5 Geometry of field space

Differential geometry is one of the most powerful mathematical tools one can use in theoretical physics. It is indispensible in the description of fields in spacetime. Indeed, we view spacetime as a differential manifold $\mathcal{M}$, and fields as tensors, spinors, or other objects, living on that manifold. The simplest kind of field, a scalar field, is just a function on $\mathcal{M}$. The electromagnetic gauge potential A_μ is a 1-form on $\mathcal{M}$. The metric $g_{\mu\nu}$ is a rank 2 tensor on $\mathcal{M}$.

But the use of differential geometry does not stop there. For example, previously we have seen how the phase space of a physical system can be viewed as a symplectic manifold. In ordinary mechanics, there are only finitely many degrees of freedom, so this manifold is finite-dimensional. However, in the case of field theory, there are infinitely many degrees of freedom. Thus, phase space is infinite-dimensional.

Usually when learning differential geometry one only deals with finite-dimensional manifolds. The standard definitions do not automatically extend to the infinite-dimensional case, and subtle and strange things can happen if one is not careful. There are different ways to define infinite-dimensional manifolds, with different names, for example: Banach manifolds, Frechet manifolds, Hilbert manifolds. If we were being perfectly rigorous in this course, we would have to pick one of these options, and check that everything we are doing still makes sense. However, there is not enough time to do this, and we wouldn't get much out of it anyway. So we will not worry too much about this, and just assume that all of the finite-dimensional intution works out.

5.1 Configuration space

A 'configuration' of a field is a specification of the value that it takes at every point in spacetime. For a scalar field, a configuration is just a function over $\mathcal{M}$. More generally, a field configuration is a section of some bundle Φ over $\mathcal{M}$, which is sometimes called the 'field bundle'. Recall that such bundles are equipped with a map $\pi : \Phi \to \mathcal{M}$, and that the preimage

$$\pi^{-1}(x) = \{y \in \Phi : \pi(y) = x\} \tag{5.1}$$

of a point $x \in \mathcal{M}$ is known as the 'fibre' over x. A 'section' of Φ is the choice of an element of each fibre. The space of all sections is denoted $\mathcal{C} = \Gamma(\Phi)$, and is called 'configuration space'. Let us use $\phi(x)$ to denote the value taken at $x \in \mathcal{M}$ by a given configuration $\phi \in \mathcal{C}$.

For an example, consider the electromagnetic gauge potential, where Φ is $T^*\mathcal{M}$, the cotangent bundle over $\mathcal{M}$. An electromagnetic field configuration is a section of the cotangent bundle $T^*\mathcal{M}$, i.e. a choice of 1-form at x for every $x \in \mathcal{M}$.

We can view configuration space as an infinite-dimensional manifold (modulo rigour), from which every other space of interest in the covariant phase space formalism is derived. Let us now try to understand the meanings of various geometric objects in configuration space.

A path in configuration space $I \to \mathcal{C}$, $t \mapsto \phi_t$, where I is some interval in $\mathbb{R}$, is simply a 1-parameter family of field configurations. Let V be the tangent vector to this path. What is the meaning of this tangent vector? To find out, suppose we have some function F on configuration space, which we evaluate on ϕ_t. Then taking a derivative with respect to t, we have

$$\frac{\mathrm{d}}{\mathrm{d}t} F(\phi_t) = \dot{\phi}_t \cdot \left.\frac{\partial F}{\partial \phi}\right|_{\phi=\phi_t}, \tag{5.2}$$

where the $\cdot$ denotes a sum over field components. But by the definition of a tangent vector, we have

$$\frac{\mathrm{d}}{\mathrm{d}t}F(\phi_t) = V(F(\phi_t)), \tag{5.3}$$

so

$$V = \dot{\phi}_t \cdot \frac{\partial}{\partial \phi}. \tag{5.4}$$

So the components of the tangent vector are just $\dot{\phi}_t$, which is the infinitesimal change in field configuration as one moves along the path ϕ_t. More generally, considering the tangent vectors to all possible paths in configuration space, a general vector on $\mathcal{C}$ may be written $\delta\phi \cdot \frac{\partial}{\partial \phi}$, where $\delta\phi$ is some infinitesimal variation of the fields. So this tells us what configuration space vectors are – they are simply field variations! The space of all field variations $\phi \to \phi + \delta\phi$ is thus the same as $\mathrm{T}\mathcal{C}$, the tangent bundle to the configuration space. The space of all field variations to a particular field configuration ϕ is $\mathrm{T}_\phi\mathcal{C}$, the tangent space to $\mathcal{C}$ at ϕ.

