

A geometric perspective on gauging procedures in the Hamiltonian formalism

Sylvain Lavau

*Euler Institute (EIMI) & Steklov Institute, St-Petersburg, Russian Federation.
Aristotle University of Thessaloniki, Thessaloniki, Greece.*

The idea of the course is to provide notions of topology and geometry to mathematical physicists, as well as present concrete applications of such notions to pure algebraists and geometers. The idea is not to fall in a too abstract presentation, and to anchor it into examples taken from physics. I propose to start from the basics and grow in complexity to reach higher grounds which are much more intricate. We first introduce differential calculus on \mathbb{R}^n and then we turn to differential geometry, in order to later understand Poisson geometry. The last part of the course deals with the applications of differential and Poisson geometry to Hamiltonian mechanics and quantization of such classical systems. We cover canonical Hamiltonian formalism under constraints, as well as the BRST formalism. These two last chapters can be understood as a mathematical interpretation and a reading guide to the well-known book of Henneaux and Teitelboim on gauge theories [[Henneaux and Teitelboim, 1992](#)].

I propose to follow a physically informed mathematical path, as I am convinced that physical ideas provide fertile ground for mathematical invention and, on the other hand, that geometry is a very natural and adapted framework shedding light on physical theories. Mathematical physicists often find their inspiration in problems and objects set up by theoretical physicists, on which they draw to develop interesting and useful mathematical objects. The latter may be somewhat ‘generalizations’ of the former, but they need not encode exactly the physics that inspired them. Theoretical physics is a playground for mathematical physicists who use nice and insightful results to develop fruitful mathematical theories. This course will follow the same line of reasoning: drawing on physical examples to present useful mathematical objects.

Contents

1	Differential calculus on \mathbb{R}^n	4
1.1	Tangent vectors and vector fields on \mathbb{R}^n	4
1.2	Cotangent vectors and differential 1-forms on \mathbb{R}^n	15
1.3	Differential forms on \mathbb{R}^n and the de Rham complex	23
1.4	De Rham cohomology and Maxwell equations	32
2	Differential calculus on smooth manifolds	37
2.1	Smooth manifolds	37
2.2	Vector bundles, pushforwards, pullbacks	42
2.3	Submanifolds in differential geometry	54
2.4	Distributions and foliations	62
2.5	Orientation of smooth manifolds and integration of differential forms	70
3	Poisson geometry	78
3.1	Poisson manifolds	78
3.2	Properties of Poisson bivectors	88
3.3	Submanifolds and reduction in Poisson geometry	99
4	A geometric perspective on the canonical Hamiltonian formalism	117
4.1	Lagrangian and Hamiltonian formalism from a geometric point of view	118
4.2	Hamiltonian under constraints	129
4.3	The Bergmann-Dirac algorithm	137
4.4	First-class and second-class constraints, gauge transformations	147
4.5	The geometry of the constraint surface	162
5	BRST formalism and quantization of constrained systems	177
5.1	Canonical quantization of a Hamiltonian system	179
5.2	First principles of the BRST formalism: the irreducible case	184
5.3	Graded geometry and its associates	200
5.4	BRST formalism for reducible constraints	216
5.5	BRST quantization of Hamiltonian systems under constraints	228
A	Mathematical background in linear algebra	237
A.1	The tensor algebra, the symmetric algebra and the exterior algebra	237
A.2	Scalar product and Hodge star operator	244
B	Supplementary and facultative material	251
B.1	Pseudo-Riemannian manifolds and Laplace-de Rham operator	251
B.2	The Poisson-sigma model	254

B.3 The symplectic structure of phase space and canonical transformations	259
B.4 Canonical formalism in field theories	270

References **275**

Main resources: the lecture notes aim at providing a self-consistent introduction to the main results of differential and Poisson geometry, and to their applications in gauge theories. In order to build this course, I mostly relied on various lecture notes and on the following sources:

[Baez and Muniain, 1994] Baez, J. and Muniain, J. P. (1994). *Gauge Fields, Knots and Gravity*, volume 4 of *Series on Knots and Everything*. World Scientific, Singapore

[Henneaux and Teitelboim, 1992] Henneaux, M. and Teitelboim, C. (1992). *Quantization of Gauge Systems*. Princeton University Press, Princeton

[Laurent-Gengoux et al., 2013] Laurent-Gengoux, C., Pichereau, A., and Vanhaecke, P. (2013). *Poisson Structures*. Grundlehren Der Mathematischen Wissenschaften. Springer-Verlag, Berlin, Heidelberg

[Lee, 2003] Lee, J. M. (2003). *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer-Verlag, New York

[Lee, 2009] Lee, J. M. (2009). *Manifolds and Differential Geometry*. Number 107 in Graduate Studies in Mathematics. American Mathematical Society, Providence, RI

[Rothe and Rothe, 2010] Rothe, H. J. and Rothe, K. D. (2010). *Classical and Quantum Dynamics of Constrained Hamiltonian Systems*, volume 81 of *World Scientific Lecture Notes in Physics*. World Scientific, Singapore

Notice that reference [Lee, 2003] is the first (2003) edition. The second (2012) edition has been widely revised and the chapters have been shuffled so that the references to the first edition do not correspond to the same in the second edition.

Notations: we will use Einstein summation convention on sums over space-time coordinates: when an index appears twice (only) in a term, and is such that it appears once as an exponent, and once as a bottom index, then one may get rid of the sum sign, and understand that the sole presence of the repeated indices symbolizes the summation. For exemple, when we write $g_{ij}e^j$ (where g_{ij} symbolizes a metric and e^j is a covector), it mathematically means $\sum_{1 \leq j \leq n} g_{ij}e^j$. We would not use this convention for summation other than space-time coordinates which are such that there is one index up and one index down. The kronecker delta will always have one index up and one index down, as in δ_j^i . We also widely use the *rationalized Planck units*, where:

$$c = 4\pi G = \hbar = \varepsilon_0 = k_B = 1$$

Acknowledgments: I would like to thank Nikita Belousov, Casey Blacker, Leonid Ryvkin, Gérard Lavau and Dimitris Makris for their insightful comments and ideas. Their help has been invaluable to improve the quality of the notes and detect the unavoidable mistakes (there are certainly some surviving in the text).

1 Differential calculus on \mathbb{R}^n

This chapter is dedicated to the study of smooth functions, vector fields and differential forms on the euclidean vector space $E = \mathbb{R}^n$. A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ that is infinitely differentiable is called *smooth* and the set of all such functions is denoted $\mathcal{C}^\infty(\mathbb{R}^n)$. It is an infinite dimensional real vector space, and since the product of two smooth functions is always smooth, it is actually an algebra over \mathbb{R} . Vector fields are derivations of this algebra, while differential 1-forms are their dual objects.

1.1 Tangent vectors and vector fields on \mathbb{R}^n

In this section we generalize the notion of tangent vector to a curve. The idea is the following: assume $n = 3$ and that we represent the trajectory of a physical object in space by a parametrized (differentiable) curve $\gamma : [0, 1] \rightarrow \mathbb{R}^3$. For every $t_0 \in [0, 1]$, the velocity vector $\dot{\gamma}(t_0)$ is often represented as a tangent vector to the curve which has the following properties:

1. it is a 3-dimensional vector based at $\gamma(t_0)$;
2. it is tangent to the curve and points towards the future, i.e. towards the points $\gamma(t_1)$, for small $t_1 > t_0$;
3. its norm is the velocity $\|\dot{\gamma}(t_0)\|$ of the physical object at time t_0 .

The direction and the norm of the tangent vector are somewhat "internal" informations because they can be represented by an abstract vector $X_{\Gamma(t_0)}$ based at the origin of an abstract 3-dimensional space, and which points in the same direction as $\dot{\gamma}(t_0)$ and has the same norm. The base point at which the velocity vector is defined however is an external information since it depends on the curve γ .

Hence, an abstraction of the velocity vector and of the data contained in the three items above can be equivalently represented by a couple $(\gamma(t_0), X_{\gamma(t_0)})$ of the product space $\mathbb{R}^3 \times \mathbb{R}^3$. The first \mathbb{R}^3 is the "position space" (or *configuration space*): it is the space in which the trajectory γ of the physical object takes values. The second \mathbb{R}^3 is the "velocity space": for any given point $x \in \mathbb{R}^3$ (of the position space), The curve $\gamma : t \rightarrow \mathbb{R}^3$ encoding the trajectory of a physical object defines a path in the position space. When t varies in $[0, 1]$, the direction and the norm of the tangent vector of γ varies, which in turn defines a path $X : t \rightarrow X_{\gamma(t)}$ in the velocity space. Thus, the path $t \rightarrow (\gamma(t), X_{\gamma(t)})$ defines a curve in $\mathbb{R}^3 \times \mathbb{R}^3$ which contains every data on the physical position of the object and its velocity.

Another way of looking at tangent vectors is the following: let $t_0 \in]0, 1[$, then the tangent vector $\dot{\gamma}(t_0)$ acts on any smooth function $f \in \mathcal{C}^\infty(\mathbb{R}^3)$ by:

$$\dot{\gamma}(t_0)(f) = \frac{d(f \circ \gamma)}{dt}(t_0) \tag{1.1}$$

In particular, if f is a function locally constant at $\gamma(t_0)$, then $\dot{\gamma}(t_0)(f) = 0$. What is not apparent on this equation is that, although the right hand side involve γ , the left hand side only depends on the velocity vector at the point $\gamma(t_0)$. Any other curve $\eta : [0, 1] \rightarrow \mathbb{R}^3$ such that $\eta(t_0) = \gamma(t_0)$ and such that $\dot{\eta}(t_0) = \dot{\gamma}(t_0)$, satisfies $\frac{d(f \circ \eta)}{dt}(t_0) = \frac{d(f \circ \gamma)}{dt}(t_0)$. This implies that we can forget about the dependency on the curve and look at elements of the tangent space as linear maps sending functions to real numbers: for any given point $x \in \mathbb{R}^3$ and any vector $X_x \in \mathbb{R}^3$, pick up a parametrized curve $\gamma : [0, 1] \rightarrow \mathbb{R}^3$ and $t_0 \in]0, 1[$ such that $\gamma(t_0) = x$, and that $\dot{\gamma}(t_0) = X_x$,




Figure 1: Usually we represent a path and its tangent vectors on the same drawing. The tangent vector $\dot{\gamma}(t)$ is based at the point $\gamma(t)$ but this is not rigorous, mathematically: the norm and the direction of $\dot{\gamma}(t)$ characterizes the tangent vector, and the base point is an external information reminding the reader that the tangent vector is attached to the point $\gamma(t)$.




Figure 2: The figure on the left represents the path γ in the "position space" \mathbb{R}^n , and the figure on the right is a possible representation of the path $X : t \mapsto X_{\gamma(t)}$ of velocity vectors tangent to the curve γ , in the "velocity space" \mathbb{R}^n (to determine the exact form of this path, one has to compute every $\dot{\gamma}(t)$). For each $t \in [0, 1]$, the vector $X_{\gamma(t)}$ has the same norm and the same direction as $\dot{\gamma}(t)$. The path $t \mapsto (\gamma(t), X_{\gamma(t)})$ in the abstract product space $\mathbb{R}^n \times \mathbb{R}^n$ contains the same data which is represented in Figure 1.

then X_x defines a linear morphism $X_x : \mathcal{C}^\infty(\mathbb{R}^3) \longrightarrow \mathbb{R}$ via Equation (1.1). Due to the properties of the time derivative, one can show that this action satisfies the following properties:

$$X_x(\lambda f + \mu g) = \lambda X_x(f) + \mu X_x(g) \tag{1.2}$$

$$X_x(fg) = X_x(f)g(x) + f(x)X_x(g) \tag{1.3}$$

for any $f, g \in \mathcal{C}^\infty(\mathbb{R}^3)$ and $\lambda, \mu \in \mathbb{R}$. The first equation characterizes the fact that $X_x :$

$\mathcal{C}^\infty(\mathbb{R}^3) \longrightarrow \mathbb{R}$ is a linear morphism, whereas the second equation implies that X_x acts as what we call a *derivation at x* . Actually, we will see that the action of the vector X_x on a function f can be identified with the directional derivative of f in the direction X_x , evaluated at the point x (see below).

Generalizing this observation to n -dimensional vector spaces gives the following definition: the *tangent space* to \mathbb{R}^n at a given point x is the vector space of linear morphisms that are derivations at x , i.e. all the maps $X_x : \mathcal{C}^\infty(\mathbb{R}^n) \longrightarrow \mathbb{R}$ satisfying Equations (1.2) and (1.3); it is denoted $T_x\mathbb{R}^n$. The following Lemma says that every directional derivative is a derivation at x :

Lemma 1.1. *Let $v \in \mathbb{R}^n$ and $x \in \mathbb{R}^n$. The linear morphism $D_{v,x} : \mathcal{C}^\infty(\mathbb{R}^n) \longrightarrow \mathbb{R}$ defined by:*

$$D_{v,x}(f) = \left. \frac{d}{dt} \right|_{t=0} f(x + tv)$$

is a derivation at x , i.e. $D_{v,x} \in T_x\mathbb{R}^n$.

Exercise 1.2. Prove this Lemma, using Equation (1.1).

