Covariant Phase Space Methods in Gravity

OR, THE GEOMETRY OF GEOMETRY

David Grabovsky

February 20, 2023

Abstract

We develop the geometrical formulation of classical mechanics. We begin by reviewing the Lagrangian and Hamiltonian mechanics of particles in the language of bundles and symplectic geometry. We then describe field theories in the same vein, paying special attention to boundary conditions. Our goal is the manifestly covariant construction of a phase space for classical relativistic field theories; the fruits of our labor are found in the application of the methods we develop to the general theory of relativity. Our purpose is to "finally" explain how classical physics works, and to understand how to geometrize the theory of spacetime geometry itself.

Contents

0 Introduction **2**

1 Particle Mechanics **2**
1.1 Philosophy 2
1.2 Lagrangian Mechanics 3
1.3 The Hamiltonian Formalism 4

2 Plan of Attack **7**

3 Covariant Phase Space **9**
3.1 Boundary Palooza 9
3.2 The Symplectic Form 10
3.3 The Hamiltonian 12

4 Gravity at Last **14**
4.1 The Action 14
4.2 The Hamiltonian 15
4.3 Gauge Symmetry 16

0 Introduction

A single field permeates space; it fluctuates, untethered, through the vast emptiness of the cosmos. The faintest shadow of a smile fills the void: somewhere far away, a student has just caught their first glimpse of Hamilton's equations in all their symplectic glory. At first dimly, then all at once with blinding clarity, the entire cotangent bundle reveals itself, adorned with forms and vector fields ethereal; the Hamiltonian flow revolves majestically, wheels within wheels; and the sheer beauty of the canonical invariants... The student sighs deeply, feels an overwhelming sense of contentment and peace, and is, at last, truly happy.

This is not the story of the student, who—unlike the field—is merely an idealization dreamed up in the twisted minds of physicists. This is the story of the field.

Discussions of the dynamics of relativistic fields often introduce the Hamiltonian formalism quickly after the Lagrangian one, and immediately set off doing computations. While the Lagrangian description is always written down in a manifestly covariant way, the Hamiltonian viewpoint encounters problems because the definitions $\mathcal{H} = \pi\dot{\phi} - \mathcal{L}$ and $\pi = \frac{\partial \mathcal{L}}{\partial \dot{\phi}}$ explicitly pick out a preferred time coordinate and slicing, and none of the computations that follow look very covariant. But the ideas behind Hamiltonian mechanics are too simple and too elegant *not* to be treated correctly, "without the dots," and in what follows we will present at least part of a proper attempt to do things right. The resulting set of techniques, called *covariant phase space* (CPS) methods, give a transparent and elegant way to understand classical dynamics. In particular, CPS provides a geometrical perspective on general relativity, which is the classical field theory of the world's underlying geometry.

Motivated by the mantra that geometry should be understood geometrically, we will begin (section 1) by reformulating traditional particle mechanics in the language of symplectic geometry. We will then (section 2) spend some time strategizing about how to upgrade this machinery to the case of fields, develop a general plan of attack, and try the plan out. We will expend most of our effort (section 3) in making the transition from Lagrangian field theory to a Hamiltonian point of view. Much care will be taken in formulating the variational problem upon which everything is based, and we will pay special attention to boundary terms. These boundary terms will return (section 4) in a crucial way when we apply CPS to the general theory of relativity and finally see the full power of the formalism.

1 Particle Mechanics

1.1 Philosophy

In kindergarten, we learn that one may use the Euler-Lagrange equations to write down the time evolution of a mechanical system; very often this is all we need to know, and kindergarten is as far as we get. But there is a more subtle idea: one can compare where the system *is* at a given instant to where it is *going*. The resulting axes of "being and becoming" (hereafter: position q and momentum p) define what we call phase space; in my opinion, the phase-space trajectory of a system gives a more distilled understanding of its dynamics.

Of course, after kindergarten we all go to high school, where we learn about the Hamiltonian H and the Poisson bracket $\{\cdot, \cdot\}$. Packaging the positions and momenta into a set of unified coordinates called ζ, we can rewrite the equations of dynamics as $\dot{\zeta} = \{\zeta, H\}$. This first-order differential equation gives us an inkling that the trajectories $\zeta(t)$ are really the integral curves of some "guiding" vector field, illustrated below for the harmonic oscillator. We will call this the *Hamiltonian vector field*, and below we develop more concretely how it actually performs the phase-space evolution of any mechanical system.

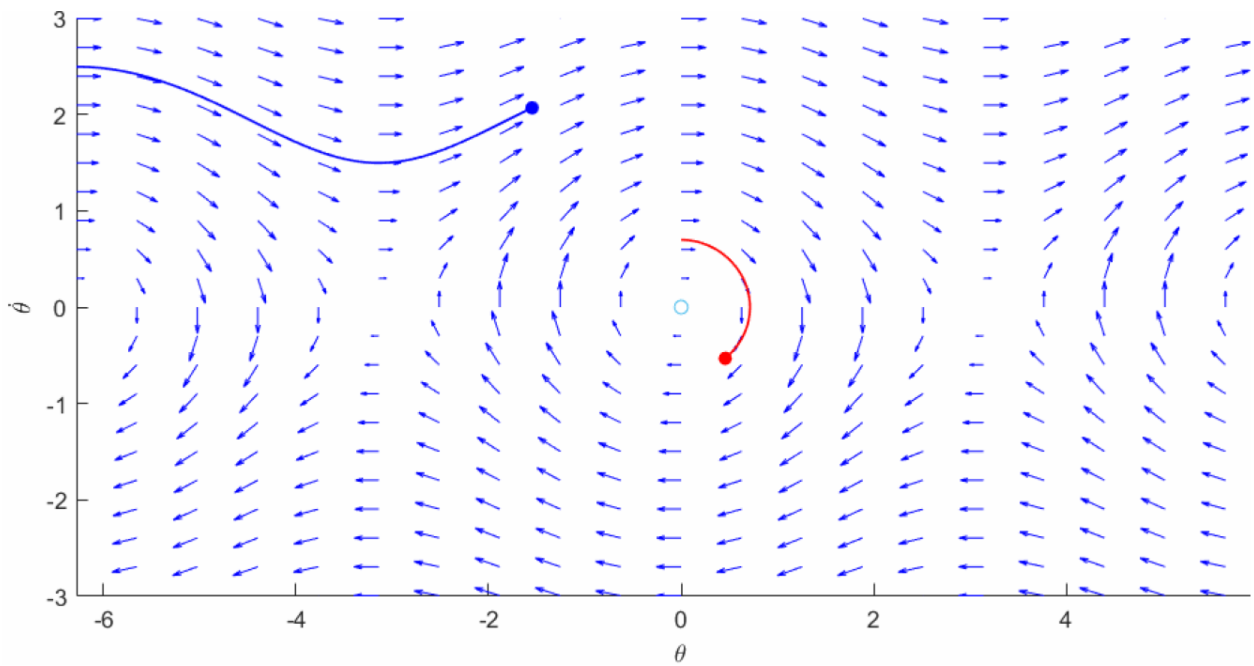

1.2 Lagrangian Mechanics

Let M be the universe, a smooth manifold of dimension n. To begin, we will imagine $M = \mathbb{R}^n$ to be space. We consider a single particle whose motion traces out a curve $q\colon \mathbb{R} \longrightarrow M$, so that its position at any $t \in \mathbb{R}$ is a point $q(t) \in M$. If the particle is at $q \in M$, we can use a set of n local coordinates[1] q^i to describe q. The particle's velocity is a tangent vector (q, v) in the tangent space T_qM to M at q. Given coordinates q^i of q, T_qM has a natural basis $\{\frac{\partial}{\partial q^i} = \partial_i\}$ in which to decompose $v = v^j \partial_j$. Now a pedantic but important aside: note that the numbers v^j are merely the coefficients at ∂_i needed to specify an arbitrary tangent vector at a point $q \in M$, and can be chosen independently of the point's coordinates q^i to yield different tangent vectors. Only when the vector lies tangent to a given curve, which we might call $q(t)$, do the components v^j actually equal the time derivatives $\dot{q}^j(t)$ of $q(t)$.