5.2 Locality and jets

In this course we will be interested in theories which can be defined *locally*. For such theories, the Lagrangian form at a point x in spacetime only depends on the fields ϕ and some number of their derivatives at x:

$$L|_{x\in\mathcal{M}} = L\bigg(\phi(x), \frac{\partial \phi}{\partial x^\mu}(x), \frac{\partial^2 \phi}{\partial x^\mu \partial x^\nu}(x), \dots\bigg) \tag{5.5}$$

On the other hand, a general field-dependent spacetime function is some map

$$f: \quad \mathcal{C} \times \mathcal{M} \to \mathbb{R}, \quad (\phi, x) \mapsto f(x)|_\phi, \tag{5.6}$$

that does not need to have this kind of local dependence. For example, f could be the shortest distance from x to $\partial\mathcal{M}$. This function depends on the metric in a very non-local way.

A useful way to formalise this locality is in terms of *jets*. Given a section ϕ of the field bundle Φ, i.e. a configuration in $\mathcal{C}$, we can compute k of its derivatives at a point $x \in \mathcal{M}$. The k-jet of ϕ at x is simply the collection of values taken by the field and these derivatives,

$$j^k_x(\phi) = \bigg(\phi(x), \frac{\partial \phi}{\partial x^\mu}(x), \frac{\partial^2 \phi}{\partial x^\mu \partial x^\nu}(x), \dots, \frac{\partial^k \phi}{\partial x^\mu \dots \partial x^\rho}(x)\bigg). \tag{5.7}$$

The k-jet space of Φ at x is the set containing all possible k-jets at x,

$$J^k_x = \{j^k_x(\phi) : \phi \in \mathcal{C}\}. \tag{5.8}$$

In other words, each point in J^k_x corresponds to a specification of the first k derivatives of the fields at x. Finally, the k-jet bundle is the disjoint union over spacetime of all the k-jet spaces,

$$J^k = \bigsqcup_{x\in\mathcal{M}} J^k_x. \tag{5.9}$$

We have the map

$$J^k \to \mathcal{M}, \quad J^k_x \ni j \mapsto x, \tag{5.10}$$

under which J^k forms a fibre bundle over spacetime. The fibre over $x \in \mathcal{M}$ is the k-jet space J^k_x. Given a field configuration $\phi \in \mathcal{C}$, let $j^k(\phi)$ be the section of the k-jet bundle J whose value in the fibre over x is $j^k_x(\phi)$.

From now on, let us drop the k on everything, and just refer to k-jets as jets, the k-jet bundle J^k as the jet bundle J, and so on.

The map $j_x : \mathcal{C} \to J_x$ enables us to define local functions on configuration space. Indeed, suppose we have some function $f : J \to \mathbb{R}$ defined on the jet bundle,

$$f = f\Big(x, \underbrace{\phi(x), \frac{\partial \phi}{\partial x^\mu}(x), \frac{\partial^2 \phi}{\partial x^\mu \partial x^\nu}(x), \dots}_{j_x(\phi)}\Big). \tag{5.11}$$

We can then define a field-dependent spacetime function $F : \mathcal{C} \times \mathcal{M} \to \mathbb{R}$ via

$$F : \quad (\phi, x) \mapsto F(x)|_\phi = f(x, j_x(\phi)). \tag{5.12}$$

In this way, at the point x, the function F only depends on the values of ϕ and its first k derivatives at x. Thus F is a local function.

We can do a similar thing to define other kinds of local objects, such as spacetime tensors that depend locally on the fields. Essentially, we can think of these objects as being defined on the jet bundle. Then we use the function $j : \mathcal{C} \to \Gamma(J)$ to go from the field configuration ϕ to a section $j(\phi)$ of the jet bundle, on which we can then evaluate the object. This will always give a local dependence on the fields and their derivatives.

It is sometimes more useful to think of local objects as being defined directly on the jet bundle, and we will frequently do this.

5.3 Variational bicomplex

Of particular interest in the covariant phase space formalism are differential forms with local field dependence. We already know how to think of local 0-forms, i.e. local functions – these are just functions on the jet bundle J. What about higher-degree forms?