Let e_1, \dots, e_n be the standard basis of \mathbb{R}^n . We denote by $\left. \frac{\partial}{\partial x^i} \right|_x$ the directional derivative at x associated to the basis vector e_i by Lemma 1.1:

$$\left. \frac{\partial}{\partial x^i} \right|_x = D_{e_i,x}$$

The notation is such that the action of $\left. \frac{\partial}{\partial x^i} \right|_x$ precisely coincides with what is expected from such a directional derivative:

$$\left. \frac{\partial}{\partial x^i} \right|_x (f) = \frac{\partial f}{\partial x^i}(x)$$

From the definitions of such elements, we deduce the following Lemma:

Lemma 1.3. *The n directional derivatives $\left. \frac{\partial}{\partial x^1} \right|_x, \dots, \left. \frac{\partial}{\partial x^n} \right|_x$ at the point x are linearly independent, and there is a one-to-one correspondence between vectors v of \mathbb{R}^n and directional derivatives $D_{v,x}$:*

$$v = v^i e_i \quad \longleftrightarrow \quad D_{v,x} = v^i \left. \frac{\partial}{\partial x^i} \right|_x$$

Proof. This result can be shown as follows: pick up a set of scalars $\lambda_1, \dots, \lambda_n$ and assume that $\sum_{i=1}^n \lambda_i \left. \frac{\partial}{\partial x^i} \right|_x = 0$. Thus in particular, applying it to the i -th coordinate function $x^i : \mathbb{R}^n \longrightarrow \mathbb{R}$ gives $\lambda_i = 0$. Also one notices that the assignment $v \longmapsto D_{v,x}$ is a linear morphism. This shows that differential derivative decompose as $v_{v,x} = v^i \left. \frac{\partial}{\partial x^i} \right|_x$. Finally, this result is used to prove that the linear map $v \longmapsto D_{v,x}$ is injective. This proves that there is a one-to-one correspondence between \mathbb{R}^n and the space of directional derivatives at x . \square

The following proposition explains why this is also true for derivations at x :

Proposition 1.4. *The n directional derivatives $\left. \frac{\partial}{\partial x^1} \right|_x, \dots, \left. \frac{\partial}{\partial x^n} \right|_x$ at the point x form a basis of $T_x\mathbb{R}^n$. In particular it means that directional derivatives at x and derivations at x are in one-to-one correspondence, and that $T_x\mathbb{R}^n$ is a n -dimensional vector space.*

Proof. We know by Lemma 1.1 that directional derivatives are derivation at x . This, together with Lemma 1.3, implies that the assignment $v \mapsto D_{v,x}$ is an injection from \mathbb{R}^n into $T_x\mathbb{R}^n$. We need only show that it is surjective. It can be proven by assigning, to each derivation X_x at x , a vector v so that its i -th coordinate coincides with $X(x^i)$: $v = X_x(x^i)e_i$. Then, showing that $D_{v,x} = X$ is just a matter of using Taylor's series expansion (the Hadamard Lemma). For a detailed proof, see Proposition 3.2 in [Lee, 2003]. \square

Thus, any tangent vector X_x at x decomposes in this basis as:

$$X_x = X_x^i \frac{\partial}{\partial x^i} \Big|_x \quad (1.4)$$

where the X_x^i are real numbers, and result from applying X_x to the i -th coordinate function $x^i : \mathbb{R}^n \rightarrow \mathbb{R}$, that is to say:

$$X_x^i = X_x(x^i)$$

The vector $v = X_x^i e_i$ of \mathbb{R}^n which has the same coordinates as X_x then induces a directional derivative $D_{v,x}$ that precisely coincides with X_x :

$$X_x = D_{X_x^i e_i, x}$$

The one-to-one correspondence between directional derivatives at x and derivations at x is summarized by the following sequence of operations:

$$X_x = X_x^i \frac{\partial}{\partial x^i} \Big|_x \quad \longrightarrow \quad v = X_x^i e_i \quad \longrightarrow \quad D_{v,x} = v^i \frac{\partial}{\partial x^i} \Big|_x = X_x$$

Here, $v^i = X_x^i$ by construction. This sequence also describes the canonical isomorphism between \mathbb{R}^n and $T_x\mathbb{R}^n$. We will often identify \mathbb{R}^n with its image under the canonical bijection $v \mapsto D_{v,x}$, and will either use the notation (x, X_x) or the notation X_x for a tangent vector in $T_x\mathbb{R}^n$, depending on how much emphasis we wish to give to the point x .

Example 1.5. In relativity, if x^μ are coordinates of a point particle in space-time, then the four-velocity, often represented by its coordinate $U^\mu = \frac{dx^\mu}{d\tau}$ (where τ is the proper time), is a tangent vector to the world line of the particle.

So far we have used a geometric perspective (tangent vectors to a curve) to determine algebraic properties that they satisfy: Equation (1.2) and (1.3). Then, we have adopted the other perspective: we started from all the linear morphisms from $C^\infty(\mathbb{R}^n)$ to \mathbb{R}^n satisfying these equations, and we have shown that they are directional derivatives. Thus we started from algebraic properties to come back to the geometric realm. We will see in the following that this alternance between geometric and algebraic perspectives are central in the discussion. This one-to-one correspondence actually allows us to transform cumbersome geometric considerations into easier algebraic computations, and conversely, to find clear geometric illustrations of algebraic obscure notions. Let us now generalize the notion of tangent vector, to the whole of \mathbb{R}^n :

Definition 1.6. *The disjoint union of all tangent spaces:*

$$T\mathbb{R}^n = \bigsqcup_{x \in \mathbb{R}^n} T_x\mathbb{R}^n$$

is called the tangent bundle of \mathbb{R}^n .

The word ‘bundle’ means that several things of the same kind have been fastened or held together. We call it a *trivial* (vector) bundle because it is homeomorphic¹ to $\mathbb{R}^n \times \mathbb{R}^n$. In the latter cartesian product, we call the space on the left the *base* and the space on the right the *fiber*. The fiber *at* x is the tangent space $T_x\mathbb{R}^n$, which can then be identified with the product $\{x\} \times \mathbb{R}^n$. An element of the tangent bundle is a couple (x, X_x) , where x is a point in \mathbb{R}^n and X_x is a tangent vector at x . The projection on the first variable:

$$\begin{aligned} \pi : T\mathbb{R}^n &\longrightarrow \mathbb{R}^n \\ (x, X_x) &\longmapsto x \end{aligned}$$

is surjective, and the pre-image of x through π is the tangent space $T_x\mathbb{R}^n$. This defines a short exact sequence:

$$0 \longrightarrow \mathbb{R}^n \simeq T_x\mathbb{R}^n \longrightarrow T\mathbb{R}^n \xrightarrow{\pi} \mathbb{R}^n \longrightarrow 0$$

$\longleftarrow \sigma$

This sequence *splits*, which means that the map π admits *sections*: continuous maps $\sigma : \mathbb{R}^n \longrightarrow T\mathbb{R}^n$ such that $\pi \circ \sigma = \text{id}_{\mathbb{R}^n}$.

Definition 1.7. We call vector fields over \mathbb{R}^n the sections of π :

$$\begin{aligned} X : \mathbb{R}^n &\longrightarrow T\mathbb{R}^n \\ x &\longmapsto (x, X_x) \end{aligned}$$

that are infinitely differentiable (or smooth) in the second variable (see Scholie 1.8). We denote by $\mathfrak{X}(\mathbb{R}^n)$ the \mathbb{R} -vector space of vector fields on \mathbb{R}^n .




Figure 3: On the left hand side, the ‘geometric’ representation of the tangent vectors to a path in \mathbb{R} . On the right hand side, the abstract representation through the tangent bundle of \mathbb{R} : over each point x there is a fiber $T_x\mathbb{R} \simeq \mathbb{R}$, and the vector field, tangent to the path at each point, is symbolized by a section (dashed curve) of the vector bundle. The ‘height’ of the section in the fiber over a given point x is equal to the modulus of the tangent vector to the path at x .

By definition, vector fields consist of the assignment to every point x of a tangent vector at x , denoted X_x , which is, additionally, required to vary smoothly over \mathbb{R}^n . We will now explain

¹We will actually see later that it is actually diffeomorphic (the notion of equivalence in the category of smooth manifolds).

what we mean by that. The tangent bundle $T\mathbb{R}^n$ is *trivial*, i.e. it is homeomorphic to the cartesian product $\mathbb{R}^n \times \mathbb{R}^n$. We already know a basis on the base: the vectors e_1, \dots, e_n ; let us define a basis on the fiber, denoting it by:

$$\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^n} \quad (1.5)$$

This notation is consistent with the notation of the basis vectors of the tangent space $T_x\mathbb{R}^n$ as in Proposition 1.4, because $T_x\mathbb{R}^n \simeq \{x\} \times \mathbb{R}^n$, so that one can make the straightforward identification:

$$\left. \frac{\partial}{\partial x^i} \right|_x \in T_x\mathbb{R}^n \quad \longleftrightarrow \quad \left(x, \frac{\partial}{\partial x^i} \right) \in \{x\} \times \mathbb{R}^n \quad (1.6)$$

Now, given a section X of the tangent bundle, its evaluation at the point x is a tangent vector X_x which can be decomposed on the standard basis of $T_x\mathbb{R}^n$ as in Equation (1.4). Using the one-to-one correspondence (1.6), this gives the following correspondence:

$$X_x \quad \longleftrightarrow \quad \left(x, X_x^i \frac{\partial}{\partial x^i} \right) \quad (1.7)$$

Then, for every $1 \leq i \leq n$, this defines an assignment:

$$\begin{aligned} X^i : \mathbb{R}^n &\longrightarrow \mathbb{R} \\ x &\longmapsto X_x^i \end{aligned}$$

This provides us with the following criterion for smoothness of sections of the tangent bundle:

Scholie 1.8. Smoothness criterion for vector fields *A section $X : \mathbb{R}^n \rightarrow T\mathbb{R}^n$ being smooth means that the applications:*

$$\begin{aligned} X^i : \mathbb{R}^n &\longrightarrow \mathbb{R} \\ x &\longmapsto X_x^i \end{aligned}$$

are smooth functions of x (i.e. they are infinitely differentiable).

It turns out that the role of the basis $\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^n}$ is not only computational, for the sake of the presentation, but has an additional very practical scope. First, the basis elements $\frac{\partial}{\partial x^i}$ can be seen as sections of the tangent bundle, through the following assignment: to every point x , $\frac{\partial}{\partial x^i}$ associates the element $\left. \frac{\partial}{\partial x^i} \right|_x$ in the tangent space $T_x\mathbb{R}^n$. Because the assignment is canonical, we still denote by $\frac{\partial}{\partial x^i}$ these sections, although the reader should remember that, rigorously, they are not the same mathematical objects as the basis vectors (1.5). Second, this set of sections forms a basis of the fiber at each point, by Proposition 1.4.

A set of smooth sections that satisfy these two criteria is called a *frame*. In the present case, the $\frac{\partial}{\partial x^i}$ are in fact *constant* sections, and thus are automatically vector fields on \mathbb{R}^n by Scholie 1.8. They provide a basis for the C^∞ -module of sections, as the following discussion illustrates. Given a vector field X , Scholie 1.8 says that the functions X^i are smooth, so that one can define an additional vector field $X^i \frac{\partial}{\partial x^i}$ on \mathbb{R}^n . By equivalence (1.7), we observe that the vector field X and the vector field $X^i \frac{\partial}{\partial x^i}$ coincide at every point x . Thus, one can identify the two vector fields and write:

$$X = X^i \frac{\partial}{\partial x^i} \quad (1.8)$$

It turns out that every vector field can be uniquely decomposed in such a way. This is why we call the functions X^i the *coordinate functions* of the vector field X .

Example 1.9. Examples of vector fields (every coordinate functions are smooth):

$$\begin{aligned}
 X &= y^2 z \frac{\partial}{\partial x} + x e^y \frac{\partial}{\partial y} + 4 \frac{\partial}{\partial z} \quad \text{in } \mathbb{R}^3 \\
 Y &= 3 y \sin(t) \frac{\partial}{\partial x} + x^3 y^8 z^3 t^9 \frac{\partial}{\partial z} + \arctan(x) \frac{\partial}{\partial t} \quad \text{in } \mathbb{R}^4 \\
 Z &= \begin{cases} 0 & \text{when } x \leq 0 \\ e^{-\frac{1}{x}} \frac{\partial}{\partial x} & \text{when } x > 0 \end{cases} \quad \text{in } \mathbb{R} \\
 E &= x^i \frac{\partial}{\partial x^i} \quad \text{in } \mathbb{R}^n, \text{ is called the } \textit{Euler vector field}
 \end{aligned}$$

Examples of objects which are *not* vector fields:

$$1. \quad e^{-\frac{1}{x}} \frac{\partial}{\partial x}, \quad 2. \quad |x^i| \frac{\partial}{\partial x^i}, \quad 3. \quad \frac{x}{y-1} \frac{\partial}{\partial z}, \quad 4. \quad t^{\frac{1}{3}} \frac{\partial}{\partial t}$$

The first object differs from the vector field Z on the negative semi-axis, and this actually makes a huge difference: although the function $x \mapsto e^{-\frac{1}{x}}$ is smooth on the right of 0 (its limit is zero), it explodes in the left of 0. This function is not smooth at 0, let alone continuous: that is why the object defined in item 1. is not a vector field. In contrast, to avoid this problem, we have imposed on the object Z to vanish for negative values of x , so that it becomes a well-defined vector field. The object in item 2. is not smooth at zero because the absolute value function is not a smooth function (although it is continuous). The third object is not a vector field because if $x \neq 0$ then the function $y \mapsto \frac{x}{y-1}$ explodes at 1. The fourth item is not a vector field because it is not differentiable in 0.