[1]Henceforth, we use points and their coordinates interchangeably. It is not strictly true that $q^i \in \mathbb{R}^n$ lies in M, but anyone upset about such things can direct their complaints to the nearest math department.

Due to a disastrous historical misunderstanding, the symbol $\dot{q}^j$ was kept as the *name* of the coefficient v^j, leaving many students to wonder why q and $\dot{q}$ are treated independently in mechanics courses. We will use the standard notation $\dot{q}^j$ instead of v^j, but one should keep in mind that $\dot{q}^j = \dot{q}^j(t)$ only when the vector being described is the velocity of $q(t)$ itself.

One often proclaims[2] that a particle's position q_0 and velocity $\dot{q}_0$ at one instant suffice to determine its entire motion $q(t)$; one concludes that $q(t)$ is determined by a second-order ODE, and then writes down Newton's law before leaving in a hurry to see if there are still cookies at department teatime. But at teatime, one realizes that q_0 and $\dot{q}_0$ determine not only $q(t)$, but also $\dot{q}(t)$. Indeed, $q(t)$ and $\dot{q}(t)$ are co-constitutive and effect each other's dynamics, so our real interest is in the combined trajectory of positions and velocities $(q(t), \dot{q}(t)) \subset TM$. In this new setting, Lagrangian mechanics basically follows as usual. We quantify local deviations of $(q(t), \dot{q}(t))$ from their physical paths by a smooth function $L\colon TM \longrightarrow \mathbb{R}$, the Lagrangian of the system. The action integral $S = \int_{t_i}^{t_f} \mathrm{d}t\, L(q(t), \dot{q}(t))$ adds up local contributions to L along a path, and we require it to be stationary on physical paths:

$$\delta S = \int_{t_i}^{t_f} \mathrm{d}t \left[\frac{\partial L}{\partial q^i} - \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial L}{\partial \dot{q}^i}\right)\right]\delta q^i + \left[\frac{\partial L}{\partial \dot{q}^i}\delta q^i\right]\Bigg|_{t_i}^{t_f} = 0 \implies \frac{\partial L}{\partial q^i} = \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial L}{\partial \dot{q}^i}\right). \tag{1.1}$$

Here the techniques of variational calculus have been upgraded: we consider $q \mapsto q^i(t)$ as scalar functions on M; these are zero-forms, and the δq^i are their exterior derivatives. As we will see later, integration by parts becomes an application of Stokes's theorem. We have assumed that δq^i vanishes at t_i and t_f; this assumption amounts to an imposition of boundary conditions that renders the variational problem $\delta S = 0$ well-posed. We understand $\partial L/\partial q^i$ and $\partial L/\partial \dot{q}$ as the partials of L with respect to the base and fiber coordinates of TM, respectively. Examples of interesting Lagrangians abound, as do generalizations that toy with the configuration manifold M; moreover, there are powerful reformulations of the variational principle in terms of geodesic spray. But let us restrain ourselves.

1.3 The Hamiltonian Formalism

In Hamiltonian mechanics, we replace the velocities $\dot{q}$ by their canonical momenta, defined by $p \equiv \partial L/\partial \dot{q} = (\partial/\partial \dot{q})(L) = \mathrm{d}L(\partial/\partial \dot{q})$. Here we have written p as the action of the tangent basis vector $\partial/\partial \dot{q}$ on L and recognized that p is a function of the tangent vectors $\dot{q}$. In other words, p is a one-form on M and is thus an element of the cotangent bundle T^*M. In coordinates, its action on $\dot{q} \in TM$ is given by

$$p(\dot{q}) = \big(p_i\, \mathrm{d}q^i\big)\left(\dot{q}^j \frac{\partial}{\partial q^j}\right) = p_i q^j\, \mathrm{d}q^i\left(\frac{\partial}{\partial q^j}\right) = p_i \dot{q}^i. \tag{1.2}$$

The key realization here is the same one that motivates Lagrangian mechanics: as the point traces out a path $q(t)$, the changing velocity vector $\dot{q}(t)$ forces $p = \frac{\partial L}{\partial \dot{q}}$ to change in

[2]This is an experimental fact; however, the Ostrogradsky instability gives a partial theoretical explanation.

time as well. Thus we might say that the physical path on M *lifts* to a trajectory $(q(t), p(t))$ in $T^*M \equiv \mathcal{M}$, which we will hereafter call *phase space*. The Hamiltonian vector field X from earlier is then just the velocity vector tangent to this phase-space trajectory; we write $X \in T\mathcal{M}$. Hamilton's equations are nothing more than the flow by X, and $\mathcal{M}$ itself is nothing more than the set of all possible integral curves of X, i.e. the space of solutions to the equations of motion. This is a devastatingly simple and concise description of classical mechanics, and we will do well to develop it further.

To make progress, we will construct the *canonical one-form*. The basic idea is that there is a natural way for the Hamiltonian vector field X to "eat itself." The natural projection $\pi\colon \mathcal{M} = T^*M \longrightarrow M$ sends $(q, p) \mapsto q$; correspondingly, its linearization $\mathrm{d}\pi\colon T\mathcal{M} = T(T^*M) \longrightarrow TM$ sends tangent vectors on $\mathcal{M}$ attached at (q, p) to tangent vectors on M attached at q. What would happen if we fed one of the resulting vectors $\mathrm{d}\pi(X) \in TM$ into the one-form p where the original vector X was attached? This is a coordinate-free construction that begins with a vector $X \in T\mathcal{M}$ and uses one-forms on $\mathcal{M}$ to yield a number $\theta(X) \equiv p(\mathrm{d}\pi(X))$. The object θ is therefore a one-form on $\mathcal{M}$; we call it the canonical or tautological one-form, or the *symplectic potential*. If X happens to be Hamiltonian, then $\mathrm{d}\pi$ sends it to the velocity $\dot{q}$ tangent to $q(t)$. In coordinates, $\theta(X) = p\dot{q}$ recovers an expression ubiquitous in mechanics; we have finally seen where it comes from!