Using a set of local spacetime coordinates, we can immediately write down a set of D 1-forms on the jet bundle:

$$\mathrm{d}x^\mu\,, \quad \mu = 1, \dots, D. \tag{5.13}$$

However, x^μ are not the only coordinates on the jet bundle. We also have coordinates for the field values and their derivatives $\phi, \partial_\mu \phi, \partial_\mu \partial_\nu \phi, \dots$, so we also have the 1-forms

$$\mathrm{d}\phi\,, \mathrm{d}(\partial_\mu \phi)\,, \mathrm{d}(\partial_\mu \partial_\nu \phi)\,, \dots. \tag{5.14}$$

Actually, when discussing these 1-forms, we will change the notation d to δ, and so write them as

$$\delta\phi, \delta(\partial_\mu \phi), \delta(\partial_\mu \partial_\nu \phi), \dots. \tag{5.15}$$

The reason for this change of notation is that we will eventually associate these 1-forms with linearised field variations. The notation then helps us to distinguish between field variations and spacetime forms.

Note that the first set of 1-forms (5.13) are aligned along spacetime directions of the bundle, while the second set (5.15) are aligned along the vertical directions of the bundle, i.e. along the fibres. For this reason, we sometimes call (5.13) 'horizontal', and (5.15) 'vertical'.

We can get higher order forms by taking exterior products of the above basis 1-forms. For example

$$\mathrm{d}x^1 \wedge \mathrm{d}x^2 = -\,\mathrm{d}x^2 \wedge \mathrm{d}x^1 \tag{5.16}$$

is a horizontal 2-form on J. It is a useful convention when taking exterior products of vertical forms to not explicitly write out the wedge symbol. This allows us to avoid confusing the two types of form. For example,

$$\delta\phi\,\delta(\partial_1\phi) = -\delta(\partial_1\phi)\,\delta\phi \tag{5.17}$$

is a vertical 2-form on J. We can also have forms which have both horizontal and vertical parts, such as

$$\mathrm{d}x^\mu\,\delta(\partial_\mu\phi) = -\delta(\partial_\mu\phi)\,\mathrm{d}x^\mu\,. \tag{5.18}$$

It is useful to distinguish between the horizontal and vertical degrees of a form. To this end, a form which involves p horizontal 1-forms and q vertical 1-forms is called a (p,q)-form. For example, (5.18) is a $(1,1)$-form.

Now let's consider exterior derivatives on the jet bundle. We denote the exterior derivative $\underline{\mathrm{d}}$. Given a function f on the jet bundle, we can now take its exterior derivative to get the 1-form

$$\underline{\mathrm{d}}f = \mathrm{d}x^\mu\,\frac{\partial f}{\partial x^\mu} + \delta\phi\cdot\frac{\partial f}{\partial\phi} + \delta(\partial_\mu\phi)\cdot\frac{\partial f}{\partial\partial_\mu\phi} + \ldots, \tag{5.19}$$

where, as previously, the $\cdot$ denotes a sum over fields. It is convenient to write this as

$$\underline{\mathrm{d}}f = \mathrm{d}f + \delta f, \tag{5.20}$$

where

$$\mathrm{d}f = \mathrm{d}x^\mu\,\frac{\partial f}{\partial x^\mu} \tag{5.21}$$

and

$$\delta f = \delta\phi\cdot\frac{\partial f}{\partial\phi} + \delta(\partial_\mu\phi)\cdot\frac{\partial f}{\partial\partial_\mu\phi} + \ldots. \tag{5.22}$$

The exterior derivatives of higher degree forms are defined in the usual way.

We call d the horizontal exterior derivative, and δ the vertical exterior derivative. One can confirm that they anticommute[2]

$$\{\mathrm{d},\delta\} = 0. \tag{5.23}$$

We also have $\mathrm{d}^2 = 0$ and $\delta^2 = 0$. Suppose α is a (p,q)-form. Then $\mathrm{d}\alpha$ is a $(p+1,q)$-form, while $\delta\alpha$ is a $(p,q+1)$-form.

We call this system of (p,q)-forms, as well as the two types of exterior derivative, a *variational bicomplex*.

[2] It is also possible to use a convention in which d and δ commute.