Example 1.10. In quantum mechanics (where space time is \mathbb{R}^4 , say), one can write the wave function ψ in polar form: $\psi = \sqrt{\rho} e^{iS}$, where ρ is a positive smooth function over \mathbb{R}^4 and where S is a real-valued smooth function (over \mathbb{R}^4). The probability density is $\rho = \psi^\dagger \psi$ and the probability current is denoted $\mathbf{j} = \frac{\rho}{m} \nabla S$, where ∇ (nabla) symbolizes the gradient (with respect to spatial coordinates). In coordinate notations this reads: $\mathbf{j} = \sum_{i=1}^3 \frac{\rho}{m} \frac{\partial S}{\partial x^i} \frac{\partial}{\partial x^i}$; the sum is made over spatial dimensions only. Since the function S is defined all over space-time and is supposedly smooth, for each fixed time t , the probability current is a vector field (on the space \mathbb{R}^3). Representing the time with the fourth coordinate x^4 , the 4-vector $\rho \frac{\partial}{\partial x^4} + \mathbf{j}$ defines a vector field over space-time \mathbb{R}^4 . Using the Schrodinger equation, one can show that the divergence of this 4-vector vanishes, which can be interpreted as a *continuity equation for the probability current*.

Example 1.11. Pick up a solution of the heat equation $\frac{\partial u}{\partial t} = \Delta u$ in a homogeneous material of thermal conductivity k (it is a real number then). Assume that this solution is smooth in the four variables (x, y, z, t) , or equivalently (x^1, x^2, x^3, x^4) . Then, for each time t , the heat flow defined as $\mathbf{q} = -k \nabla u = -\sum_{i=1}^3 k \frac{\partial u}{\partial x^i} \frac{\partial}{\partial x^i}$ is a vector field on \mathbb{R}^3 (or at least the part of \mathbb{R}^3 where the material is). It indicates at every point of space in which direction the heat flows.

It is important to notice at this point that, when x varies over M , the direction and the norm of the tangent vector X_x varies (it can even vanish at some point !). Hence one sees that a vector field has no "direction" per se, but it is assigned one direction at each point of \mathbb{R}^n . Tangent vectors at x were directional derivatives at x ; what would be the similar perspective for vector fields?

For every function $f \in \mathcal{C}^\infty(\mathbb{R}^n)$, the assignment $x \mapsto X_x(f)$ defines a function from \mathbb{R}^n to \mathbb{R} . We call this function $X(f)$ and it satisfies, at every point:

$$X(f)(x) = X_x(f) = X_x^i \frac{\partial f}{\partial x^i}(x)$$




Figure 4: Example of a vector field with two points where it vanishes: one from which the vector field ‘flows out’, and one where it ‘flows in’.

In particular, if f is a constant function, $X(f) = 0$. Because X is smooth, the coordinate functions $X^i : x \mapsto X_x^i$ are smooth functions of x , as are the derivatives $\frac{\partial f}{\partial x^i}$. Then, the function $X(f)$ is a smooth function. Then the vector field X can be seen as an endomorphism of the (infinite dimensional) vector space $\mathcal{C}^\infty(\mathbb{R}^n)$, also denoted X :

$$\begin{aligned} X : \mathcal{C}^\infty(\mathbb{R}^n) &\longrightarrow \mathcal{C}^\infty(\mathbb{R}^n) \\ f &\longmapsto X(f) = X^i \frac{\partial f}{\partial x^i} \end{aligned}$$

This is consistent with the remark that the vector field X can be written as $X^i \frac{\partial}{\partial x^i}$ as explained in Equation (1.8). From this discussion, one sees that the vector field X can be seen as a directional derivative in the direction of X or, said differently, along the *integral curves of X* , i.e. those paths γ in \mathbb{R}^n such that X is always tangent to γ : at each time t , $X_{\gamma(t)} = \dot{\gamma}(t)$. A vector field being a family of tangent vectors indexed over the points of \mathbb{R}^n , they inherit the derivation property of tangent vectors Equation (1.3).

A vector field X induces a *derivation* of the algebra of smooth functions $\mathcal{C}^\infty(\mathbb{R}^n)$, i.e. an endomorphism of the (infinite dimensional) vector space $\mathcal{C}^\infty(\mathbb{R}^n)$ that satisfies the following identity:

$$\text{for every } f, g \in \mathcal{C}^\infty(\mathbb{R}^n) \quad X(fg) = X(f) \cdot g + f \cdot X(g) \quad (1.9)$$

where \cdot symbolizes the multiplication of function in $\mathcal{C}^\infty(\mathbb{R}^n)$. While Equation (1.3) was valid pointwise (because we were working with tangent vectors, defined at a point), Equation (1.9) is valid independently of the point. We denote by $\text{Der}(\mathcal{C}^\infty(\mathbb{R}^n))$ the space of all derivations of $\mathcal{C}^\infty(\mathbb{R}^n)$. Conversely, one can show that any derivation of $\mathcal{C}^\infty(\mathbb{R}^n)$ is a vector field, in the sense of Definition 1.7:

Proposition 1.12. *Vector fields on \mathbb{R}^n are in one-to-one correspondence with derivations of $\mathcal{C}^\infty(\mathbb{R}^n)$:*

$$\mathfrak{X}(\mathbb{R}^n) \simeq \text{Der}(\mathcal{C}^\infty(\mathbb{R}^n))$$

Proof. We have shown that every vector field is a derivation, and we just need to show that a derivation is a vector field. Let $\mathcal{X} \in \text{Der}(\mathcal{C}^\infty(\mathbb{R}^n))$ and define a section of the tangent bundle $X : \mathbb{R}^n \rightarrow T\mathbb{R}^n$ by:

$$X_x(f) = \mathcal{X}(f)(x)$$

This equation makes sense because $\mathcal{X}(f)$ is a smooth function, and the right hand side is its evaluation in x . The object X_x is then a derivation at x , i.e. an element of $T_x\mathbb{R}^n$. One needs only to prove that the assignment $x \mapsto X_x$ is smooth. This is shown for example in Proposition 4.7 in [Lee, 2003] and in Proposition 8.15 in the 2012 edition. \square

This proposition is important because it shows that there is a correspondence between the geometric perspective (vector fields on \mathbb{R}^n) and the algebraic perspective (derivations of $\mathcal{C}^\infty(\mathbb{R}^n)$). We have said that passing from one point of view to the other allows to make sense or make things easier. Let us illustrate this strategy by showing that the algebraic perspective is adapted to define a Lie bracket on the space of vector fields $\mathfrak{X}(\mathbb{R}^n)$:

Definition 1.13. A Lie algebra is a (real, possibly infinite dimensional) vector space \mathfrak{g} , equipped with a bilinear operation $[\cdot, \cdot]$ called the Lie bracket, which satisfies the following identities:

$$\begin{array}{ll} \text{skew-symmetry} & [x, y] = -[y, x] \\ \text{Jacobi identity} & [x, [y, z]] = [[x, y], z] + [y, [x, z]] \end{array}$$

for every $x, y, z \in \mathfrak{g}$

Remark 1.14. The Jacobi identity is often written under the following, equivalent but more symmetric, form:

$$[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0$$

The form presented in Definition 1.13 is useful because it makes clear that "the Lie bracket is a derivation of itself". Here, by *derivation of \mathfrak{g}* we mean any endomorphism $\delta : \mathfrak{g} \rightarrow \mathfrak{g}$ such that:

$$\delta([x, y]) = [\delta(x), y] + [x, \delta(y)]$$

Then, notice that to every element x of a Lie algebra $(\mathfrak{g}, [\cdot, \cdot])$, we can associate a derivation of \mathfrak{g} via the adjoint action of x on \mathfrak{g} :

$$\begin{array}{ccc} \text{ad} : \mathfrak{g} & \longrightarrow & \text{Der}(\mathfrak{g}) \\ x & \longmapsto & \text{ad}_x : y \longmapsto [x, y] \end{array}$$

The Jacobi identity ensures that ad_x is a derivation of \mathfrak{g} , for every x . The image $\text{ad}(\mathfrak{g}) \subset \text{Der}(\mathfrak{g})$ forms what is called the space of *inner derivation of \mathfrak{g}* , sometimes denoted $\text{inn}(\mathfrak{g})$.

Example 1.15. One can always equip an associative algebra (A, \cdot) with a Lie algebra structure, by setting:

$$[a, b] = a \cdot b - b \cdot a$$

for every $a, b \in A$. In particular, the space of $n \times n$ matrices $\mathcal{M}_n(\mathbb{R})$ (equivalently, the space $\text{End}(E)$ of a n -dimensional vector space) is an associative algebra, so that we can define a Lie bracket on it.

Vector fields are derivations of the (infinite-dimensional) space $\mathcal{C}^\infty(\mathbb{R}^n)$. Derivations are special cases of endomorphisms. However, the composition of two vector fields is *not* a derivation of $\mathcal{C}^\infty(\mathbb{R}^n)$, as the following computation shows:

$$X(Y(fg)) = X(Y(f)g + fY(g)) = X(Y(f))g + fX(Y(g)) + (X(f)Y(g) + X(g)Y(f))$$

The latter parenthesis prevents the composite $X \circ Y$ to be a derivation of $\mathcal{C}^\infty(\mathbb{R}^n)$. Hence, the space of derivations of $\mathcal{C}^\infty(\mathbb{R}^n)$ is not stable under composition. However, inspired by Example 1.15, let us define the following operation on the space of vector fields:

$$[X, Y](f) = X(Y(f)) - Y(X(f)) \tag{1.10}$$

The right-hand side is a smooth function so the left-hand side is a smooth function as well, but one needs to show that the bracket $[X, Y]$ is still a derivation of $\mathcal{C}^\infty(\mathbb{R}^n)$, that is to say: the space $\mathfrak{X}(\mathbb{R}^n) \simeq \text{Der}(\mathcal{C}^\infty(\mathbb{R}^n))$ is stable under the action of this bracket. The proof of the following proposition is left as an exercise:

Proposition 1.16. *The \mathbb{R} -vector space $\mathfrak{X}(\mathbb{R}^n)$ equipped with this operation is a Lie algebra.*

Exercise 1.17. Prove that Equation (1.10) defines a Lie bracket, i.e. that it is bilinear (with respect to real numbers), skew-symmetric, and that it satisfies the Jacobi identity. By expanding the Lie bracket, prove that the Lie bracket of two vector fields is still a vector field (i.e. a derivation of $\mathcal{C}^\infty(\mathbb{R}^n)$).




Figure 5: Metaphorical picture of the bracket of vector fields. If one follows for small times first the integral curve of Y , then the integral curve of X , one arrives at the point x_1 . Whereas, if one had followed the integral curve of X first, and then that of Y , one arrives at the point x_2 . The Lie bracket $[X, Y]$ is the vector field whose integral curve links x_1 to x_2 . This discussion can be made rigorous if one makes the time of walking along integral curves tend to 0. See for example page 47 of [Baez and Muniain, 1994], where however the last equation is wrong: the sign should be the opposite, as consequently should be the vector field $[X, Y]$ on figure 10.

A subtle remark has to be made here. We have seen that every vector field can be decomposed on the basis of vectors $\frac{\partial}{\partial x^i}$. However, this basis does *not* form a basis of the \mathbb{R} -vector space $\mathfrak{X}(\mathbb{R}^n)$, which is actually infinite dimensional as a real vector space. More precisely:

Scholie 1.18. Algebraic characterization of $\mathfrak{X}(\mathbb{R}^n)$ *The space $\mathfrak{X}(\mathbb{R}^n)$ is, at the same time:*

1. *an infinite dimensional \mathbb{R} -vector space;*
2. *a $C^\infty(\mathbb{R}^n)$ -module of finite rank.*

The notion of module over a ring is the generalization of the notion of a vector space over a field. Given a ring (R, \circ) , we say that a vector space E is a R -module² if there is an action \cdot of R on E which satisfies the following axioms:

$$\begin{aligned} r \cdot (x + y) &= r \cdot x + r \cdot y & (r + s) \cdot x &= r \cdot x + s \cdot x \\ r \cdot (s \cdot x) &= (r \circ s) \cdot x & 1_R \cdot x &= x \end{aligned}$$

where $r, s \in R$, $x, y \in E$, and where 1_R is the identity of the ring. The reader can check that these axioms are the same axioms that the scalars have to satisfy when acting on a vector space. In our case, the field is \mathbb{R} and the ring is $C^\infty(\mathbb{R}^n)$. Then, when we say that $\mathfrak{X}(\mathbb{R}^n)$ is a \mathbb{R} -vector space we understand that vector fields can be added, and multiplication by real scalars is well-defined. When we say that $\mathfrak{X}(\mathbb{R}^n)$ is a $C^\infty(\mathbb{R}^n)$ -module, we mean that vector fields can be added, and that multiplication by *smooth functions* is well-defined. Notice that, since constant functions can be identified with real scalars, the fact that $\mathfrak{X}(\mathbb{R}^n)$ is a \mathbb{R} -vector space is a consequence of the fact that it is a $C^\infty(\mathbb{R}^n)$ -module.

Now, the dimension of a vector space is the minimal number of independent vectors that generate the space (using only multiplication by real scalars and addition). The rank of a module is the maximum number of elements which are linearly independent under the action of the ring. In our case, every vector field X decomposes on the elements $\frac{\partial}{\partial x^i}$ as $X = X^i \frac{\partial}{\partial x^i}$, where the X^i are *smooth functions* (we see the module structure emerging). Moreover, those constant sections are linearly independent over $C^\infty(\mathbb{R}^n)$ because by definition the identity $X^i \frac{\partial}{\partial x^i} = 0$ implies $X^i = 0$. Thus, $\mathfrak{X}(\mathbb{R}^n)$ is a $C^\infty(\mathbb{R}^n)$ -module of rank n . What is crucial in the present discussion is that the generators of a module need not coincide with a basis of the underlying vector space, because *the multiplication with a ring element generate much different elements than the multiplication with a scalar*. Indeed, one can explicitly compute how the ring of smooth functions $C^\infty(\mathbb{R}^n)$ acts on a vector field X via multiplication: let $f \in C^\infty(\mathbb{R}^n)$, then we define the vector field fX to be the unique vector field whose coordinate functions are fX^i , where here we understand the product of two functions. In other words:

$$fX = (f \cdot X^i) \frac{\partial}{\partial x^i}$$

so that $(fX)^i = f \cdot X^i$, where \cdot symbolizes the multiplication of function in $C^\infty(\mathbb{R}^n)$. Pointwise, this vector field satisfies $(fX)_x = f(x)X_x$. We see how the structure of $C^\infty(\mathbb{R}^n)$ -module only needs n -generators to be defined.