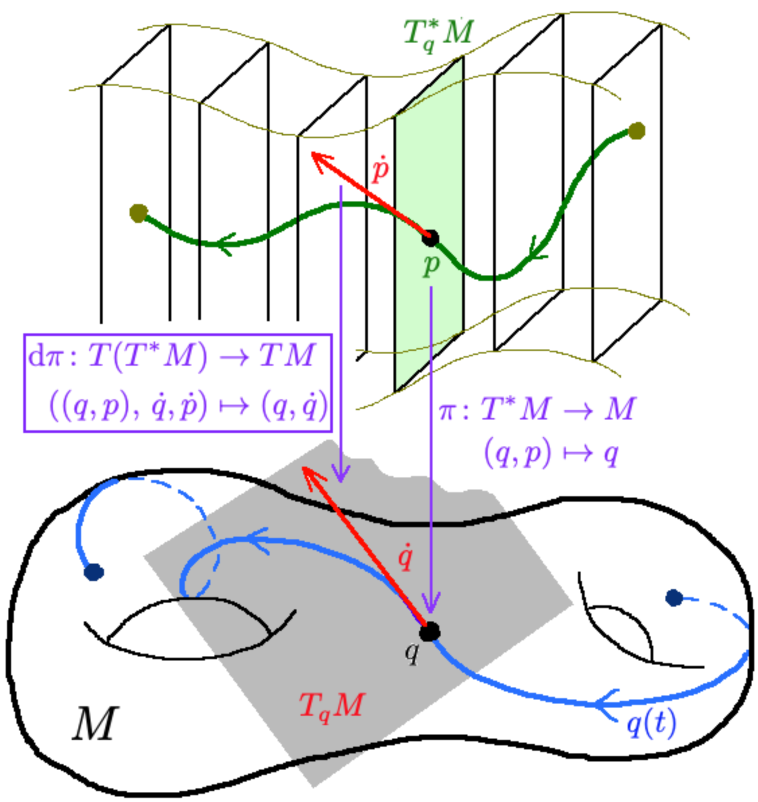

Hamiltonian mechanics is usually set up using a function H rather than a vector field X: how do we get the former from the latter? It is easier to first convert X into a one-form, since vectors and one-forms are dual, and thence to a function. We thus seek an isomorphism $T\mathcal{M} \longrightarrow T^*\mathcal{M}$, and in fact such an isomorphism is given by the two-form $\omega \equiv \mathrm{d}\theta$. Two-forms have two slots for vectors; sticking X into the first slot leaves the other slot unfilled, so the object $i_X\omega$ is a one-form, as desired. That $i_X\omega$ is unique to X can be seen in coordinates:

$$\theta = p_\alpha \mathrm{d}q^\alpha \in \Omega^1(\mathcal{M}) \implies \omega = \mathrm{d}\theta = \mathrm{d}p_\alpha \wedge \mathrm{d}q^\alpha \in \Omega^2(\mathcal{M}). \tag{1.3}$$

One shows that ω is nondegenerate by proving that its kernel is trivial: $i_X\omega$ does not vanish identically if $X \neq 0$. (The proof observes that the matrix of ω has nonzero determinant.) Note also that ω is closed, since $\mathrm{d}\omega = \mathrm{d}(\mathrm{d}\theta) = 0$. Any closed, nondegenerate two-form ω is called a *symplectic form*, and its existence of is the key to Hamiltonian mechanics on $\mathcal{M}$.

With some more work, it can proven that $i_X\omega$ is not only unique to X, but is also exact: it can be expressed as a differential $\mathrm{d}H$, where the function H is the Hamiltonian! In coordinates, we find that the equation $i_{X_H}\omega = \mathrm{d}H$ expresses Hamilton's equations:

$$\mathrm{d}H = \left(\frac{\partial H}{\partial q^\alpha}, \frac{\partial H}{\partial p_\alpha}\right) \implies X_H = \left(\frac{\partial H}{\partial p_\alpha}, -\frac{\partial H}{\partial q^\alpha}\right) \stackrel{!}{=} (\dot{q}, \dot{p}). \tag{1.4}$$

Mathematical treatments of this subject usually reverse the story above.[3] They begin with a function H and compute its differential $\mathrm{d}H$. Then ω, being nondegenerate, has an inverse[4] that yields a unique X_H. In this sense a vector field is *defined* to be Hamiltonian if its contraction with the symplectic form is exact, while one whose contraction is merely closed is called a *symplectic vector field*. These generate symplectomorphisms, or what physicists call canonical transformations. All Hamiltonian vector fields are symplectic, and viewing them as such is what Hamilton-Jacobi theory does in seeing Hamiltonian evolution as a type of canonical transformation. Some intuition might help here: H gives the total energy of a system. Since time evolution conserves energy, X_H should lie orthogonal to the gradient $\mathrm{d}H$. This is exactly what the calculation $\omega(X_H, X_H) = \mathrm{d}H(X_H) = 0$ encodes.

This formalism may seem heavy-handed, but it begets spectacularly concise and transparent computations. For instance, here is a one-line proof of Liouville's theorem:

$$X \text{ symplectic} \implies \mathcal{L}_X\omega = \mathrm{d}i_X\omega + i_X\mathrm{d}\omega = \mathrm{d}(\text{closed}) + i_X(0) = 0 \implies \mathcal{L}_{X_H}\omega^n = 0. \tag{1.5}$$

One way to understand this is by analogy to electromagnetism. Both the vector potential A and the symplectic potential θ are one-forms, and their exterior derivatives $F = \mathrm{d}A$ and $\omega = \mathrm{d}\theta$ may be thought of as a "field strength" or "gauge curvature." In some sense, ω

[3]If the discussion above has been slightly artificial, it is because I have introduced symplectic geometry inside out. Mathematicians are therefore advised to read this section while standing on their heads.

[4]In coordinates, one notices that this inverse is the Poisson bracket! The "algebraic" view from $\{\cdot,\cdot\}$ and H, developed in physics courses, is equivalent to the "geometric" view from ω and X_H, developed here.

curves phase space by generating the vector fields that guide particles along their winding paths. As to the proof, we have taken the Lie derivative[5] of ω using Cartan's magic formula $\mathcal{L}_X = \mathrm{d}i_X + i_X\mathrm{d}$. The equation $\mathcal{L}_{X_H}\omega = 0$ may then be understood in analogy to the Killing equation $\mathcal{L}_X g = 0$: the symplectic form (respectively, the metric) is conserved under Hamiltonian flow (resp. flow by isometries). Since $\omega \in \Omega^2(\mathcal{M}^{2n})$ is nondegenerate, the top form ω^n is nonvanishing and is therefore a volume form. Liouville's theorem states that the symplectic structure of $\mathcal{M}$ conserves ω, and therefore leaves phase-space volume invariant!

The main takeaway here is that the Hamiltonian vector field X_H, which directly performs physics on $\mathcal{M}$, is the main hero of our story. X_H is supported by its sidekick ω, whose existence guarantees Hamiltonian mechanics on $\mathcal{M}$. As a consequence, we can dispense entirely with cotangent bundles: it is enough for phase space to be a *symplectic manifold*.

2 Plan of Attack

Before we discuss symplectic geometry for field theories, we caution that fields are quite different from particles. Traditionally, fields are introduced as continuum generalizations of many-particle systems; the value of a field $\phi(x)$ is the density of particles at x. But behind this scaffolding, the mathematical definitions tell a different story. A particle's trajectory $q\colon \mathbb{R} \longrightarrow M$ takes in a single parameter (e.g. proper time) and produces positions in M. Meanwhile, a field takes all of the spacetime coordinates as parameters, and its output lies in a distinct manifold C of possible field values: $C = \mathbb{R}$ for a real scalar, and so on.

