment explicitly on why it is needed).

Returning to the calculation, the Atiyah-Bott bracket is

$$\{\phi_1^l, \phi_2^{\bar{l}}\}_{AB} = \sum_{x \in l \cap \bar{l}} B_{ij}(X_{lx}^i \phi_1)(X_{\bar{l}x}^j \phi_2) \epsilon(x), \quad (8)$$

where the intersection number $\epsilon(x)$ is taken relative to $(lx, \bar{l}x)$. This is in a form that obscures the (anti-) symmetry, which is hidden in the intersection numbers and the gauge invariance. We can handle this with the following trick. By gauge invariance, we must get the same result if we use \bar{l} to lift ϕ_1 and l to lift ϕ_2 (i.e. both ways must give the bracket of ϕ_1^l and $\phi_2^{\bar{l}}$). In other words

$$\{\phi_1^l, \phi_2^{\bar{l}}\}_{AB} = \{\phi_1^{\bar{l}}, \phi_2^l\}_{AB} = \frac{1}{2} \left(\{\phi_1^l, \phi_2^{\bar{l}}\}_{AB} + \{\phi_1^{\bar{l}}, \phi_2^l\}_{AB} \right). \quad (9)$$

Using the same setup (contractible set U , a connection A vanishing on U , etc.) we have

$$\{\phi_1^{\bar{l}}, \phi_2^l\}_{AB} = \sum_{x \in l \cap \bar{l}} B_{ij}(X_{lx}^i \phi_1)(X_{\bar{l}x}^j \phi_2) \epsilon'(x).$$

Now the intersection number $\epsilon'(x)$ is taken relative to $(\bar{l}x, lx)$, which is just the reverse of $(lx, \bar{l}x)$ (used for $\epsilon(x)$). Thus $\epsilon'(x) = -\epsilon(x)$, and so taking the average as in (9) we get

$$\{\phi_1^l, \phi_2^{\bar{l}}\}_{AB} = \frac{1}{2} \sum_{x \in l \cap \bar{l}} B_{ij}(X_{lx}^i \phi_1 X_{\bar{l}x}^j \phi_2 - X_{\bar{l}x}^i \phi_1 X_{lx}^j \phi_2) \epsilon(x) \quad (10)$$

$$= \frac{1}{2} \sum_{x \in l \cap \bar{l}} \epsilon(x) \langle X_{lx} \wedge X_{\bar{l}x}, d\phi_1 \otimes d\phi_2 \rangle. \quad (11)$$

Recall that we have one point of intersection x coming from each pivot point $a_i \cap \bar{a}_i$. In this case x must be associated with the same end in both l and \bar{l} (otherwise U would have to contain pieces of a_i and \bar{a}_i forming a closed loop, meaning U would not be contractible). Consequently $X_{\bar{l}x} = X_{lx}$ for these points, and the wedge means that these terms drop out. Recall that there is also one intersection point for each pair of ends α, β in l such that $\alpha > \beta$ (the intersection point being $\alpha \cap \bar{\beta}$). Therefore the sum becomes a sum over $\alpha > \beta$. The intersection number can be carefully checked in the general case for a pair of loops; a little thought shows it is $+1$ exactly when $lx > \bar{l}x$ and -1 when $\bar{l}x > lx$. Finally we get the Atiyah-Bott bracket:

$$\{\phi_1^l, \phi_2^{\bar{l}}\}_{AB} = \frac{1}{2} \sum_{\alpha > \beta} \langle X_\alpha \wedge X_\beta, d\phi_1 \otimes d\phi_2 \rangle. \quad (12)$$

This is the same expression as for the Fock-Rosly bracket, except (?) for the factor of $1/2$.

In the attached figures are relevant pictures for the one-holed torus and 3-holed sphere, showing the graph l , the deformed graph \bar{l} , and the contractible open set U . For some nice applications of Fock and Rosly's result to Goldman functions and integrable systems, see the paper by Audin [2].

References

- [1] M. F. Atiyah and R. Bott, *The Yang-Mills equations over Riemann surfaces*, Phil. Trans. R. Soc. London **308** (1982), 523–615.
- [2] M. Audin, *Lectures on gauge theory and integrable systems*, Gauge Theory and Symplectic Geometry (J. Hurtubise and F. Lalonde, eds.), Kluwer Academic Publishers, Netherlands, 1997, pp. 1–48.
- [3] V. Fock and A. Rosly, *Poisson structure on moduli of flat connections on Riemann surfaces and r -matrices*, preprint (1992).
- [4] L. Jeffrey, *Flat connections on oriented 2-manifolds*, Bull. London Math. Soc. **37** (2005), 1–14.