5.4 Solution space

As described in the previous section, when dealing with covariant field theories it is most convenient to think of the Lagrangian L as a top form on spacetime that depends locally on the fields ϕ. In the context of the variational bicomplex, we should think of L as a $(D,0)$-form, because it is a horizontal (i.e. spacetime) form of degree D, but only a function of the fields, i.e. a vertical form of degree 0.

When we take a variation of the Lagrangian with respect to the field change $\phi \to \phi + \delta\phi$, this is essentially the same as taking the vertical exterior derivative. We end up with a $(D,1)$-form, which, as we showed in the last section, may be written

$$\delta L = E \cdot \delta\phi + \mathrm{d}\theta \tag{5.24}$$

by appropriate application of the product rule. In this expression E is a spacetime D-form that depends locally on the fields, so it is a $(D,0)$-form, and $\theta = \theta(\phi, \delta\phi)$ is a spacetime $(D-1)$-form that depends locally on the fields ϕ, and linearly on one field variation $\delta\phi$, so it is a $(D-1,1)$-form. One can check that the degrees of all the terms in the above equation are the same.

The equations of motion $E(\phi) = 0$ are obeyed only on a subspace $\mathcal{S} \subset \mathcal{C}$ of the full configuration space. This space $\mathcal{S}$ is called the *solution space*, and configurations in the solution space are sometimes called *on-shell*. On-shell, we clearly have

$$\delta L = \mathrm{d}\theta\,. \tag{5.25}$$

In the case of a theory without gauge symmetry, the solution space $\mathcal{S}$ is the phase space of the theory. In the case of a theory *with* gauge symmetry, we need to perform a symplectic reduction on the solution space, and the result is the phase space of the theory.

In either case we need to define a presymplectic form for the solution space, and this is what we will do next.

5.5 Presymplectic current and presymplectic form

The $(D-1,1)$-form θ is sometimes called the *presymplectic potential density*. Consider its vertical exterior derivative

$$\omega = \delta\theta. \tag{5.26}$$

This is a $(D-1,2)$-form, sometimes called the *presymplectic current*. On the space of solutions, we have

$$\mathrm{d}\omega = \mathrm{d}(\delta\theta) = -\delta(\mathrm{d}\theta) = -\delta^2 L = 0. \tag{5.27}$$

Thus ω is on-shell closed.

Let us now pick a Cauchy surface Σ in spacetime. Since ω is a $(D-1)$-form in spacetime, we can integrate it over Σ to obtain

$$\Omega_\Sigma = \int_\Sigma \omega. \tag{5.28}$$

Ω_Σ is a 2-form on configuration space $\mathcal{C}$. Moreover, it is an exact form $\Omega_\Sigma = \delta\Theta_\Sigma$, where

$$\Theta_\Sigma = \int_\Sigma \theta. \tag{5.29}$$

The pullback of Ω_Σ to the space of solutions $\mathcal{S} \subset \mathcal{C}$ is the presymplectic 2-form.

There is now the question of which Cauchy surface we should use to define the presymplectic 2-form. Suppose we have two Cauchy surfaces $\Sigma_{1,2}$ which together bound a region $\mathcal{U}$ of spacetime. Note that any two Cauchy surfaces which share a boundary $\partial\Sigma_1 = \partial\Sigma_2$ have this property. Then on-shell we have by Stokes' theorem

$$\Omega_{\Sigma_1} - \Omega_{\Sigma_2} = \int_{\Sigma_1} \omega - \int_{\Sigma_2} \omega = \int_{\mathcal{U}} \mathrm{d}\omega = 0. \tag{5.30}$$

Thus the presymplectic 2-form is the same on both surfaces. So the presymplectic 2-form can only change if the Σ_1, Σ_2 do not share a boundary. If that is true, then there is an additional codimension 1 surface $\mathcal{I}$ interpolating between $\partial\Sigma_1$ and $\partial\Sigma_2$ such that $\Sigma_1, \Sigma_2, \mathcal{I}$ bound a region $\mathcal{U}$, and we have

$$\Omega_{\Sigma_1} - \Omega_{\Sigma_2} = \int_{\Sigma_1} \omega - \int_{\Sigma_2} \omega = \underbrace{\int_{\mathcal{U}} \mathrm{d}\omega}_{=0} - \int_{\mathcal{I}} \omega. \tag{5.31}$$

So in this case the presymplectic 2-form changes by the integral of the presymplectic current over $\mathcal{I}$, which is sometimes called the presymplectic flux through $\mathcal{I}$.