It is natural to unify spacetime and field values in a C-bundle of *field histories* over M. The concept of such *field bundles* leads naturally to the study of jet bundles, which in turn are just thicc tangent bundles.[6] The action principle then has to account for both "horizontal" and "vertical" variations in M and C respectively. These are packaged neatly in a *variational bicomplex*, which we will not construct; however, we cannot do without its basic ingredients. We distinguish between the exterior derivative δ on C (to be thought of as a field variation) and the exterior derivative d on M (to be thought of as the usual differential). Most of the objects we deal with will be forms of different ranks on M and C.

Heeding this warning, we will follow Harlow and Wu (2019) in constructing a symplectic form and Hamiltonian on the phase space of a field theory. We will follow six steps:

1. The *configuration space* $\widetilde{\mathcal{C}}$ is the set of fields ϕ^a satisfying the boundary conditions, on which we will elaborate. Classically, this is the space in which we consider local variations; quantum-mechanically, this is the space on which we compute path integrals.

[5] In Soviet Russia, $\mathcal{L}_X$ was called the "fisherman's derivative." The idea was that differential-geometric objects flow down a river, while a fisherman sits on the bank and differentiates them as they go by.

[6] Tangent bundles are the natural setting for ODEs, which are geometrized in the notion of "geodesic spray." One introduces jet bundles when there are more coordinates with respect to which to be tangent. In this case one speaks of the "prolongation" of jet bundles to geometrize PDEs.

2. The *pre-phase space* $\widetilde{\mathcal{P}}$ is the set of fields ϕ^a for which the equations of motion hold, in line with our description of $\mathcal{M}$ as the set of phase-space trajectories. This definition does not pick out a preferred time slice, nor does it choose (ϕ, π)-style coordinates.

3. We will construct the *pre-symplectic form* $\widetilde{\Omega}$. We will think highly of ourselves and nod pretentiously, only to be promptly removed from our pedestal. $\widetilde{\Omega}$ will be closed by fiat, but it will be degenerate, and will not be an isomorphism as described above.

4. The problem is gauge redundancy: two nearby field configurations $\phi, \phi' \in \widetilde{\mathcal{P}}$ may represent the same physical state. A vector $Z \in T\widetilde{\mathcal{P}}$ pointing from ϕ to ϕ' is a "degeneracy direction" in $\widetilde{\mathcal{P}}$, and Hamiltonian evolution by Z is "fake" in the sense that $i_Z\widetilde{\Omega} = 0$. We call such Z *zero modes* of $\widetilde{\Omega}$, and together they form a Lie algebra $\mathcal{Z}$ that generates a group $G \subset \mathrm{Diff}(\widetilde{\mathcal{P}})$ of physical equivalences or gauge symmetries.

5. We perform *symplectic reduction* to get rid of the zero modes, schematically writing $\mathcal{P} = \widetilde{\mathcal{P}}/G$ and $\widetilde{\Omega}/G$. The first expression is well-defined as written: we "glue" all physically equivalent field configurations to each other along the degeneracy directions Z.[7] The second expression "does the same thing" to the action of Ω on vector fields, but first technically requires us to quotient the algebra of vector fields $T\mathcal{P}$ by G.

6. Having obtained a true phase space and symplectic form for fields, it is straightforward to construct a Hamiltonian for the theory from a Hamiltonian vector field. Nevertheless, there are extra subtleties due to the difference between spacetime and field space.

Let us illustrate step (3), which will occupy most of our time, in the case of particle mechanics on a spatial manifold. We continue a computation we began in §1.2:

$$\delta S = \int_{t_i}^{t_f} \mathrm{d}t \left[\frac{\partial L}{\partial q} - \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial L}{\partial \dot{q}}\right)\right]\delta q + \left[\frac{\partial L}{\partial \dot{q}}\delta q\right]\Bigg|_{t_i}^{t_f} = \int_M (E\,\mathrm{d}t)\,\delta q + \int_{\partial M} (p\,\delta q). \tag{2.1}$$

We package the equations of motion as a one-form $E\,\mathrm{d}t$ and keep the oft-forgotten boundary term. We write the same expression pretentiously: M is the manifold of parameters, here just the time interval $[t_i, t_f]$, taken in by the $(0+1)$-dimensional "field" q. Evaluation at the endpoints t_i and t_f is really an integral over the boundary of the parameter manifold; the boundary integrand is the symplectic potential from above! By Stokes's theorem, we have

$$\int_{\partial M} (p\,\delta q) = \int_{\partial M} \theta = \int_M \omega. \tag{2.2}$$

[7] This construction was brought to you by the Frobenius gang. The closure of the Zs under the Lie bracket guarantees that $\mathcal{Z}$ generates a distribution on $\widetilde{\mathcal{P}}$ spanned by the zero modes. By Frobenius's theorem, this foliates $\widetilde{\mathcal{P}}$ into gauge orbits, along which the gluing proceeds and whereby the quotient is well-defined.

Thanks to Stokes's theorem and the fact that spacetime is the configuration space for particles, the symplectic form can be found hiding in the action principle defining mechanics. As we will find shortly, essentially the same procedure will work for fields, with a few differences.

3 Covariant Phase Space

After much procrastination and moralizing, we will construct the symplectic form on the phase space of a relativistic classical field theory. We'll begin by writing down the action. It will be useful to view the Lagrangian as a top-dimensional form on M, suitable for integration on spacetime, instead of as a scalar function; the difference is immaterial by Poincarè duality. We write $L = \mathscr{L}\varepsilon_{M}$, where ε_{M} is the standard volume form on M (we restrict to orientable spacetimes), and where $\mathscr{L} : \widetilde{\mathcal{C}} \longrightarrow \mathbb{R}$ carries the same information as L.[8] The action is

$$S = \int_M L(\phi, \chi) + \int_{\partial M} \ell(\phi, \chi). \tag{3.1}$$

We write ϕ for all of the dynamical fields, i.e. those whose variations and equations of motion interest us, and χ for all of the non-dynamical or background fields, e.g. a fixed background metric. We allow M to have boundary and include ℓ in the action. If $\partial M \subset M$, the second term is technically included in the first as its pullback to ∂M, but we want to see explicitly what happens to various boundary terms. We will henceforth abuse notation in this way.

3.1 Boundary Palooza

Naïvely, we would set $\delta S = 0$. But here we must be careful: every well-posed variational problem comes with boundary conditions. As shown below, we decompose $\partial M = \Gamma \sqcup \Sigma_- \sqcup \Sigma_+$ into spatial (Γ) and temporal ($\Sigma_\pm$) boundaries. We fix the value of ϕ or its derivatives on Γ (this specifies the theory itself), but *not* on $\Sigma_\pm$ (this is more akin to preparing a state *within* the theory).[9] Thus the action is only stationary up to a term defined on $\Sigma_\pm = \Sigma_+ - \Sigma_-$ (the minus sign denotes orientation), where a "flux of δS" is permitted. Let us see this in action (pun intended):

$$\begin{aligned} S = \int_M L + \int_{\partial M} \ell \implies \delta S &= \int_M (E \wedge \delta\phi) + \int_{\partial M} (\theta + \delta\ell) = \\ &= \int_M (E \wedge \delta\phi) + \int_{\Sigma_\pm} (\theta + \delta\ell) + \int_\Gamma (\theta + \delta\ell) \implies \\ 0 = \int_M (E \wedge \delta\phi) = \int_\Gamma (\theta + \delta\ell) \implies E = 0, \quad &(\theta + \delta\ell)\Big|_\Gamma = \mathrm{d}C. \end{aligned} \tag{3.2}$$

[8] Notice that the distinctions between TM and T^*M have vanished in the absence of coordinates: there is only one true configuration space, and phase space sits *inside* it, not *over* it as in the case of particles.