However, $\mathfrak{X}(\mathbb{R}^n)$ is an infinite dimensional vector space. This can be shown by contradiction. Assume there exists a finite number of vector fields X_1, \dots, X_r which form a basis of $\mathfrak{X}(\mathbb{R}^n)$ (as a real vector space), that is: every vector field X would be uniquely written as $X = \sum_{s=1}^r \lambda_s X_s$, where the λ_s are real numbers. Then, given a smooth function f , there exists real scalars μ_1, \dots, μ_r such that $fX = \sum_{s=1}^r \mu_s X_s$. On the other hand, multiplying $\sum_{s=1}^r \lambda_s X_s$ by f gives $fX = \sum_{s=1}^r f \lambda_s X_s$. By unicity of the decomposition, $f \lambda_s = \mu_s$ for every $s = 1, \dots, r$, which is impossible most of the time because f need not be constant. The demonstration may be a bit too much abstract. The idea of the proof is that a finite number of elements cannot form a set

²Actually, in full generality we only require E to be an abelian group. A vector space is an abelian group with respect to the addition.

of generators for all the vector fields, because multiplication by any function offers much more freedom and variability that can be encoded by a mere finite dimensional vector space.

The Lie algebra structure on $\mathfrak{X}(\mathbb{R}^n)$ is defined on top of the vector space structure. Thus, Scholie 1.18 explains why $\mathfrak{X}(\mathbb{R}^n)$ is a real Lie algebra of infinite dimension, although only a finite number of constant sections $\frac{\partial}{\partial x^i}$ is needed to generate all the vector fields (using the ring multiplication). This additionally explains why the Lie bracket is bilinear with respect to the scalars, but not with respect to the smooth functions. More precisely, since every vector field can be decomposed on the frame $\frac{\partial}{\partial x^i}$, a small computation shows that the Lie bracket of X and Y reads:

$$[X, Y] = \left(X(Y^i) - Y(X^i) \right) \frac{\partial}{\partial x^i} \quad (1.11)$$

where we recall that X^i and Y^i are the i -th coordinate functions associated to X and Y , respectively. The Einstein summation convention has been used. Then, for any smooth function $f \in \mathcal{C}^\infty(\mathbb{R}^n)$:

$$[X, fY] = f[X, Y] + X(f)Y \quad (1.12)$$

where the term on the right hand side has to be understood as the multiplication of the function $X(f)$ with the vector field Y . Equation (1.12) shows that the Lie bracket defined in Equation (1.10) is not linear with respect to the functions, as expected since it should only be linear with respect to real numbers.

Remark 1.19. We conclude this section by introducing an alternative notation for the constant vector fields $\frac{\partial}{\partial x^i}$, that may also be denoted ∂_i :

$$\partial_i \equiv \frac{\partial}{\partial x^i}$$

The position of the index i is indeed at the bottom because one should formally consider the fractional notation $\frac{\partial}{\partial x^i}$ as a fraction of fractions: $\frac{\frac{a}{b}}{\frac{c}{d}}$, where the index i is at the top of the denominator, at the place occupied by the element c . Since the latter fraction can be written as $\frac{ad}{bc}$, and that the element c is at the bottom, this justifies that we place the index i at the bottom of the notation ∂_i . I emphasize that keeping track of the position of indices is central in differential geometry when we work in coordinates. Moreover the above informal reasoning will have some relevance later in the text. Using Equation (1.11), we deduce that the commutator of two vectors of the frame vanishes:

$$[\partial_i, \partial_j] = 0 \quad (1.13)$$

1.2 Cotangent vectors and differential 1-forms on \mathbb{R}^n

Now let us turn to the elements dual to tangent vectors and vector fields. Given some point $x \in \mathbb{R}^n$ we call the *cotangent space* and we write $T_x^*\mathbb{R}^n$ the dual of the tangent space at x :

$$T_x^*\mathbb{R}^n = (T_x\mathbb{R}^n)^*$$

Elements of this dual space are called *cotangent vectors* at x . They are linear forms on $T_x\mathbb{R}^n$ and there is a canonical bijection between $(\mathbb{R}^n)^*$ and $T_x^*\mathbb{R}^n$: since a basis of $T_x\mathbb{R}^n$ is given by the vectors $\partial_i|_x = \frac{\partial}{\partial x^i}|_x$, using Equation (A.11) one obtains a dual basis of $T_x^*\mathbb{R}^n$ whose elements are denoted $dx^i|_x$. In particular one has:

$$dx^i|_x(\partial_j|_x) = \delta_j^i \quad (1.14)$$

Thus, for any tangent vector X_x , one has $dx^i|_x(X_x) = X_x^i$. In particular, one observes that it is as if the tangent vector X_x had been fed with the coordinate function $x^i : \mathbb{R}^n \mapsto \mathbb{R}$ that associates a point to its i -th coordinate:

$$dx^i|_x(X_x) = X_x(x^i)$$

One can now extend the notion of cotangent vectors to the one of covector fields, following what has been said in Section 1.1.

Definition 1.20. We define the cotangent bundle $T^*\mathbb{R}^n$ to be the union of all cotangent spaces:

$$T^*\mathbb{R}^n = \bigsqcup_{x \in \mathbb{R}^n} T_x^*\mathbb{R}^n$$

It has several properties: it is a trivial vector bundle over \mathbb{R}^n , i.e. it is diffeomorphic to $\mathbb{R}^n \times (\mathbb{R}^n)^*$. As expected from the definition of cotangent spaces, the *fiber* or $T_x^*\mathbb{R}^n$ is $(\mathbb{R}^n)^*$, the dual space of the fiber of $T\mathbb{R}^n$. Points in $T^*\mathbb{R}^n$ are couples (x, ξ_x) , where ξ_x is a notation for cotangent vectors at x ; they decompose on the basis of $T_x^*\mathbb{R}^n$ as $\xi_{x,i}dx^i|_x$, where the $\xi_{x,i}$ are the coordinates of ξ_x (they are real numbers). The projection on the first variable, denoted $\tau : T^*\mathbb{R}^n \rightarrow \mathbb{R}^n$, admits sections:

Definition 1.21. We call covector fields – or differential 1-forms – over \mathbb{R}^n the sections of τ :

$$\begin{aligned} \xi : \mathbb{R}^n &\longrightarrow T^*\mathbb{R}^n \\ x &\longmapsto (x, \xi_x) \end{aligned}$$

that are infinitely differentiable (or smooth) in the second variable (see Scholie 1.22). We denote by $\Omega^1(\mathbb{R}^n)$ the \mathbb{R} -vector space of covector fields/differential 1-forms on \mathbb{R}^n .

Since the cotangent bundle is trivial (i.e. it is diffeomorphic to a cartesian product), one can define a standard basis on its fiber. The fiber \mathbb{R}^n of the tangent bundle is already equipped with a standard basis: the generators $\partial_i = \frac{\partial}{\partial x^i}$. The dual basis would form a basis of the fiber $(\mathbb{R}^n)^*$ of the cotangent bundle; let us denote this basis by:

$$dx^1, \dots, dx^n$$

and we call it the *dual coframe* to the given frame. This notation is consistent with the notation of the basis vectors of the cotangent spaces. Indeed, since $T_x^*\mathbb{R}^n \simeq \{x\} \times (\mathbb{R}^n)^*$, we can make the following identification:

$$dx^i|_x \in T_x^*\mathbb{R}^n \quad \longleftrightarrow \quad (x, dx^i) \in \{x\} \times (\mathbb{R}^n)^* \quad (1.15)$$

Thus, the basis vectors dx^1, \dots, dx^n can also be seen as *constant* sections of the cotangent bundle: the differential 1-forms dx^i associates, to every point x , the cotangent vector $dx^i|_x$ via the above correspondence. Every cotangent vector ξ_x defined at the point x can be decomposed on the dual basis defined in Equation (1.14) as $\xi_x = \xi_{x,i}dx^i|_x$. Then, because of the injection of $T_x^*\mathbb{R}^n$ into $T^*\mathbb{R}^n$, the one-to-one equivalence defined in Equation (1.15) defines an equivalence:

$$\xi_x \quad \longleftrightarrow \quad (x, \xi_{x,i}dx^i)$$

where Einstein summation convention has been used. Then, for every $1 \leq i \leq n$, this defines an assignment:

$$\begin{aligned} \xi_i : \mathbb{R}^n &\longrightarrow \mathbb{R} \\ x &\longmapsto \xi_{i,x} \end{aligned}$$

A priori the coordinate functions ξ_i of a random section ξ are not smooth, unless the section is smooth, i.e. unless it is a covector field. As for vector fields, this actually provides a first criterion for smoothness of covector fields (see Scholie 1.22).

Furthermore, this enables us to understand how sections of $T^*\mathbb{R}^n$ act on vector fields. Recall that we have the following identity, by definition of the dual basis on the fiber of $T^*\mathbb{R}^n$:

$$dx^i(\partial_j) = \delta_j^i$$

By construction, the constant sections dx^i are $\mathcal{C}^\infty(\mathbb{R}^n)$ -linear: $dx^i(X^j\partial_j) = X^j dx^i(\partial_j) = X^i$. Then, given a section ξ of $T^*\mathbb{R}^n$ and a vector field X , one has:

$$\xi(X) = \xi_i X^j dx^i(\partial_j) = \xi_i \cdot X^i \quad (1.16)$$

where the Einstein summation convention has been used, and where \cdot symbolizes the multiplication of function in $\mathcal{C}^\infty(\mathbb{R}^n)$. The term on the right of Equation (1.16) is a product of functions, thus the term on the left is a function as well. Evaluating both terms in a point x gives:

$$\xi(X)(x) = \xi_{x,i} X_x^i = \xi_x(X_x)$$

where the term in the middle is a sum of products of real numbers. A priori the function $\xi(X) : x \mapsto \xi_x(X_x)$ is not smooth, unless ξ is a smooth section of $T^*\mathbb{R}^n$, i.e. unless it is a covector field. This observation provides the second criterion for smoothness of covector fields:

Scholie 1.22. Smoothness criteria for covector fields *A section $\xi : \mathbb{R}^n \longrightarrow T^*\mathbb{R}^n$ being smooth means:*

1. *that the components functions ξ_i :*

$$\begin{aligned} \xi_i : \mathbb{R}^n &\longrightarrow \mathbb{R} \\ x &\longmapsto \xi_{x,i} \end{aligned}$$

are smooth functions of x (i.e. they are infinitely differentiable);

2. *or that, equivalently³, for every vector field X , the function:*

$$\begin{aligned} \xi(X) : \mathbb{R}^n &\longrightarrow \mathbb{R} \\ x &\longmapsto \xi_x(X_x) \end{aligned}$$

is smooth.

Example 1.23. An example of a covector field in \mathbb{R}^2 :

$$\xi = (2xy \cos(x) - x^2y \sin(x))dx + x^2 \cos(x)dy$$

We will see in Section 1.3 that such a differential 1-form is actually the differential of the function $f(x, y) = x^2y \cos(x)$.

Example 1.24. A physically oriented example consists of the connections A_μ . Actually they correspond to a differential 1-form (taking values in a Lie algebra) $A = A_\mu dx^\mu$.

³The equivalence of the two criteria is shown by noticing that: 1. the second one is implied by the first one using Equation (1.16) and smoothness of vector fields, and 2. the first one is implied by the second one if one picks up $X = \partial_i$ for every $1 \leq i \leq n$.

Hence there are at least two way at looking at covector fields (= differential 1-forms): one is to see them as smooth sections of the cotangent bundle, and in that case they are smooth if and only if item 1. of Scholie 1.22 is satisfied. Another way of looking at covector fields is to see them as being linear morphisms on the space of vector fields, landing in the smooth functions, that is to say:

$$\Omega^1(\mathbb{R}^n) \simeq \text{Hom}(\mathfrak{X}(\mathbb{R}^n), \mathcal{C}^\infty(\mathbb{R}^n)) \quad (1.17)$$

The homomorphisms here have to be understood as homomorphisms of $\mathcal{C}^\infty(\mathbb{R}^n)$ -modules. More precisely, in that case, a covector field ξ can be seen as a linear morphism:

$$\begin{aligned} \xi : \mathfrak{X}(\mathbb{R}^n) &\longrightarrow \mathcal{C}^\infty(\mathbb{R}^n) \\ X &\longmapsto \xi(X) \end{aligned}$$

The fact that this map lands in the smooth functions for every choice of vector field is precisely the content of item 2. of Scholie 1.22. Although the latter perspective is often the most used, the former one is useful to have a glimpse of the geometrical meaning of differential 1-forms.