This will be very important when we come to consider asymptotic symmetries. If we want to view the evolution from Σ_1 to Σ_2 as a symmetry, then we need to require that $\Omega_{\Sigma_1} = \Omega_{\Sigma_2}$ – because symmetries must conserve the presymplectic 2-form. Thus, such symmetries need to satisfy $\int_{\mathcal{I}} \omega = 0$, i.e. need to have zero presymplectic flux through $\mathcal{I}$. This restricts which transformations $\Sigma_1 \to \Sigma_2$ are symmetries, in a way that depends on properties of the fields at $\mathcal{I}$.

5.6 Possible ambiguities

The above construction as we have described it is not completely well-defined. There are two possible sources of ambiguity in the definition of the presymplectic 2-form.

1. First, we've only used the bulk Lagrangian form L. But L is only defined up to the addition of an exact $(D, 0)$-form, so $L \to L + \mathrm{d}K$ for some $(D-1, 0)$-form K. This is because we can simultaneously modify the boundary lagrangian by $l \to l - K$, and the action does not change:

$$S = \int_{\mathcal{M}} L + \int_{\partial\mathcal{M}} l \to \int_{\mathcal{M}} (L + \mathrm{d}K) + \int_{\mathcal{M}} (l - K) = S + \int_{\partial\mathcal{M}} K - \int_{\partial\mathcal{M}} K = S. \tag{5.32}$$

Because the action does not change, the dynamics will be completely unaltered. However, the vertical exterior derivative of the Lagrangian does change:

$$\delta L \to \delta L + \delta(\mathrm{d}K) = E \cdot \delta\phi + \mathrm{d}(\theta - \delta K)\,. \tag{5.33}$$

The equations of motion thus stay the same, but $\theta \to \theta - \delta K$. Then

$$\Theta_\Sigma \to \Theta_\Sigma - \int_\Sigma \delta K. \tag{5.34}$$

This seems like it could be an ambiguity, but actually, the presymplectic 2-form itself does not change:[3]

$$\Omega_\Sigma = \delta\Theta_\Sigma \to \delta\Theta_\Sigma + \int_\Sigma \underbrace{\delta^2}_{=0} K. \tag{5.35}$$

[3] Note that the vertical derivative does not necessarily commute with spacetime integration. Indeed, one can check that $\delta \int_S = (-1)^k \int_S \delta$, where k is the codimension of S.

So there is no problem here.

2. Second, consider again the fundamental relation

$$\delta L = E \cdot \delta\phi + \mathrm{d}\theta\,. \tag{5.36}$$

We have been using this to define the $(D-1,1)$-form θ. However, this relation is also satisfied if we change $\theta \to \theta + \mathrm{d}\alpha$ for some $(D-2,1)$-form α, since

$$\mathrm{d}\theta \to \mathrm{d}(\theta + \mathrm{d}\alpha) = \mathrm{d}\theta + \underbrace{\mathrm{d}^2}_{=0}\alpha. \tag{5.37}$$

So θ is only determined up to the addition of a spacetime-exact form. Under this ambiguity we have

$$\Theta_\Sigma \to \Theta_\Sigma + \int_\Sigma \mathrm{d}\alpha = \Theta_\Sigma + \int_{\partial\Sigma} \alpha, \tag{5.38}$$

and therefore

$$\Omega_\Sigma \to \Omega_\Sigma + \delta\Big(\int_{\partial\Sigma} \alpha\Big) = \Omega_\Sigma + \int_{\partial\Sigma} \delta\alpha. \tag{5.39}$$

So in this case the presymplectic 2-form itself *does* change, by a term at the boundary $\partial\Sigma$ of the Cauchy surface. Thus, we have a genuine ambiguity in the formulation.

5.7 Boundary conditions

Usually, $\mathcal{I}$ will coincide with the boundary $\partial\mathcal{M}$. Recall from the last lecture that we need to impose boundary conditions at $\partial\mathcal{M}$, and include a boundary action, in order to have a well-defined variational principle. We can use these to fix the ambiguity in the presymplectic 2-form.