[9] This is a feature of Lorentz signature, which makes many dynamical equations of motion wavelike. These hyperbolic PDEs solve a Cauchy problem, whose well-posedness requires the boundary conditions above.

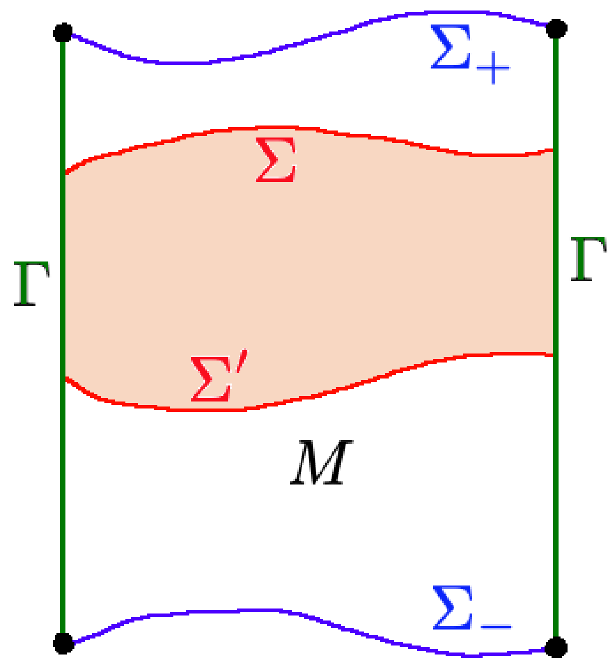

The calculation begins as in §2: we integrate by parts, isolate the equations of motion (labeled E), and pick up a boundary term to become part of the symplectic potential. There is also the variation $\delta\ell$. We split up the boundary as described above and require the first and third terms to vanish in accordance with our variational principle. The vanishing of the first term yields the equations of motion, which is a good sign. We might think to set $(\theta + \delta\ell)\big|_\Gamma = 0$ to make the third term vanish, but this is too strong: if $(\theta + \delta\ell)\big|_\Gamma$ is merely exact, then by Stokes's theorem

$$\int_\Gamma (\theta + \delta\ell) = \int_\Gamma \mathrm{d}C = \int_{\partial\Gamma} C, \tag{3.3}$$

and since $\partial\Gamma \subset \Sigma_\pm$, a contribution to δS from C is allowed by our boundary conditions!

3.2 The Symplectic Form

Now we apply a standard trick in physics: identify a quantity that naïvely vanishes, but which upon closer inspection is nonzero, and give it a fancy name. We obtain such a quantity by moving $\mathrm{d}C$ to the left-hand side in $(\theta + \delta\ell)\big|_\Gamma = \mathrm{d}C$:

$$\Psi \equiv \theta + \delta\ell - \mathrm{d}C \implies \Psi\Big|_\Gamma = (\theta + \delta\ell - \mathrm{d}C)\Big|_\Gamma = 0. \tag{3.4}$$

Ψ is our old friend θ, dressed up with the proper boundary terms: it is the symplectic potential. We then hurriedly define the *pre-symplectic current*[10] and *pre-symplectic form*:

$$\omega \equiv \delta\Psi\Big|_{\widetilde{\mathcal{P}}} = \delta(\theta - \mathrm{d}C)\Big|_{\widetilde{\mathcal{P}}}, \qquad \widetilde{\Omega} \equiv \int_\Sigma \omega, \tag{3.5}$$

where we have used $\delta^2 = 0$ to eliminate ℓ, and where Σ is a Cauchy surface.

Several comments, in no particular order, are by now long overdue:

- All objects in sight are differential forms on both M and $\widetilde{\mathcal{C}}$, where they may have different ranks. The table below summarizes their names and natures. Each form's "natural

[10]Physicists use "current" for densities whose integral is some kind of charge, and for objects dual to or sourcing some kind of potential. Here, different authors throw around a profusion of terms, some even calling ω a "symplectic potential current." Lacking a philosophy degree, I have chosen the most common name.

habitat" is its domain of definition, but they all have arbitrary smooth extensions into M. Such is the nature of globalization in today's compact but connected world.

Form	Habitat	Rank on M	Rank on $\widetilde{\mathcal{C}}$
L	M	d	0
ℓ	∂M	$d-1$	0
θ	∂M	$d-1$	1
C	Γ	$d-2$	1
ω	M	$d-1$	2

- Both Ψ and ω are "pinned down" at Γ, where they vanish by construction. C was defined only on Γ, so we extend it to M; however, only its value on Γ is physical.

- While it looks like ω does not depend on ℓ, it does so indirectly through C: the equation $(\theta + \delta\ell)\big|_{\Gamma} = \mathrm{d}C$ gives a consistency condition governing the kinds of boundary Lagrangians allowed by the boundary conditions.

- ω is *biclosed*, meaning that $\delta\omega = \mathrm{d}\omega = 0$. The first equality is obvious from the definition $\omega = \delta\Psi$, while the second follows from a one-line calculation:

$$\mathrm{d}\omega = \mathrm{d}\delta(\theta - \mathrm{d}C) = -\delta\mathrm{d}\theta = -\delta(\delta L - E \wedge \delta\phi) = \delta E \wedge \delta\phi = 0. \tag{3.6}$$

  Here we have used that δ and d anticommute (by definition), used $\mathrm{d}^2 = 0$ to get rid of C, and further substituted $\delta L = E \wedge \delta\phi + \mathrm{d}\theta$[11] in order to use δ^2 to get rid of L. In the last step, we recall that ω is really the restriction of Ψ to $\widetilde{\mathcal{P}}$, where $E \equiv 0$.

- The integration over Σ to define $\widetilde{\Omega}$ is a technicality that "undoes" the fact that ω is a $(d-1)$-form on M. In fact $\widetilde{\Omega}$ is independent of Σ; to see this, integrate ω over two Cauchy surfaces Σ, Σ' as shown above. Let M' and Γ' be the regions of M and Γ between Σ and Σ'. Then by Stokes's theorem and the previous properties,

$$\int_{\Sigma-\Sigma'} \omega = \int_{M'} \mathrm{d}\omega - \int_{\Gamma'} \mathrm{d}\omega = \int_{M'} 0 - \int_{\Gamma'} \mathrm{d}(0) = 0. \tag{3.7}$$

- One should understand the symplectic potential Ψ as an analog of the first variation δS, and the pre-symplectic current ω as an analog of the second variation "$\delta^2 S$."

Let us summarize our current state of affairs. Starting from a well-posed variational problem, we have constrained the boundary terms and constructed a pre-symplectic form $\widetilde{\Omega}$ by differentiating the symplectic potential θ, up to a total spacetime derivative $\mathrm{d}C$. After the quotient procedure described above, we obtain a bona fide phase space and a true symplectic form: we will henceforth assume that this procedure has been carried out.

[11] This is an example of integration by parts but without the integration, also called the product rule.

3.3 The Hamiltonian

All that remains is to specify a Hamiltonian vector field and function on phase space. Technically, we are free to specify any Hamiltonian function at all and be done, but this is rather reductionist. One might try to do better by developing a covariant notion of the Legendre transform, but while this may be possible, it is unclear how to treat both L and ℓ. As we will see, an approach based on spacetime symmetries will recover the correct form and yield additional physical insight that the Legendre transform would have missed.