What is the meaning of covector fields/differential 1-forms? An explanation can be the following: a differential 1-form $\xi : \mathbb{R}^n \longrightarrow T^*\mathbb{R}^n$ defines, at every point x , a linear form $\xi_x : T_x\mathbb{R}^n \longrightarrow \mathbb{R}$ on the tangent space at x . As a map from a n -dimensional space to a 1-dimensional space, the kernel of this linear form is an hyperplane H_x of $T_x\mathbb{R}^n$, that is: a $n - 1$ -dimensional subspace. This hyperplane separates the n -dimensional space $T_x\mathbb{R}^n$ in two (n -dimensional) open half-spaces. The linear form additionally defines a ‘positive’ half-space H_x^+ and a ‘negative’ half-space H_x^- : the former consists of all tangent vectors X_x such that $\xi_x(X_x) > 0$, while the latter consists of all tangent vectors X_x such that $\xi_x(X_x) < 0$. The hyperplane H_x is a separator between these two half-spaces since $\xi_x|_{H_x} = 0$. Since the covector ξ_x is a linear morphism from $T_x\mathbb{R}^n$ to \mathbb{R} , its *level sets* are $(n - 1)$ -dimensional affine subspaces defined as follows:

$$H_{x,t} = \{X_x \mid \xi_x(X_x) = t\}$$

for every $t \in \mathbb{R}$. The notation is consistent with the definition of H_x because $H_{x,0} = H_x$. In particular, the positive half-space and the negative half-space are union of level sets:

$$H_x^+ = \bigcup_{t>0} H_{x,t} \quad \text{and} \quad H_x^- = \bigcup_{t<0} H_{x,t}$$

The level sets define a partition of $T_x\mathbb{R}^n$ by parallel affine subspaces. The main point here is that a linear form is entirely described from its level sets. Smoothly varying the linear form ξ_x then has the consequence of smoothly changing its level sets and in particular: their inclination and their respective distance. Thus a differential 1-form can be seen as a smooth assignment, to every point x , of a partition of $T_x\mathbb{R}^n$ by parallel affine subspace. Smoothness of this assignment means that the partition (of the fiber \mathbb{R}^n) smoothly varies when the base point varies.

These hyperplanes have a geometric significance: let $\xi^\#$ be the vector field corresponding to ξ through the musical isomorphism $\#$ (where we assume the metric to be the euclidean metric on \mathbb{R}^n). Then the hyperplane $H_x = \text{Ker}(\xi_x)$ defines the tangent space to the transversal to the integral curve of $\xi^\#$. In other words, the tangent vector $\xi_x^\#$ is orthogonal to H_x . The distance between the hyperplanes is an alternative – though equivalent – measure of the length of $\xi^\#$: the hyperplanes are closer to one another at the points where $\xi^\#$ has a small modulus, and they are more distant to one another at the points where $\xi^\#$ has a bigger modulus.

Before moving to the next section, we would go for a quick excursion through the realm of *vector bundles* (over \mathbb{R}^n). The idea of a vector bundle is the following: given a k -dimensional vector space, \mathbb{R}^k say, we attach a copy of such a vector space at each and every point of the




Figure 6: Visual representation of the geometrical meaning of differential forms. When the base point varies smoothly, the partition of the tangent space by $(n - 1)$ -dimensional affine subspaces varies smoothly: the inclination and the relative distance of the affine hyperspaces is smoothly modified. The hyperspace H_x (resp. H_y) is orthogonal to the tangent vector $\xi_x^\#$ (resp. $\xi_y^\#$), and can be seen as the tangent space to the transversal to the integral curve of $\xi^\#$ at x (resp. y).

space \mathbb{R}^n . This form an enormous space denoted E for example, that we require to be sufficiently well-defined (to be clear: it should be a topological space, i.e. a space along with a topology of open sets). The topology on E is chosen so that at least locally, in the neighborhood of any point, say U , E looks like $U \times \mathbb{R}^k$. Trivial vector bundles are precisely those that have this structure globally, i.e. those of the form $\mathbb{R}^n \times \mathbb{R}^k$. It turns out that every vector bundle defined over \mathbb{R}^n has this property. The precise statement is the following:

Definition 1.25. A (trivial) vector bundle of rank k (over \mathbb{R}^n) is a topological space E together with a surjective continuous map $\pi : E \rightarrow \mathbb{R}^n$, satisfying the two following conditions:

1. for every $x \in \mathbb{R}^n$, the preimage $\pi^{-1}(x) \subset E$ is a k -dimensional vector space, called the fiber of E at x and denoted E_x ;
2. there exists a homeomorphism $\Phi : E \rightarrow \mathbb{R}^n \times \mathbb{R}^k$ (called a trivialization of E), making the following triangle commutative:



Figure 7: Wherever it makes sense, the kernel of the differential $d\xi$ defines the tangent space to the transversal to the vector field $\xi^\#$.

$$\begin{array}{ccc}
 E & \xrightarrow{\Phi} & \mathbb{R}^n \times \mathbb{R}^k \\
 \pi \searrow & & \swarrow pr_1 \\
 & \mathbb{R}^n &
 \end{array}$$

where $pr_1 : \mathbb{R}^n \times \mathbb{R}^k \rightarrow \mathbb{R}^n$ is the projection on the first variable, and such that for every $y \in \mathbb{R}^n$, the restriction of Φ to E_y is a linear isomorphism from E_y to $\{y\} \times \mathbb{R}^k \simeq \mathbb{R}^k$.

Remark 1.26. The second item means that for every $u \in E$, one has the following identity:

$$\pi(u) = pr_1 \circ \Phi(u)$$

Notice that, in full generality (i.e. on a smooth manifold), the second item should hold only locally (see e.g. Chapter 5 of [Lee, 2003], Chapter 10 in the 2012 edition). The fact that \mathbb{R}^n is contractible implies that every vector bundle is trivial, and that we wrote this second item from the global perspective.

One should really think of a vector bundle as a bunch of vector spaces stacked together and labeled by points. The set underlying any vector bundle is the disjoint union of its fiber:

$$E = \bigsqcup_{x \in \mathbb{R}^n} E_x$$

There is a natural topology on a disjoint union, that we call the ‘disjoint union topology’: it is the finest topology that makes the injective functions $\phi_x : E_x \hookrightarrow E$ continuous. More precisely, with respect to this topology, $U \subset E$ is open if and only if $\phi_x^{-1}(U)$ is open in E_x for every $x \in \mathbb{R}^n$. Assuming that every E_x is homeomorphic to \mathbb{R}^k with its standard topology, the disjoint union $E = \bigsqcup_{x \in \mathbb{R}^n} E_x$ equipped with its disjoint union topology is then homeomorphic to the product of topological space $\mathbb{R}^n \times \mathbb{R}^k$, where \mathbb{R}^n has the *discrete topology* and the product has the product topology. Hence, the disjoint union underlying every vector bundle over \mathbb{R}^n is homeomorphic to $\mathbb{R}^n \times \mathbb{R}^k$, with respect to topologies that we do not like though (because we are not interested in working on \mathbb{R}^n with the discrete topology). What additional property does a vector bundle have then, that the mere underlying disjoint union does not have? The answer is that it is equipped with a ‘vector bundle topology’ – certainly coarser than the disjoint union topology – such that there is an homeomorphism between $E = \bigsqcup_{x \in \mathbb{R}^n} E_x$ equipped with its vector bundle topology and $\mathbb{R}^n \times \mathbb{R}^k$, but here \mathbb{R}^n has its standard topology (which is not discrete!). This is why a vector bundle is much more than its underlying set, the disjoint union of all its fibers.

A *section* of a vector bundle E (over \mathbb{R}^n) is a continuous map $\sigma : \mathbb{R}^n \rightarrow E$ satisfying the following identity:

$$\pi \circ \sigma = \text{id}_{\mathbb{R}^n}$$

In other words, $\sigma(x) \in E_x$ for every x . A section can be symbolically represented as a n -dimensional surface in E , that is projectable onto \mathbb{R}^n . Given a section σ , if the map $\Phi \circ \sigma$ is smooth we call σ a *smooth function*. The space of smooth sections of E then consists of smooth functions from \mathbb{R}^n to E and is denoted $\Gamma(E)$ (sometimes, also denoted $\Gamma(\mathbb{R}^n, E)$ or $\mathcal{C}^\infty(\mathbb{R}^n, E)$). As was explained in Scholie 1.18, these spaces are real vector spaces of infinite dimension, and $\mathcal{C}^\infty(\mathbb{R}^n)$ -modules of finite rank, k to be precise, as is shown by the following paragraph.

Assume that we have k smooth sections $\sigma_1, \dots, \sigma_k$ that are fiberwise linearly independent, i.e. for every x , the vectors $\sigma_1(x), \dots, \sigma_k(x)$ form a basis of E_x . Then, we call such a family a *frame* for E . Since every vector bundle over \mathbb{R}^n is trivial, one can pickup constant orthonormal frames, i.e. for every $1 \leq i \leq k$, the smooth map $\Phi \circ \sigma_i : \mathbb{R}^n \rightarrow \mathbb{R}^n \times \mathbb{R}^k$ is constant, and thus defines a vector $f_i \in \mathbb{R}^k$, so that f_1, \dots, f_k forms a basis of \mathbb{R}^k . Under an intelligent choice of sections, this basis can be made orthonormal. A frame forms a set of generator of the sections of E , with respect to the $\mathcal{C}^\infty(\mathbb{R}^n)$ -module structure on $\Gamma(E)$. That is why the rank of this module coincides with the number of vector in the frame, which is the same as the dimension of the fiber: k .

Famous examples of vector bundles are the tangent bundle $T\mathbb{R}^n$ (with fiber $T_x\mathbb{R}^n \simeq \mathbb{R}^n$) and the cotangent bundle $T^*\mathbb{R}^n$ (with fiber $T_x^*\mathbb{R}^n \simeq (\mathbb{R}^n)^*$). Smooth sections of $T\mathbb{R}^n$ are vector fields and smooth sections of $T^*\mathbb{R}^n$ are differential 1-forms:

$$\mathfrak{X}(\mathbb{R}^n) = \Gamma(T\mathbb{R}^n) \quad \text{and} \quad \Omega^1(\mathbb{R}^n) = \Gamma(T^*\mathbb{R}^n)$$

A frame for $T\mathbb{R}^n$ is the family of constant vector fields $\partial_1, \dots, \partial_n$, whereas a frame for $T^*\mathbb{R}^n$ (what we had called a *coframe*) is made of the constant covector fields dx^1, \dots, dx^n . Moreover, drawing on the material presented in Section A.1, one can construct the following other vector bundles:

$$\wedge^m T\mathbb{R}^n = \bigsqcup_{x \in \mathbb{R}^n} \wedge^m T_x\mathbb{R}^n \quad \text{and} \quad \wedge^m T^*\mathbb{R}^n = \bigsqcup_{x \in \mathbb{R}^n} \wedge^m T_x^*\mathbb{R}^n$$

The notation is transparent: the fiber at a given point x is the m -th exterior power of $T_x\mathbb{R}^n$ (or $T_x^*\mathbb{R}^n$, respectively). These are trivial vector bundles (as is every vector bundle over \mathbb{R}^n). Smooth sections of $\wedge^m T\mathbb{R}^n$ are called m -multivector fields, whereas smooth sections of $\wedge^m T^*\mathbb{R}^n$ are called differential m -forms, and are denoted as follows:

$$\mathfrak{X}^m(\mathbb{R}^n) = \Gamma(\wedge^m T\mathbb{R}^n) \quad \text{and} \quad \Omega^m(\mathbb{R}^n) = \Gamma(\wedge^m T^*\mathbb{R}^n)$$

In particular $\mathfrak{X}^0(\mathbb{R}^n) = \Omega^0(\mathbb{R}^n) = \mathcal{C}^\infty(\mathbb{R}^n) \simeq \mathfrak{X}^n(\mathbb{R}^n) \simeq \Omega^n(\mathbb{R}^n)$, and $\mathfrak{X}^1(\mathbb{R}^n) = \mathfrak{X}(\mathbb{R}^n)$. When $m \geq 1$, a frame for $\wedge^m T\mathbb{R}^n$ consists of the constant sections $\partial_{i_1} \wedge \dots \wedge \partial_{i_m}$, whereas a frame for $\wedge^m T^*\mathbb{R}^n$ is given by constant sections of the form $dx^{i_1} \wedge \dots \wedge dx^{i_m}$, for $1 \leq i_1 < \dots < i_m \leq n$.

Let us now find criteria for smoothness of sections of $\wedge^\bullet T^*\mathbb{R}^n$. A (not necessarily smooth but at least continuous) section η of $\wedge^m T^*\mathbb{R}^n$ decomposes on this basis as:

$$\eta = \sum_{1 \leq i_1 < \dots < i_m \leq n} \bar{\eta}_{i_1 \dots i_m} dx^{i_1} \wedge \dots \wedge dx^{i_m} \quad (1.18)$$

We denote the coordinate functions in the basis $dx^{i_1} \wedge \dots \wedge dx^{i_m}$ where we assume that $1 \leq i_1 < \dots < i_m \leq n$ as $\bar{\eta}_{i_1 \dots i_m}$. However, usually the Einstein summation convention (in which the indices i_k vary from 1 to n and are not ordered) is much more practical. To use it, one needs to do a bit of gymnastics. First define the following functions:

$$\text{for every } 1 \leq i_1 < \dots < i_m \leq n \quad \eta_{i_1 \dots i_m} = \frac{1}{m!} \bar{\eta}_{i_1 \dots i_m}$$

Then, for every choice of non-ordered indices $i_1, \dots, i_m \in \{1, \dots, n\}$, there is a unique permutation $\sigma \in S_m$ such that $i_{\sigma(1)} < i_{\sigma(2)} < \dots < i_{\sigma(m)}$. In other words, the permutation σ rearrange the indices so that they come in order. For such a permutation, we define the function $\eta_{i_1 \dots i_m}$ as follows:

$$\eta_{i_1 \dots i_m} = (-1)^\sigma \eta_{i_{\sigma(1)} \dots i_{\sigma(m)}} = \frac{(-1)^\sigma}{m!} \bar{\eta}_{i_{\sigma(1)} \dots i_{\sigma(m)}}$$

Then one can write under Einstein summation convention:

$$\eta = \eta_{i_1 \dots i_m} dx^{i_1} \wedge \dots \wedge dx^{i_m} \quad (1.19)$$

Exercise 1.27. By using the antisymmetry of the wedge product, prove that Equation (1.19) gives back Equation (1.18).