Since we are assuming $\mathcal{M}$ has a Cauchy surface, it must be globally hyperbolic, so we can foliate it by Cauchy surfaces and write $\mathcal{M} = \mathbb{R}_t \times \Sigma$. Let Σ_t be the surface Σ at a fixed value of t, which we take as a time label. Consider the evolution between an initial time t_0 and a final time t_1, which takes place in the submanifold $[t_0, t_1] \times \Sigma$. The action in this region is

$$S_{t_0,t_1} = \int_{[t_0,t_1]\times\Sigma} L + \int_{\partial([t_0,t_1]\times\Sigma)} l \tag{5.40}$$

$$= \int_{[t_0,t_1]\times\Sigma} L + \int_{[t_0,t_1]\times\partial\Sigma} l + \int_{\Sigma_{t_1}} l - \int_{\Sigma_{t_0}} l. \tag{5.41}$$

Note that $[t_0, t_1] \times \partial\Sigma \subset \partial\mathcal{M}$.

We can now increase the extent to which we require the variational principle to be well-defined. In particular, we require it to be well-defined between any arbitrary times t_0 and t_1. The variation of the action is

$$\delta S_{t_0,t_1} = \int_{[t_0,t_1]\times\Sigma} E\cdot\delta\phi + \int_{[t_0,t_1]\times\partial\Sigma} (\theta + \delta l) + \int_{\Sigma_{t_1}} (\theta + \delta l) - \int_{\Sigma_{t_0}} (\theta + \delta l). \tag{5.42}$$

It is instructive to compare $\delta S_{t_0,t_1}$ to the variation of the action s_{t_0,t_1} for a general Hamiltonian mechanical system with generalised coordinates and momenta q_a, p^a. The action for such a system between times t_0 and t_1 is

$$s_{t_0,t_1} = \int_{t_0}^{t_1} \mathrm{d}t \left(\sum_a p^a \dot{q}_a - H(p,q) \right), \tag{5.43}$$

where $H(p, q)$ is the Hamiltonian, and its variation may be written (after integrating by parts)

$$\delta s_{t_0,t_1} = \int_{t_0}^{t_1} \sum_a \left[\delta p^a \left(\dot{q}_a - \frac{\partial H}{\partial p^a} \right) - \delta q_a \left(\dot{p}^a + \frac{\partial H}{\partial q_a} \right) \right] + \left[\sum_a p^a \delta q_a \right]_{t=t_1} - \left[\sum_a p^a \delta q_a \right]_{t=t_0} . \quad (5.44)$$

Let us go on-shell, i.e. assume Hamilton's equations of motion are satisfied. Then we are just left with

$$\delta s_{t_0,t_1} = \left[\sum_a p^a \delta q_a \right]_{t=t_1} - \left[\sum_a p^a \delta q_a \right]_{t=t_0} . \quad (5.45)$$

Note that these remaining two terms are the symplectic potential $\Theta = \sum_a p^a \, \mathrm{d}q_a$ at t_1 minus the symplectic potential at t_0 (evaluated against the variation δq_a).

In order for the Hamiltonian interpretation of the action S_{t_0,t_1}, we need an analogous result to hold. In fact, we can use this to determine what the "true" symplectic potential is, and this turns out to fix the ambiguity.

Going on-shell (i.e. setting $E = 0$), we have

$$\delta S_{t_0,t_1} = \int_{[t_0,t_1]\times\partial\Sigma} (\theta + \delta l) + \int_{\Sigma_{t_1}} (\theta + \delta l) - \int_{\Sigma_{t_0}} (\theta + \delta l). \quad (5.46)$$

In order for this to match with (5.45), it should only have a dependence on the fields at $\Sigma_{t_1}, \Sigma_{t_0}$. But there is an integral over $[t_0, t_1] \times \partial\Sigma$ that seems to ruin this.