In flat spacetime, the Hamiltonian generates time evolution and can be thought of as a "time charge." Moreover, there is a whole stress tensor's worth of charges that generate spacetime translations, rotations, and boosts. Our covariant framework will treat all of these equally, producing a "Hamiltonian,"[12] or more properly a *diffeomorphism Noether charge* for each one. This is tricky: if $\xi \in TM$ generates a diffeomorphism of M, the dynamical fields ϕ will flow along ξ together with all of the non-dynamical gunk (e.g. a fixed background metric) living on M, so we cannot claim to have evolved *only* the dynamical fields. We need a way to implement the flow of ξ on phase space—that is, to construct a vector field $X_\xi \in T\mathcal{P}$ that flows the dynamical fields $\phi \in \mathcal{P}$ by ξ and "ignores" the non-dynamical fields. Also, not just any diffeomorphism will do: fields can only meaningfully evolve along directions ξ that are *symmetries* of the theory: the action must be constructed in such a way that any non-dynamical fields enter in combinations that remain invariant under ξ-flow. (For example, if the metric is non-dynamical, its invariance forces ξ to be an isometry.) An additional technicality is that ξ must respect the boundary conditions: it cannot move the spatial boundary Γ, so we also require that $n_\mu \xi^\mu = 0$, where n_μ is the unit normal to Γ.[13]

Thus field evolution is dictated by the spacetime symmetries respected by the action. When we ask for a Hamiltonian vector field, we seek to implement evolution along ξ on phase space; thus our challenge is to construct, for each symmetry generator, a vector field $X_\xi \in T\mathcal{P}$ that flows *only* the dynamical fields by ξ. How does one flow? The answer, by definition, is to take a Lie derivative! Thus the (spacetime) variation of ϕ under the flow of the symmetry generator ξ is $\delta_\xi \phi = \mathcal{L}_\xi \phi$. Similarly, the (phase-space) variation of ϕ under the flow of the Hamiltonian vector field X_ξ is $\delta_{X_\xi}\phi = \mathcal{L}_{X_\xi}\phi = X_\xi(\phi)$, where the last equality follows because ϕ is a scalar on $\mathcal{P}$. We claim that the correct Hamiltonian vector field is

$$X_\xi = \int_M \mathcal{L}_\xi \phi^a \frac{\delta}{\delta \phi^a} \in T\mathcal{P} \iff X_\xi^a = \mathcal{L}_\xi \phi^a. \tag{3.8}$$

This vector field is constructed from the basis vectors $\delta/\delta\phi^a$ on configuration space; its *coefficients* are $\mathcal{L}_\xi \phi^a$. This makes precise the notion of "importing" the flow by ξ into phase

[12] And a Momentonian, an Angletonian, and a Boostonian...

[13] This requirement is due to considerations involving L. Considerations involving ℓ introduce extra conditions, which we omit; this is due to considerations involving brevity and sanity.

space. The dynamics of ϕ in phase space are determined by the integral curves of X_ξ:[14]

$$\delta_{X_\xi}\phi^a = \mathcal{L}_{X_\xi}\phi^a = X_\xi(\phi^a) \overset{!}{=} \mathcal{L}_\xi\phi^a = \delta_\xi\phi^a. \tag{3.9}$$

Here we can see the dynamical fields being singled out and flowed. As a consequence, we find that $\mathcal{L}_{X_\xi}\phi^a = \mathcal{L}_\xi\phi^a$ captures the idea that ϕ transforms covariantly under a symmetry operation. This equation is short and looks completely impenetrable: it must be important, and it is imperative that we immediately generalize it beyond recognition. We say that any spacetime tensor $T(\phi,\chi)$ is *covariant* under ξ if $\mathcal{L}_{X_\xi}T = \mathcal{L}_\xi T$. It is quite natural to require the covariance of L and ℓ, whence it follows[15] that θ and C are also covariant.

We are in excellent shape: we have a symplectic form and a Hamiltonian vector field. By symplectic geometry (see §1.3), Hamilton's equations are guaranteed to hold, so $i_{X_\xi}\Omega = \delta H_\xi$ defines the Hamiltonian for ξ-evolution. This is rather useless, however, unless we can actually tell what H_ξ is. Our goal will be to compute $i_{X_\xi}\Omega$ and bumble around with the result until we turn it into $\delta(\text{stuff})$; we will declare victory by setting $(\text{stuff}) = H_\xi$. Reveling in the beauty and the treachery of hindsight, we define the *Noether current* J_ξ by

$$J_\xi = i_{X_\xi}\theta - i_\xi L, \tag{3.10}$$

which is a field-theoretic version of $H = p\dot{q} - L$. We now proceed to compute $i_{X_\xi}\Omega$:

$$\begin{aligned}
i_{X_\xi}\omega &\overset{(1)}{=} i_{X_\xi}\delta(\theta - \mathrm{d}C) \overset{(2)}{=} \mathcal{L}_{X_\xi}\theta - \delta i_{X_\xi}\theta - \mathrm{d}\big(\mathcal{L}_{X_\xi}C - \delta i_{X_\xi}C\big) \overset{(3)}{=} \\
&\overset{(3)}{=} \mathcal{L}_\xi\theta - (\delta J_\xi + i_\xi\delta L) - \mathrm{d}(\text{stuff}) \overset{(4)}{=} i_\xi\mathrm{d}\theta + \mathrm{d}i_\xi\theta - \delta J_\xi - i_\xi\delta L - \mathrm{d}(\text{stuff}) \overset{(5)}{=} \\
&\overset{(5)}{=} \cancel{i_\xi\delta L} + \mathrm{d}i_\xi\theta - \delta J_\xi - \cancel{i_\xi\delta L} - \mathrm{d}\big(\mathcal{L}_{X_\xi}C - \delta i_{X_\xi}C\big) \overset{6}{=} \\
&\overset{6}{=} -\delta J_\xi - \mathrm{d}\big(-i_\xi\theta + \mathcal{L}_{X_\xi}C - \delta i_{X_\xi}C\big) \overset{(7)}{\Longrightarrow} \\
-i_{X_\xi}\Omega &\overset{(7)}{=} \int_\Sigma \delta J_\xi + \int_{\partial\Sigma}\big(-i_\xi\theta + \mathcal{L}_{X_\xi}C - \delta i_{X_\xi}C\big) \overset{(8)}{=} \\
&\overset{(8)}{=} \delta\int_\Sigma J_\xi + \int_{\partial\Sigma}\big(-i_{X_\xi}\theta + i_\xi\mathrm{d}C + \cancel{\mathrm{d}i_\xi C} - \delta i_{X_\xi}C\big) \overset{(9)}{=} \\
&\overset{(9)}{=} \delta\int_\Sigma J_\xi + \int_{\partial\Sigma}\big(i_\xi(\mathrm{d}C - \theta) - \delta i_{X_\xi}C\big) \overset{(10)}{=} \delta\bigg(\int_\Sigma J_\xi + \int_{\partial\Sigma}\big(i_\xi\ell - i_{X_\xi}C\big)\bigg).
\end{aligned} \tag{3.11}$$

Here we did the following: (1) expanded the definition of ω; (2) used Cartan's formula on $\mathcal{P}$ and that $\delta\mathrm{d} = \mathrm{d}\delta$; (3) used the covariance of θ, substituted the definition of J_ξ, and got lazy with the last term; (4) used Cartan's formula on M; (5) used the equations of motion via $\delta L = E\wedge\delta\phi + \mathrm{d}\theta = \mathrm{d}\theta$; (6) did algebra; (7) recalled the definition of Ω and used Stokes's theorem; (8) used the covariance of C on $\partial\Sigma\subset\Gamma$ and Cartan's formula on M; (9) used Stokes's theorem and $\mathrm{d}^2 = 0$; (10) used $(\theta + \delta\ell)\big|_\Gamma = \mathrm{d}C$ and commuted δ past i_ξ.