The section η is at least continuous so the functions $\eta_{i_1 \dots i_m}$ are continuous functions on \mathbb{R}^n and, as for covector fields (see Scholie 1.22), they are smooth if and only if η is a smooth section, i.e. if and only if $\eta \in \Omega^m(\mathbb{R}^n)$ (η is a differential m -form). Another criterion for smoothness of η is obtained by using Equation (A.17); when fed with m vector fields, η gives the following continuous function:

$$\begin{aligned} \eta(X_1, \dots, X_m) &= \eta_{i_1 \dots i_m} dx^{i_1} \wedge \dots \wedge dx^{i_m}(X_1, \dots, X_m) \\ &= \eta_{i_1 \dots i_m} \det \begin{pmatrix} X_1^{i_1} & X_2^{i_1} & \dots & X_{m-1}^{i_1} & X_m^{i_1} \\ X_1^{i_2} & X_2^{i_2} & \dots & X_{m-1}^{i_2} & X_m^{i_2} \\ \dots & \dots & \dots & \dots & \dots \\ X_1^{i_{m-1}} & X_2^{i_{m-1}} & \dots & X_{m-1}^{i_{m-1}} & X_m^{i_{m-1}} \\ X_1^{i_m} & X_2^{i_m} & \dots & X_{m-1}^{i_m} & X_m^{i_m} \end{pmatrix} \end{aligned} \quad (1.20)$$

Since the X_i are vector fields, Scholie 1.8 implies that their coordinate function are infinitely differentiable, which implies that the above determinant, as a product of smooth functions of x , is a smooth function over \mathbb{R}^n . Then, it implies that $\eta(X_1, \dots, X_m)$ is a smooth function if and only if the coordinate functions $\eta_{i_1 \dots i_m}$, i.e. if and only if η is a differential m -form. The situation can be summarized as follows:

Scholie 1.28. Smoothness criteria for differential m -forms A section $\eta : \mathbb{R}^n \longrightarrow \wedge^m T^*\mathbb{R}^n$ being smooth means:

1. that the coordinate functions $\eta_{i_1 \dots i_m}$ are smooth functions of x ;
2. or that, equivalently, for every vector fields X_1, \dots, X_m , the continuous function $\eta(X_1, \dots, X_m)$ defined in Equation (1.20) is smooth.

Exercise 1.29. For any given choice of m indices $j_1, \dots, j_m \in \{1, \dots, n\}$, show that Equation (1.20) applied to $\partial_{j_1}, \dots, \partial_{j_m}$ gives:

$$\eta(\partial_{j_1}, \dots, \partial_{j_m}) = m! \eta_{j_1 \dots j_m} = (-1)^\sigma \bar{\eta}_{j_{\sigma(1)} \dots j_{\sigma(m)}}$$

where σ is the unique permutation of m elements such that $j_{\sigma(1)} < j_{\sigma(2)} < \dots < j_{\sigma(m)}$.

The properties of the wedge product on the exterior algebra $\wedge^\bullet T^*\mathbb{R}^n$ are transported to the differential forms. So in particular $dx^i \wedge dx^j = -dx^j \wedge dx^i$, and for any $\eta \in \Omega^k(\mathbb{R}^n)$ and $\mu \in \Omega^l(\mathbb{R}^n)$, the object $\eta \wedge \mu$ is a differential $k+l$ -form, and:

$$\eta \wedge \mu = (-1)^{kl} \mu \wedge \eta \quad (1.21)$$

This turns the graded vector space $\Omega^\bullet(\mathbb{R}^n)$ into a (*graded*) *commutative graded algebra*.

1.3 Differential forms on \mathbb{R}^n and the de Rham complex

Scholie 1.22 shows us that an obvious family of covector fields would be those induced by smooth functions. For every $f \in C^\infty(\mathbb{R}^n)$, let us define the covector field formally denoted df , by the following identity:

$$df(X) = X(f) \quad (1.22)$$

The left hand side is a smooth function by item 2. of Scholie 1.22, as is the right hand side. We call the covector field $df : \mathfrak{X}(\mathbb{R}^n) \longrightarrow C^\infty(\mathbb{R}^n)$ the *differential of the function f* . Not every covector field is the differential of a function. For example, there is no smooth function $f : \mathbb{R}^2 \longrightarrow \mathbb{R}$ such that the globally defined covector field $\xi = xdy - ydx$ would be the differential of⁴. Those covector fields that are of the form df for some smooth function f , and thus satisfy Equation (1.22), are called *exact differential 1-forms*. Recall that the basis vectors of the fibre of the cotangent bundle are denoted dx^i ; this is not a coincidence, because dx^i is the differential of the coordinate function $x^i : \mathbb{R}^n \longmapsto \mathbb{R}$, and its action on a vector field gives:

$$dx^i(X) = X(x^i) = X^i$$

which is the i -th coordinate function of X .

Exercise 1.30. Show that the covector field defined in Example 1.23 is actually an exact differential 1-form by finding a function f from which it is the differential of.

Let us now compute the coordinates of df for some given f , by applying Equation (1.22) to every generator ∂_i :

$$df(\partial_i) = \frac{\partial f}{\partial x^i}$$

Hence, the covector field df decomposes as follows in the dual coframe:

$$df = \frac{\partial f}{\partial x^i} dx^i$$

In other words, the coordinate functions of df coincide with the components of the gradient of f . This is not a coincidence, because we have the following result:

⁴In polar coordinates (r, θ) , this covector field reads $r^2 d\theta$, from which we understand that it cannot be written under the form df .

Proposition 1.31. *Given a smooth function $f \in C^\infty(\mathbb{R}^n)$, there exists a unique vector field on \mathbb{R}^n , denoted $\overrightarrow{\text{grad}}(f)$, such that:*

$$df(X) = g(\overrightarrow{\text{grad}}(f), X) \quad \text{for every } X \in \mathfrak{X}(\mathbb{R}^n)$$

where g is the standard euclidean metric on the fiber of $T\mathbb{R}^n$.

Since the tangent bundle is *trivial*, it is diffeomorphic to the cartesian product $\mathbb{R}^n \times \mathbb{R}^n$. The metric g appearing in the statement of the proposition is the Euclidean metric defined on the fiber. Thus, on the basis vectors $\partial_1, \dots, \partial_n$, it satisfies $g(\partial_i, \partial_j) = 1$ if $i = j$ and 0 otherwise. Although it is not apparent in the proposition, the metric does not depend on the base point. The metric is bilinear so, for $X = X^i \partial_i$ and $Y = Y^j \partial_j$ two vector fields on \mathbb{R}^n , one has:

$$g(X, Y) = g(X^i \partial_i, Y^j \partial_j) = X^i Y^j g(\partial_i, \partial_j) = \sum_{i=1}^n X^i Y^i$$

Notice that we did not use the Einstein summation convention in the rightmost term because the two indices are both exponentiated. It can alternatively be written under this convention as $X^i Y_i$, given that we lowered the second index via the formula $Y_i = g_{ij} Y^j$.

Remark 1.32. Proposition 1.31 is a particular case of a much more general result that states that a pseudo-Riemannian metric on a manifold M defines an isomorphism between TM and T^*M .

Let us now turn to the question of ‘dualizing’ the Lie bracket, so that we obtain an operator on $T^*\mathbb{R}^n$ that encodes it. Let us first rewrite Equation (1.10) using exact differential 1-forms:

$$df([X, Y]) = X(df(Y)) - Y(df(X)) \quad (1.23)$$

for every $X, Y \in \mathfrak{X}(\mathbb{R}^n)$. Although this equation is satisfied for exact covector fields, it does not mean that it is satisfied for all covector fields:

$$\xi([X, Y]) \stackrel{?}{=} X(\xi(Y)) - Y(\xi(X)) \quad (1.24)$$

We would like to measure ‘how far’ a given vector field ξ is from satisfying Equation (1.24). This can be done by passing the term on the left-hand side to the right-hand side, so that we can evaluate the difference between $X(\xi(Y)) - Y(\xi(X))$ and $\xi([X, Y])$. To this end, we set (formal notation) $d\xi$ to be the obstruction of a covector field ξ to satisfy Equation (1.24):

$$d\xi(X, Y) = X(\xi(Y)) - Y(\xi(X)) - \xi([X, Y]) \quad (1.25)$$

A covector field satisfies Equation (1.24) if and only if $d\xi = 0$, when evaluated on any two vector fields. We call such covector fields *closed differential 1-forms*. In particular, exact forms are closed.

Notice that since the right-hand side of Equation (1.25) is a smooth function, the object on the left-hand side formally noted $d\xi(X, Y)$ is a smooth function as well. Then, since $d\xi(X, Y) = -d\xi(Y, X)$, $d\xi$ defines a skew-symmetric operator that, when fed with two vector fields X and Y , gives a smooth function $d\xi(X, Y)$ whose evaluation at the point x reads:

$$\begin{aligned} d\xi : \mathfrak{X}(\mathbb{R}^n) \times \mathfrak{X}(\mathbb{R}^n) &\longrightarrow C^\infty(\mathbb{R}^n) \\ (X, Y) &\longmapsto d\xi(X, Y) : x \longmapsto X_x(\xi(Y)) - Y_x(\xi(X)) - \xi_x([X, Y]_x) \end{aligned}$$

This is consistent with the definitions of the objects so far, because e.g. $\xi(Y)$ is a smooth function, on which the derivation at x , $X_x : C^\infty(\mathbb{R}^n) \longrightarrow \mathbb{R}$, acts, hence the term $X_x(\xi(Y))$ is a real number. Although the Lie bracket of two vector fields is not $C^\infty(\mathbb{R}^n)$ bilinear, one can check that the map $d\xi$ is.

Exercise 1.33. Prove that $d\xi$ is $C^\infty(\mathbb{R}^n)$ bilinear, i.e. that $d\xi(fX + gY, Z) = f \cdot d\xi(X, Z) + g \cdot d\xi(Y, Z)$ for every $f, g \in C^\infty(\mathbb{R}^n)$ and $X, Y, Z \in \mathfrak{X}(\mathbb{R}^n)$, and vice versa with respect to the second variable.

Then, it is sufficient to know how $d\xi$ acts on the couples of basis vectors (∂_i, ∂_j) to know how it acts on any couple of vector fields. Using Equation (1.13), Equation (1.25) becomes:

$$d\xi(\partial_i, \partial_j) = \frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \quad (1.26)$$

The fact that $d\xi(\partial_i, \partial_i) = 0$ is consistent with the fact that $d\xi$ is a skew-symmetric operator. The observation made in Equation (1.26) induces the following result:

Proposition 1.34. *Given a differential 1-form ξ , the skew-symmetric operator $d\xi$ can be seen as a section of the vector bundle $\wedge^2 T^* \mathbb{R}^n$, and reads:*

$$d\xi = \frac{1}{2} \left(\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \right) dx^i \wedge dx^j \quad (1.27)$$

where the Einstein summation convention (on the two indices i and j !) has been used.

Proof. When one applies the right-hand side of this formula to two vector fields X and Y , one obtains:

$$\begin{aligned} \frac{1}{2} \left(\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \right) dx^i \wedge dx^j(X, Y) &= \frac{1}{2} \left(\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \right) (dx^i \otimes dx^j - dx^j \otimes dx^i)(X, Y) \\ &= \frac{1}{2} \left(\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \right) (X^i \cdot Y^j - X^j \cdot Y^i) \\ &= \left(X^i \frac{\partial \xi_j}{\partial x^i} \right) \cdot Y^j - \left(Y^i \frac{\partial \xi_j}{\partial x^i} \right) \cdot X^j \\ &= X(\xi_j \cdot Y^j) - Y(\xi_j \cdot X^j) - \xi_j \cdot X(Y^j) + \xi_j \cdot Y(X^j) \\ &= X(\xi(Y)) - Y(\xi(X)) - \xi([X, Y]) \end{aligned}$$

Where the symbol \cdot has been used to symbolize and emphasize the product of two smooth functions. Since indices which are summed over can be relabelled at one's convenience, we have done this between the second line and the third line. The two supplementary terms added on the right in the fourth line compensate the addition of the terms $\xi_j \cdot X(Y^j) - \xi_j \cdot Y(X^j)$ which were necessary to form the terms $X(\xi_j \cdot Y^j) - Y(\xi_j \cdot X^j)$. Passing from the fourth line to the fifth and last line used Equation (1.11). \square

Remark 1.35. The right hand side of Equation (1.27) contains redundant terms, since:

$$\left(\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \right) dx^i \wedge dx^j = - \left(\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \right) dx^j \wedge dx^i = \left(\frac{\partial \xi_i}{\partial x^j} - \frac{\partial \xi_j}{\partial x^i} \right) dx^j \wedge dx^i$$

The factor $\frac{1}{2}$ precisely compensates such redundancy, so that (1.27) can be rewritten:

$$d\xi = \sum_{1 \leq i < j \leq n} \left(\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j} \right) dx^i \wedge dx^j \quad (1.28)$$

Since the bivectors $dx^i \wedge dx^j$ for $i < j$ form a family of generators of $\wedge^2 T^* \mathbb{R}^n$, the coordinates functions of $d\xi$ in this basis are the $\frac{\partial \xi_j}{\partial x^i} - \frac{\partial \xi_i}{\partial x^j}$, and *not* one-half of it.

Exercise 1.36. Using Equation (A.19), check that applying Equation (1.27) or (1.28) to the couple (∂_i, ∂_j) (beware of the range of the sums!) gives back Equation (1.26).