However, there would be no problem if we could write

$$\int_{[t_0,t_1]\times\partial\Sigma} (\theta + \delta l) = F_{t_1} - F_{t_0} \quad (5.47)$$

for some function F_t that depends only on the fields on Σ_t.

$$\delta S_{t_0,t_1} = F_{t_1} + \int_{\Sigma_{t_1}} (\theta + \delta l) - F_{t_0} - \int_{\Sigma_{t_0}} (\theta + \delta l), \quad (5.48)$$

and by comparison with (5.45) we would find that the true presymplectic potential is $F_t + \int_{\Sigma_t} (\theta + \delta l)$. In fact, by the form of the left hand side of (5.47), F_t can only depend on the fields at $\partial\Sigma_t$, and in particular it should be the integral of some local $(D-2, 1)$-form C over $\partial\Sigma_t$, so

$$F_t = \int_{\partial\Sigma_t} C. \quad (5.49)$$

Then the symplectic potential would be

$$\Theta_\Sigma = \int_\Sigma (\theta + \delta l - \mathrm{d}C). \quad (5.50)$$

(The minus sign in front of $\mathrm{d}C$ comes from accounting for orientations.)

Let us note that it is possible to choose C appropriately in such a way that (5.47) and (5.49) are equivalent to

$$\theta + \delta l = \mathrm{d}C \, . \quad (5.51)$$

Note that this only needs to hold when pulled back to $[t_0, t_1] \times \partial\Sigma$, and on-shell $E = 0$.

In general, it would be too much to ask for (5.51) to hold without some additional conditions at the boundary $[t_0, t_1] \times \partial\Sigma$. But *unless* (5.51) holds, the variational principle will not be well-defined, and the Hamiltonian interpretation will not apply. Thus, we must have boundary conditions such that (5.51) holds. Different choices of boundary condition lead to different C.

This observation is enough to fix the ambiguity noted above. To see this, note that under the ambiguity transformation $\theta \to \theta + \mathrm{d}\alpha$, we have $C \to C + \alpha$. Thus, (5.50) is invariant under this transformation, and so unambiguous. Comparing with what we claimed for the symplectic potential previously,

$$\Theta_\Sigma = \int_\Sigma \theta, \tag{5.52}$$

we see there are two extra terms. The first is $\int_\Sigma \delta l$ – we can just ignore this term. The reason is that when computing the presymplectic form $\Omega_\Sigma = \delta\Theta_\Sigma$, this term contributes

$$\delta\Big(\int_\Sigma \delta l\Big) = \int_\Sigma \underbrace{\delta^2}_{=0} l = 0. \tag{5.53}$$

So it doesn't affect Ω_Σ.

The other term is $\int_\Sigma \mathrm{d}C$. In order for (5.50) to agree with what we originally claimed for the presymplectic potential, we need to be able to set this term to zero. And, in fact, we can do this: we can just apply the transformation $\theta \to \theta + \mathrm{d}\alpha$ with $\alpha = -C$. Under this transformation we have $C \to 0$. The overall value of (5.50) does not change – but now its form agrees with what we had previously.

This then gives us the way to fix the ambiguity. We pick θ (within the class of possible choices under $\theta \to \theta + \mathrm{d}\alpha$) in such a way that $\theta + \delta l|_{[t_0,t_1]\times\partial\Sigma} = \mathrm{d}C = 0$, when the boundary conditions and equations of motion are satisfied. This then gives the correct presymplectic form. Any other non-trivially distinct choice of θ would result in a $\mathrm{d}C \neq 0$, and so an incorrect presymplectic form.

For an example, let us consider the scalar field. The bulk Lagrangian is

$$L = -\frac{1}{2}\Big(\eta^{\mu\nu}\partial_\mu\phi\partial_\nu\phi + m^2\phi^2\Big)\,\mathrm{d}^D x\,. \tag{5.54}$$

We showed in the last lecture that we can pick $l = 0$, and $\theta = \delta\phi * \mathrm{d}\phi$. This is consistent with boundary conditions are such that ϕ is fixed at $\partial\Sigma$, and we see that $\theta + \delta l = 0$ at $[t_0, t_1] \times \partial\Sigma$. Thus, $\mathrm{d}C = 0$, and so the correct, unambiguous symplectic potential is

$$\Theta_\Sigma = \int_\Sigma \delta\phi * \mathrm{d}\phi\,. \tag{5.55}$$

5.8 Covariant phase space

The space of the solutions equipped with the 2-form Ω_Σ is called the covariant phase space, or in the case when Ω_Σ is presymplectic it is sometimes called the covariant pre-phase space. The presymplectic case occurs when there are gauge symmetries, and we need to carry out a reduction procedure over these gauge symmetries in order to get the true covariant phase space. We will discuss this in the next section, as well as some specific examples. As we will see, the resulting symplectic space will turn out to be equivalent to the canonical phase space.