[14]Forbidding notation aside, this is just a souped up form of the equation for integral curves, $\dot{\gamma} = X(\gamma)$.
[15]This is technical and turns out to almost even be true!

Thus we find that in addition to J_ξ, the Hamiltonian contains a boundary term:

$$H_\xi = \int_\Sigma J_\xi + \int_{\partial\Sigma} \big(i_\xi \ell - i_{X_\xi} C\big). \tag{3.12}$$

Notice that ℓ and C are only evaluated on $\partial\Sigma \subset \Gamma$, so as promised earlier their extensions into M do not affect dynamics. This formula for the symmetry charges of a field theory completes the development of the CPS formalism.

4 Gravity at Last

Having completed the formal development of classical field theory the way it should be done, let us consider an application to the most beautiful field theory of all: Einstein's general theory of relativity. Here we will be brief and omit intermediate calculations, focusing instead on results and their interpretation. We will also focus on the case of vacuum.

4.1 The Action

Pure gravity begins with a smooth pseudo-Riemannian manifold (M, g) with boundary $(\partial M, \gamma)$. The action is the sum of the Einstein-Hilbert and Gibbons-Hawking-York terms and has the metric tensor field g as its dynamical variable:

$$\begin{aligned} S = S_{\text{EH}} + S_{\text{GHY}} &= \int_M L + \int_{\partial M} \ell = \\ &= \frac{1}{16\pi G}\int_M (R - 2\Lambda)\varepsilon_{_M} + \frac{1}{8\pi G}\int_{\partial M} K \varepsilon_{\partial M}. \end{aligned} \tag{4.1}$$

Here $G = G_N$ is Newton's constant, $R = R^\mu_\mu$ is the scalar curvature of (M, g), Λ is the cosmological constant, $K = K^\mu_\mu$ is the trace of the extrinsic curvature, and ε is the volume form, either on M or restricted to ∂M. After a computation, we obtain

$$\delta L = E^{\mu\nu}\delta g_{\mu\nu} + \mathrm{d}\Theta, \tag{4.2}$$

where $E^{\mu\nu}$ are the Einstein field equations

$$E^{\mu\nu} = -\frac{1}{16\pi G}\left(R^{\mu\nu} - \frac{1}{2}R g^{\mu\nu} - \Lambda g^{\mu\nu}\right)\varepsilon_{_M}. \tag{4.3}$$

The precipitation of the field equations from S is certainly encouraging. The boundary term of the variation is given in terms of the "main" part of the symplectic potential:

$$\Theta = i_\theta \varepsilon_{_M}, \qquad \theta^\mu = \frac{1}{16\pi G}\big(g^{\mu\alpha}\nabla^\nu \delta g_{\alpha\nu} - g^{\alpha\beta}\nabla^\mu \delta g_{\alpha\beta}\big). \tag{4.4}$$

A similar computation yields the variation of the boundary Lagrangian:

$$\delta\ell = \frac{1}{16\pi G}\big[(K\gamma^{\mu\nu} - K^{\mu\nu})\delta g_{\mu\nu} + \big(g^{\alpha\beta}n^\lambda\nabla_\lambda - n^\alpha\nabla^\beta\big)\delta g_{\alpha\beta} - D_\mu(\gamma^{\mu\nu}n^\alpha\delta g_{\nu\alpha})\big]\varepsilon_{\partial M}, \tag{4.5}$$

where D_μ is the covariant derivative and n^μ is the unit normal to ∂M. We throw these ingredients into a boiling pot and leave the index soup to simmer; when we return, we find

$$\begin{aligned}(\Theta + \delta\ell)\Big|_\Gamma &= -\frac{1}{16\pi G}(K\gamma^{\mu\nu} - K^{\mu\nu})\delta g_{\mu\nu} + \mathrm{d}C = \\ &= \frac{1}{2}T^{\mu\nu}_{\mathrm{BY}}\delta g_{\mu\nu} + \mathrm{d}C \overset{!}{=} \mathrm{d}C, \\ C = i_c\varepsilon_{\partial M}, &\quad c^\mu = -\frac{1}{16\pi G}\gamma^{\mu\nu}n^\alpha\delta g_{\nu\alpha}.\end{aligned} \tag{4.6}$$

Thus gravity presents us with an example where the subtle term $\mathrm{d}C$ is nonzero. In addition, $\Theta + \delta\ell$ generates an extra term, which we recognize as the so-called *Brown-York boundary stress tensor* $T^{\mu\nu}_{\mathrm{BY}}$, contracted with the variation of the metric. This combination must vanish in order to preserve the well-posedness of the variational problem. Thus we have two options: we can set $\delta g_{\mu\nu}\big|_\Gamma = 0$ (only the pullback to Γ is required, since δg appears here contracted with boundary quantities). Fixing the metric at Γ in this way corresponds to Dirichlet boundary conditions; this choice is common in AdS/CFT. Alternatively, we can set $T^{\mu\nu}_{\mathrm{BY}} = 0$, which is analogous to imposing Neumann boundary conditions.

4.2 The Hamiltonian

With this calculation done, we can proceed past the symplectic form to the Hamiltonian for gravity. We first consider the Noether currnt J_ξ. Wald proved that in general relativity,

$$J_\xi = \mathrm{d}Q_\xi, \qquad Q_\xi = -\frac{1}{16\pi G}\star\mathrm{d}\xi. \tag{4.7}$$

Wald called Q_ξ the "Noether charge," but in fact Q_ξ should have been called the Noether potential, whereas H_ξ is the Noether charge.[16] Here ξ is our diffeomorphism generator, regarded as a one-form,[17] and $\star$ is the Hodge star. Thus by Stokes's theorem we have

$$H_\xi = \int_{\partial\Sigma}\left(-\frac{\star\,\mathrm{d}\xi}{16\pi G} + i_\xi\ell - i_{X_\xi}C\right), \tag{4.8}$$

where ℓ and C are given above.

[16] See a previous footnote complaining about terminology surrounding currents. It would be nice to have a more solid notion of charges, currents, and potentials and their interrelated dualities.

[17] Mathematicians would write $\xi^\flat$, where $\flat$ and $\sharp$ are the *musical isomorphisms* between TM and T^*M. These isomorphisms raise and lower indices, but without the indices.

After some more fiddling around in coordinates, we obtain

$$H_\xi = -\int_{\partial\Sigma} \tau^\mu \xi^\nu T^{\mathrm{BY}}_{\mu\nu} \varepsilon_{\partial\Sigma}, \tag{4.9}$$

where τ is the unit normal to $\partial\Sigma$. (We view $\partial\Sigma$ as the boundary of its past in Γ, whence $\varepsilon_{\partial M} = -\tau \wedge \varepsilon_{\partial\Sigma} \implies \varepsilon_{M} = \tau \wedge n \wedge \varepsilon_{\partial\Sigma}$.) Therefore Hamiltonian evolution in general relativity is generated by a boundary term, determined by ξ and the boundary stress tensor.