From Proposition 1.34 we deduce the very important (always true) observation:

Corollary 1.37. *Exact differential 1-forms are closed.*

Proof. We have already seen a proof of such a result by comparing Equations (1.23) and (1.25), but let us use here a more computational approach. Let ξ be an exact differential 1-form. Then there exists f a smooth function on \mathbb{R}^n such that $\xi = df$. In particular it means that $\xi_i = \partial_i f$. Then:

$$d\xi = \frac{1}{2} \left(\frac{\partial^2 f}{\partial x^i \partial x^j} - \frac{\partial^2 f}{\partial x^j \partial x^i} \right) dx^i \wedge dx^j = 0$$

Thus, ξ is a closed form. □

Remark 1.38. We will see later that in the three-dimensional space \mathbb{R}^3 , Corollary 1.37 is equivalent to the following identity:

$$\overrightarrow{\text{curl}}(\overrightarrow{\text{grad}}(f)) = 0$$

Now that we have an explicit formula for the operator $d\xi$, one may ask: which closed differential 1-forms are also exact? That is to say: which covector fields ξ satisfying Equation (1.24) are actually the differential of a function f , i.e. are such that $\xi = df$? Drawing on Proposition 1.31, this question has an equivalent interpretation in terms of vector fields: which vector field X on \mathbb{R}^n such that $\overrightarrow{\text{curl}}(X) = 0$ (whatever that means in dimension higher than 3) can be written as the gradient of a function f ? Indeed, the standard euclidean metric g on the fiber of the tangent bundle defines an isomorphism \tilde{g} between the fiber of $T\mathbb{R}^n$ and $T^*\mathbb{R}^n$ (see Section A.2). The following Lemma is a particular case of Poincaré's Lemma:

Lemma 1.39. (Part of) Poincaré Lemma *Every closed differential 1-form defined on \mathbb{R}^n (for $n \geq 2$) is an exact form. That is to say, for every $\xi \in \Omega^1(\mathbb{R}^n)$ such that $d\xi = 0$, there exists a smooth function $f \in C^\infty(\mathbb{R}^n)$ such that $\xi = df$.*

Proof. Let $\xi = \xi_i dx^i \in \Omega^1(\mathbb{R}^n)$, then define the following function:

$$f(x) = \int_0^1 x^i \xi_i(tx) dt$$

This function is smooth because the ξ_i are smooth functions by Scholie 1.22. Differentiating f at a given x with respect to the k -th variable, and seeing the function $x \mapsto \xi_i(tx)$ as the composite function $x \mapsto tx \mapsto \xi_i(tx)$, one obtains:

$$\begin{aligned} \partial_k f(x) &= \int_0^1 \partial_k(x^i \xi_i(tx)) dt \\ &= \int_0^1 \delta_k^i \xi_i(tx) dt + \int_0^1 x^i \partial_k(\xi_i(tx)) dt \\ &= \int_0^1 \xi_k(tx) dt + \int_0^1 x^i t \partial_k \xi_i(tx) dt \\ &= \int_0^1 \xi_k(tx) dt + \int_0^1 tx^i \partial_i \xi_k(tx) dt \\ &= \int_0^1 \frac{d}{dt}(t \xi_k(tx)) dt \\ &= 1 \cdot \xi_k(x) - 0 \cdot \xi_k(0) \end{aligned}$$

Here we have used the convention that $\partial_k(\xi_i(tx))$ is the derivative in the k -th variable of the function $x \mapsto \xi_i(tx)$ evaluated at x , whereas $\partial_k \xi_i(tx)$ is the derivative of the function ξ , evaluated at tx . This explains why a factor t appears on the third line. To pass to the fourth line

we used Equation (1.26), whose left-hand side is zero because ξ is closed. Thus, we obtain that $\partial_k f = \xi_k$, so that $df = \xi$. \square

Remark 1.40. Actually, Poincaré's Lemma is more general: it applies to every differential p -forms, and does not necessarily assume that they are defined globally but only on star-shaped open subsets of \mathbb{R}^n .

Remark 1.41. When $\mathbb{R}^n = \mathbb{R}^3$, using the one-to-one correspondence between the fiber of the tangent space and the fiber of the cotangent space, Lemma 1.39 is equivalent to saying that every irrotational vector field X (i.e. such that $\overrightarrow{\text{curl}}(X) = 0$) is conservative (i.e. it is the gradient of a function f).

Let us recall what we have so far: we have shown that for every smooth function, there is a differential 1-form df satisfying Equation (1.22). We have additionally shown that for every differential 1-form ξ , there is a differential 2-form $d\xi$ satisfying Equation (1.25). Additionally, by Corollary 1.37, every exact differential 1-form is closed, and by Lemma 1.39, every closed form is exact. Recalling that $\mathcal{C}^\infty(\mathbb{R}^n) = \Omega^0(\mathbb{R}^n)$, we can summarize the situation by the following sequence of spaces:

$$\Omega^0(\mathbb{R}^n) \xrightarrow{d} \Omega^1(\mathbb{R}^n) \xrightarrow{d} \Omega^2(\mathbb{R}^n) \quad (1.29)$$

Given Equations (1.23) and (1.25), the map d can be understood as the *dual* of the Lie bracket: whereas the Lie bracket is a bilinear map from $\mathfrak{X}^2(\mathbb{R}^n)$ to $\mathfrak{X}^1(\mathbb{R}^n)$, the map $d : \xi \mapsto d\xi$ is a linear map from $\Omega^1(\mathbb{R}^n)$ to $\Omega^2(\mathbb{R}^n)$.

We have also seen that there is a strong relationship between the map d and the gradient of a function and the curl of a vector field. For example, we have seen that the identity $\overrightarrow{\text{curl}}(\overrightarrow{\text{grad}}(f)) = 0$ is a reformulation of Corollary 1.37. How does the divergence of a vector field enters in the picture? The same question arises for the Laplacian of a function. We are tempted to extend the sequence (1.29) to the right to account for those operators. This is will be the topic of the rest of this subsection. We need first to introduce a few abstract material:

Definition 1.42. A chain complex (of vector spaces) is a graded vector space $E = (E_i)_{i \in \mathbb{Z}}$ equipped with a family of linear morphisms $d = (d_i : E_i \rightarrow E_{i+1})_{i \in \mathbb{Z}}$:

$$\dots \xrightarrow{d_{-3}} E_{-2} \xrightarrow{d_{-2}} E_{-1} \xrightarrow{d_{-1}} E_0 \xrightarrow{d_0} E_1 \xrightarrow{d_1} E_2 \xrightarrow{d_2} \dots$$

such that $d_{i+1} \circ d_i = 0$. We call the linear operator d the differential of the chain complex.

Remark 1.43. In general we do not bother writing all the indices on the maps d_i and we write d instead, being understood that $d|_{E_i} = d_i$. In that case $d_{i+1} \circ d_i = 0$ becomes:

$$d^2 = 0$$

Moreover, the graded vector space may be graded above or below, or may be only positively/negatively graded, etc.

Let us now show how the sequence (1.29) can be extended to the right:

$$\Omega^0(\mathbb{R}^n) \xrightarrow{d} \Omega^1(\mathbb{R}^n) \xrightarrow{d} \Omega^2(\mathbb{R}^n) \xrightarrow{d} \dots \xrightarrow{d} \Omega^{n-1}(\mathbb{R}^n) \xrightarrow{d} \Omega^n(\mathbb{R}^n) \xrightarrow{d} 0 \quad (1.30)$$

where all vector spaces of degree higher than n , on the right, are understood to be null vector spaces. Recall that each $\Omega^m(\mathbb{R}^n)$ is the space of smooth sections of the vector bundle $\bigwedge^m T^*\mathbb{R}^n$. It admits as a set of $\mathcal{C}^\infty(\mathbb{R}^n)$ -linearly independent generators the elements:

$$\left\{ dx^{i_1} \wedge \dots \wedge dx^{i_m} \mid 1 \leq i_1 < i_2 < \dots < i_m \leq n \right\}$$

Drawing on what has been said in the discussion following Scholie 1.22 – in particular Equality (1.17) – one can equivalently see $\Omega^m(\mathbb{R}^n)$ as the space of alternating m $\mathcal{C}^\infty(\mathbb{R}^n)$ -multi-linear forms on $\mathfrak{X}(\mathbb{R}^n)$ taking values in the smooth functions:

$$\Omega^m(\mathbb{R}^n) \simeq \text{Hom}(\mathfrak{X}^m(\mathbb{R}^n), \mathcal{C}^\infty(\mathbb{R}^n))$$

The homomorphisms here have to be understood as homomorphisms of $\mathcal{C}^\infty(\mathbb{R}^n)$ -modules. It means that for any given smooth section $\eta \in \Omega^m(\mathbb{R}^n)$ and any family of vectorfields X_1, \dots, X_m , the element $\eta(X_1, \dots, X_m)$ is a smooth function. This is a smoothness criterion for differential p -forms.

We have defined the linear morphism $d_0 : \Omega^0(\mathbb{R}^n) \mapsto \Omega^1(\mathbb{R}^n)$ in Equation (1.22), and we have defined the linear morphism $d_1 : \Omega^1(\mathbb{R}^n) \mapsto \Omega^2(\mathbb{R}^n)$ in Equation (1.25). In both case we have written df or $d\xi$ but it should be rigorously understood as d_0f and $d_1\xi$ if one wants to establish a differential whose notation is consistent with Definition 1.42. In the following we will write d instead of d_m because the latter notation is too cumbersome. Generalizing Equation (1.25) to any number of vector field $m \geq 1$, let us define the linear map $d : \Omega^m(\mathbb{R}^n) \mapsto \Omega^{m+1}(\mathbb{R}^n)$ (should be understood as d_m then) from its action on any section $\eta \in \Omega^m(\mathbb{R}^n)$:

$$\begin{aligned} d\eta(X_1, \dots, X_m, X_{m+1}) &= \sum_{i=1}^{m+1} (-1)^{i-1} X_i(\eta(X_1, \dots, \widehat{X}_i, \dots, X_{m+1})) \\ &+ \sum_{1 \leq i < j \leq m+1} (-1)^{i+j} \eta([X_i, X_j], X_1, \dots, \widehat{X}_i, \dots, \widehat{X}_j, \dots, X_{m+1}) \end{aligned} \quad (1.31)$$

where the notation $(X_1, \dots, \widehat{X}_i, \dots, X_{m+1})$ means that the vector field X_i has been removed from the list of vector fields. In other words, for $2 \leq i \leq m$:

$$(X_1, \dots, \widehat{X}_i, \dots, X_{m+1}) = (X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_{m+1})$$

whereas for $i = 0$ we obtain (X_2, \dots, X_{m+1}) and for $i = m + 1$ we obtain (X_1, \dots, X_m) . In a similar fashion we have:

$$([X_i, X_j], X_1, \dots, \widehat{X}_i, \dots, \widehat{X}_j, \dots, X_{m+1}) = ([X_i, X_j], X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_{j-1}, X_{j+1}, \dots, X_{m+1})$$

with similar exceptions for $i = 0$ and $j = m + 1$. First notice that both terms on the right are smooth functions: $\eta(X_1, \dots, \widehat{X}_i, \dots, X_{m+1})$ is a smooth function, on which the vector field X_i acts; and one can check that there are only m vector fields in the term $\eta([X_i, X_j], X_1, \dots, \widehat{X}_i, \dots, \widehat{X}_j, \dots, X_{m+1})$, making it a smooth function too. Thus, the right hand side is infinitely differentiable, which make the left-hand side infinitely differentiable.

Exercise 1.44. Check that Equation (1.31) gives back Equation (1.25) when $m = 1$.

Let us now give a formula for $d\eta$ in the basis of generators dx^i . To do this, evaluate Equation (1.31) on m given constant sections taken out of $\partial_1, \dots, \partial_n$, so that the last term involving the Lie bracket vanishes by Equation (1.13).

Proposition 1.45. *The action of the operator $d : \Omega^\bullet(\mathbb{R}^n) \longrightarrow \Omega^{\bullet+1}(\mathbb{R}^n)$ on a differential m -form $\eta = \eta_{i_1 \dots i_m} dx^{i_1} \wedge \dots \wedge dx^{i_m}$ is given in local coordinates by:*

$$d\eta = \partial_i(\eta_{i_1 \dots i_m}) dx^i \wedge dx^{i_1} \wedge \dots \wedge dx^{i_m} \quad (1.32)$$

where Einstein summation convention is assumed.