The point of these manipulations has been to demonstrate that:

- Calculations are not only possible in the CPS formalism, they are readily accessible. These calculations are "easy" in the sense that they are straightforward coordinate manipulations of differential forms, and they reduce the messy work of over a century of GR into a few pages of algebra at most.
- Although we have not shown this explicitly, this formalism quickly recovers all of the results of the ADM formalism in a manifestly covariant framework. It has also been used by Wald to interpret the entropy of stationary black holes as the Noether charge corresponding to the horizon-generating Killing field.
- Nontrivial physics hides on the boundary, so we have been justly rewarded for our care.

4.3 Gauge Symmetry

Perhaps one point that remains contentious is the issue of gauge symmetry. Here I wish to be extremely pedantic. Upon asking, "what are the gauge symmetries in gravity?" one is usually met with indistinct muttering, curtly told "diffeomorphisms," and then promptly thrown into a firestorm involving first- and second-class constraints, ADM-type language, Noether's first and second theorems, the Bianchi identities, asymptotic symmetries, and so on. It is mind-numbing. Let us quote Appendix B from Carroll's *Spacetime and Geometry*:

> You will often hear it proclaimed that GR is a "diffeomorphism invariant" theory. What this means is that, if the universe is represented by a manifold (M, g) and $\phi\colon M \longrightarrow M$ is a diffeomorphism, then the sets (M, g) and $(M, \phi^* g)$ represent the same physical situation... [these] two purportedly distinct configurations in GR are actually "the same," related by a diffeomorphism.

Carroll goes on to comment that *every* physical theory must be diffeomorphism-invariant in the sense that no coordinate system is preferred to any other; this principle of *general covariance* is universal and confers no physical meaning. It is, however, a principal motivation for CPS methods, which seek to describe field theories without coordinates. Of course this is not what physicists typically mean by the diffeomorphism invariance of gravity.

A *diffeomorphism* is a map of smooth manifolds and does not see Riemannian structure: for instance, the flat unit disk is diffeomorphic to the upper half-sphere, but the former is flat while the latter is curved. It is emphatically *not* true that diffeomorphisms of a spacetime preserve its metric; neither g nor R nor $S_{\text{EH}} + S_{\text{GHY}}$ are invariant under the spacetime group $\text{Diff}(M)$. These objects *are* curvature invariants, but a generic spacetime has no nontrivial isometries and therefore g is unique in the class of metrics to which it is equivalent. A judicious choice of boundary conditions may well produce a variational problem where the (on-shell) action is invariant under some continuous subgroup of $\text{Diff}(M)$, but generically the boundary conditions will produce a metric that prevents this from happening.[18]

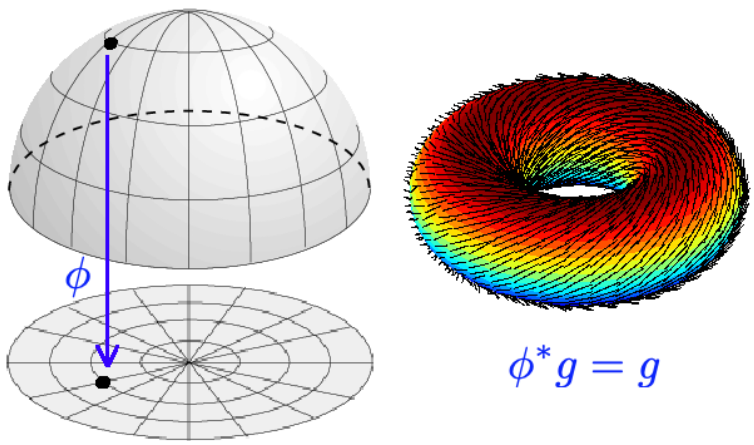

This is why Carroll pulls back the metric in his definition of physical equivalence: if both the points of M and the values of its metric are acted on by ϕ simultaneously, then M flows together with its own geometry, or "rotates into its own shape." Now, any diffeomorphism $\phi\colon (M, g) \longrightarrow (N, g')$ that satisfies $g' = \phi^* g$ is called an *isometry*, so Carroll has noticed that *every diffeomorphism of M can be upgraded to an isometry* by pulling back the metric. Carroll says that (M, g) and $(M, \phi^* g)$ are related by a diffeomorphism, but what he means is that they are related by an isometry that arises from a diffeomorphism. This does not conflict with the statement that (M, g) generically has trivial isometry group, because here ϕ is an isometry between two *distinct* spacetimes (with the same underlying smooth manifold), while the isometry group of a spacetime refers more properly to the *autoisometries* of a single spacetime (M, g), where it is required that $\phi^* g = g$.

The language of CPS sheds some light on this situation. First, recall that gauge symmetry was the reason that $\widetilde{\Omega}$ was degenerate. We may therefore *define* the gauge group of a theory as the group generated by vector fields in the kernel of $\widetilde{\Omega}$. In gravity, these vector fields implement the transformation $g \mapsto \phi^* g$ infinitesimally. These "gauge isometries" are exactly what we removed in moving from $\widetilde{\mathcal{P}}$ to $\mathcal{P}$: a rose by any other name would smell as sweet, a donut gazed upon from another angle would taste as delicious, and Minkowski space viewed in a boosted frame would appear just as flat.

[18] Any field theory with a fixed background metric may have an action that respects spacetime symmetries "a priori," whether or not the equations of motion are satisfied. But general relativity is the dynamical theory of spacetime itself: there is no prior geometry, so the action can have spacetime symmetries only *after* the Einstein field equations have been solved uniquely for the metric using appropriate boundary data.

What, then, should we make of the diffeomorphism charges H_ξ? Naïvely, since "GR is Diff(M)-invariant," there should be a charge H_ξ for every possible diffeomorphism generator ξ. But we know better; most of these ξ are fake, or rather they generate gauge isometries that just look at the spacetime from different angles. Once the metric has been uniquely determined by the field equations and the boundary conditions, we may decide whether (M, g) has any isometries; if it does, the physical charges H_ξ represent the various conserved quantities in the spacetime. It seems surprising that the charges generating the theory's evolution can only be determined after the field evolution is complete and the metric is known. Perhaps this is due to a misguided understanding of the issues raised here.

In any case, we will conclude as we started: with a single field, which molds and shapes the spacetime in which it evolves, with a puzzle, and with a smile.

References

1. **Lee:** Smooth Manifolds.
2. **Landau & Lifshitz:** Mechanics (I), Field Theory (II).
3. **Arnold:** Mathematical Methods of Classical Mechanics.
4. **Marsden:** Foundations of Mechanics.
5. **José and Saletan:** Classical Dynamics.
6. **nLab entries:** field; phase space.
7. **Carroll:** Spacetime and Geometry.
8. **Wald:** General Relativity.
9. **Crnkovic-Witten:** "Covariant description of geometrical theories."
10. **Lee-Wald:** Local Symmetries and Constraints.
11. **Harlow-Wu:** Covariant Phase Space with Boundaries.
12. **Peierls:** "The Commutation Laws of Relativistic Field Theory."