Proof. Write $\eta = \eta_{i_1 \dots i_m} dx^{i_1} \wedge \dots \wedge dx^{i_m}$, where $\eta_{i_1 \dots i_m}$ is a smooth function by Scholie 1.28 and where Einstein summation convention on contracted indices is assumed. Let $\partial_{j_1}, \dots, \partial_{j_{m+1}}$ play the role of X_1, \dots, X_{m+1} , then Equation (1.31) gives:

$$\begin{aligned}
d\eta(\partial_{j_1}, \dots, \partial_{j_{m+1}}) &= \sum_{k=1}^{m+1} (-1)^{k-1} \partial_{j_k} (\eta(\partial_{j_1}, \dots, \widehat{\partial_{j_k}}, \dots, \partial_{j_{m+1}})) \\
&\quad + \sum_{1 \leq k < l \leq m+1} (-1)^{k+l} \eta(\underbrace{[\partial_{j_k}, \partial_{j_l}]}_{=0}, \partial_{j_1}, \dots, \widehat{\partial_{j_k}}, \dots, \widehat{\partial_{j_l}}, \dots, \partial_{j_{m+1}}) \\
&= \sum_{k=1}^{m+1} (-1)^{k-1} \partial_{j_k} (m! \eta_{i_1 \dots i_m} \delta_{j_1}^{i_1} \dots \delta_{j_{k-1}}^{i_{k-1}} \delta_{j_{k+1}}^{i_{k+1}} \dots \delta_{j_{m+1}}^{i_{m+1}}) \\
&= m! \sum_{k=1}^{m+1} (-1)^{k-1} \partial_{j_k} (\eta_{j_1 \dots j_{k-1} j_{k+1} \dots j_{m+1}}) \tag{1.33}
\end{aligned}$$

where Exercise 1.29 justifies that $m!$ pops out, and where we passed from the second line to the penultimate one by using Equation (1.20). Since the top left hand side of Equation (1.33) is $d\eta$ evaluated on $\partial_{j_1}, \dots, \partial_{j_{m+1}}$, the same Exercise 1.29 implies that it is equal to $(m+1)! (d\eta)_{j_1 \dots j_{m+1}}$. Thus we have:

$$(m+1)! (d\eta)_{j_1 \dots j_{m+1}} = m! \sum_{k=1}^{m+1} (-1)^{k-1} \partial_{j_k} (\eta_{j_1 \dots j_{k-1} j_{k+1} \dots j_{m+1}})$$

Multiplying on the left and on the right by $dx^{j_1} \wedge \dots \wedge dx^{j_{k-1}} \wedge dx^{j_k} \wedge dx^{j_{k+1}} \wedge \dots \wedge dx^{j_{m+1}}$ and contracting the indices, this implies that $d\eta$ reads:

$$\begin{aligned}
d\eta &= \frac{1}{m+1} \sum_{k=1}^{m+1} (-1)^{k-1} \partial_{j_k} (\eta_{j_1 \dots j_{k-1} j_{k+1} \dots j_{m+1}}) dx^{j_1} \wedge \dots \wedge dx^{j_{k-1}} \wedge dx^{j_k} \wedge dx^{j_{k+1}} \wedge \dots \wedge dx^{j_{m+1}} \\
&= \frac{1}{m+1} \sum_{k=1}^{m+1} \partial_{j_k} (\eta_{j_1 \dots j_{k-1} j_{k+1} \dots j_{m+1}}) dx^{j_k} \wedge dx^{j_1} \wedge \dots \wedge dx^{j_{k-1}} \wedge dx^{j_{k+1}} \wedge \dots \wedge dx^{j_{m+1}} \\
&= \partial_{j_k} (\eta_{j_1 \dots j_{k-1} j_{k+1} \dots j_{m+1}}) dx^{j_k} \wedge dx^{j_1} \wedge \dots \wedge dx^{j_{k-1}} \wedge dx^{j_{k+1}} \wedge \dots \wedge dx^{j_{m+1}} \tag{1.34}
\end{aligned}$$

where we have used Equation (1.21) between the first line and the second line, and where we use Einstein summation convention on repeated indices. But then, they are dummy indices and it does not change anything that we write the m indices $j_1, \dots, j_{k-1}, j_{k+1}, \dots, j_{m+1}$ as i_1, \dots, i_m , and j_k as i , at the condition that they appear contracted with themselves in the formula. That is to say, Equation (1.34) can alternatively be written as:

$$d\eta = \partial_i (\eta_{i_1 \dots i_m}) dx^i \wedge dx^{i_1} \wedge \dots \wedge dx^{i_m}$$

which is the required result. □

Exercise 1.46. Check that Equation (1.32) gives back Equation (1.27) when $m = 1$.

Proposition 1.45 allows us to prove the following proposition in a very elegant way:

Proposition 1.47. *The $C^\infty(\mathbb{R}^n)$ -linear morphism $d : \Omega^\bullet(\mathbb{R}^n) \longrightarrow \Omega^{\bullet+1}(\mathbb{R}^n)$ is a differential, i.e. $d \circ d = 0$.*

Proof. We know already by Corollary 1.37 that $d^2f = 0$ for any smooth function. Then, let $\eta = \eta_{i_1 \dots i_m} dx^{i_1} \wedge \dots \wedge dx^{i_m} \in \Omega^m(\mathbb{R}^n)$ be a differential m -form, for $m \geq 1$, and apply twice Equation (1.32):

$$\begin{aligned} d^2(\eta) &= d(\partial_i(\eta_{i_1 \dots i_m}) dx^i \wedge dx^{i_1} \wedge \dots \wedge dx^{i_m}) \\ &= \partial_j \partial_i(\eta_{i_1 \dots i_m}) dx^j \wedge dx^i \wedge dx^{i_1} \wedge \dots \wedge dx^{i_m} \end{aligned}$$

But the element $\partial_j \partial_i(\eta_{i_1 \dots i_m})$, symmetric under a permutation $i \leftrightarrow j$, is contracted with an element $dx^j \wedge dx^i \wedge dx^{i_1} \wedge \dots \wedge dx^{i_m}$, which is skew-symmetric under a permutation $i \leftrightarrow j$. Thus their contraction is zero. \square

Remark 1.48. There is an alternative proof, much more computational, that relies exclusively on the expression of the differential given in Equation (1.31). This proof resembles the proof that one could invoke to show the consistency of the Chevalley-Eilenberg differential in the cohomology theory of Lie algebra (recall that the space of vector fields is a Lie algebra!). Doing this alternative proof is a very good training to understand how differential forms interact with vector fields.

Thus, the graded vector space of differential forms $\Omega^\bullet(\mathbb{R}^n) = (\Omega^m(\mathbb{R}^n))_{0 \leq m \leq n}$, equipped with the differential d is a chain complex. We call it the *de Rham complex* and the differential d is called the *de Rham differential*. It is bounded below and above and it is understood in this complex that for every $m \leq -1$ and every $m \geq n$, $d_m = 0$ (see sequence (1.30)). We conclude this subsection by stating a unicity result that we do not prove, but which is worth knowing:

Proposition 1.49. *The de Rham differential is the unique $C^\infty(\mathbb{R}^n)$ -linear morphism $d : \Omega^\bullet(\mathbb{R}^n) \longrightarrow \Omega^{\bullet+1}(\mathbb{R}^n)$ which satisfies all three following properties:*

1. on smooth functions (i.e. 0-forms), $df(X) = X(f)$;
2. $d \circ d = 0$;
3. for every $\eta \in \Omega^k(\mathbb{R}^n)$ and $\mu \in \Omega^l(\mathbb{R}^n)$:

$$d(\eta \wedge \mu) = (d\eta) \wedge \mu + (-1)^k \eta \wedge (d\mu) \tag{1.35}$$

Proof. This is Theorem 12.14 in [Lee, 2003]. See also the paragraph at the top of page 313 to understand the equivalence between our definition of the de Rham differential and Lee's definition. \square

Remark 1.50. Notice that Equation (1.35) implies that d is a *graded derivation* of the commutative graded algebra $(\Omega^\bullet(\mathbb{R}^n), \wedge)$, turning it into a *differential commutative graded algebra*, abbreviated *cdga* (notice the inversion of the letters in the abbreviation).

Example 1.51. The vector calculus identities. In three dimensional euclidean space \mathbb{R}^3 , Proposition (1.47) will translate under an unexpected form. Recall what we said in Remark 1.38: that exact 1-forms are closed translates as the following identity:

$$\overrightarrow{\text{curl}}(\overrightarrow{\text{grad}}(f)) = 0$$

Let us explain this identity from the perspective of differential forms. We saw in Proposition 1.31 that the gradient of a function f is the image through the musical isomorphism $\sharp : T^*\mathbb{R}^n \longrightarrow T\mathbb{R}^n$ of the differential df via the formula:

$$\overrightarrow{\text{grad}}(f) = (df)^\sharp$$

Let us pursue this analogy.

Assume we work in three dimensional euclidean space \mathbb{R}^3 with standard coordinates x, y, z (so that, for this discussion, x is a coordinate and not a point). We will not use Einstein summation convention either. Let ξ be a differential 1-form: $\xi = \xi_x dx + \xi_y dy + \xi_z dz$. Then Equation (1.28) tells us that:

$$d\xi = \left(\frac{\partial \xi_y}{\partial x} - \frac{\partial \xi_x}{\partial y} \right) dx \wedge dy + \left(\frac{\partial \xi_z}{\partial y} - \frac{\partial \xi_y}{\partial z} \right) dy \wedge dz + \left(\frac{\partial \xi_x}{\partial z} - \frac{\partial \xi_z}{\partial x} \right) dz \wedge dx$$

We recognize the coordinates of the $\overrightarrow{\text{curl}}$ of the vector field $\xi^\# = \xi_x \frac{\partial}{\partial x} + \xi_y \frac{\partial}{\partial y} + \xi_z \frac{\partial}{\partial z}$. Since $d\xi$ is a 2-form, one only needs the Hodge star operator $\star : \Omega^2(\mathbb{R}^3) \rightarrow \Omega^1(\mathbb{R}^3)$ and the musical isomorphisms to reconstruct the desired relation:

$$(\star d\xi)^\# = \overrightarrow{\text{curl}}(\xi^\#)$$

Equivalently, for every vector field X , one has:

$$\overrightarrow{\text{curl}}(X) = (\star d(X^\flat))^\#$$

Next, pick up a differential 2-form $\eta = \eta_{xy} dx \wedge dy + \eta_{yz} dy \wedge dz + \eta_{zx} dz \wedge dx$. Let write η_z instead of η_{xy} , η_x instead of η_{yz} and η_y instead of η_{zx} , for a reason that will soon be transparent. The differential of this 2-form is a 3-form, which, under some simple permutations of dx, dy and dz , can be written as:

$$d\eta = \left(\frac{\partial \eta_x}{\partial x} + \frac{\partial \eta_y}{\partial y} + \frac{\partial \eta_z}{\partial z} \right) dx \wedge dy \wedge dz$$

We recognize, in the parenthesis, the divergence of the vector field $(\star \eta)^\# = \eta_x \frac{\partial}{\partial x} + \eta_y \frac{\partial}{\partial y} + \eta_z \frac{\partial}{\partial z}$. Then, we have the following identity:

$$\star d\eta = \text{div}((\star \eta)^\#) \tag{1.36}$$

The left-hand side is indeed a smooth function because $\star(\Omega^3(\mathbb{R}^3)) = \Omega^0(\mathbb{R}^3) = \mathcal{C}^\infty(\mathbb{R}^3)$. Equivalently, for every vector field X , one has:

$$\text{div}(X) = \star d \star (X^\flat) \tag{1.37}$$

Notice that Equation (1.36) is equivalent to Equation (1.37) because in dimension 3, Equation (A.28) tells us that $\star \star \eta = \eta$ for any 2-form η .

Now let us check that the vector calculus identities in \mathbb{R}^3 amount to $d^2 = 0$. Let f be a smooth function on \mathbb{R}^3 , then:

$$\begin{aligned} \overrightarrow{\text{curl}}(\overrightarrow{\text{grad}}(f)) &= \overrightarrow{\text{curl}}((df)^\#) \\ &= (\star d((df)^\flat))^\# \\ &= (\star \underbrace{d(df)}_{=0})^\# \end{aligned}$$

To pass from the second line to the third line, we used the fact that $\#$ and \flat are inverse to one another. We thus obtain the infamous identity $\overrightarrow{\text{curl}}(\overrightarrow{\text{grad}}(f)) = 0$. Since the Hodge star operator and $\#$ are isomorphisms, we conclude that:

$$\overrightarrow{\text{curl}}(\overrightarrow{\text{grad}}(f)) = 0 \iff d^2 f = 0$$

Now, turning to the next identity: let X be a vector field on \mathbb{R}^3 . Then:

$$\begin{aligned} \operatorname{div}(\overrightarrow{\operatorname{curl}}(X)) &= \operatorname{div}((\star d(X^\flat))^\sharp) \\ &= \star d \star ((\star d(X^\flat))^\sharp)^\flat \\ &= \star d \star (\star d(X^\flat)) \\ &= \star \underbrace{d(d(X^\flat))}_{=0} \end{aligned}$$

Since the Hodge star operator is an isomorphism, we deduce that:

$$\operatorname{div}(\overrightarrow{\operatorname{curl}}(X)) = 0 \iff d \circ d(X^\flat) = 0$$

Hence, the two most famous vector calculus identities $\overrightarrow{\operatorname{curl}} \circ \overrightarrow{\operatorname{grad}} = 0$ and $\operatorname{div} \circ \overrightarrow{\operatorname{curl}} = 0$ are nothing but Proposition 1.47 applied to \mathbb{R}^3 . Thus, since on \mathbb{R}^3 with the euclidean metric $\star^{-1} = \star$, we have the following commutative diagram:

$$\begin{array}{ccccccccc} \mathcal{C}^\infty(\mathbb{R}^3) & \xrightarrow{\overrightarrow{\operatorname{grad}}} & \mathfrak{X}(\mathbb{R}^3) & \xrightarrow{\overrightarrow{\operatorname{curl}}} & \mathfrak{X}(\mathbb{R}^3) & \xrightarrow{\operatorname{div}} & \mathcal{C}^\infty(\mathbb{R}^3) & & \\ \downarrow \operatorname{id} & & \downarrow \flat & & \downarrow \star \flat & & \downarrow \star & & \\ 0 & \longrightarrow & \Omega^0(\mathbb{R}^3) & \xrightarrow{d} & \Omega^1(\mathbb{R}^3) & \xrightarrow{d} & \Omega^2(\mathbb{R}^3) & \xrightarrow{d} & \Omega^3(\mathbb{R}^3) \longrightarrow 0 \end{array}$$

1.4 De Rham cohomology and Maxwell equations

Let (E, d) be a chain complex of vector spaces. Then every map $d_i : E_i \rightarrow E_{i+1}$ has a kernel and an image. We say that an element $x \in E_i$ is *closed* when $d_i x = 0$, whereas it is *exact* when $x = d_{i-1} y$ for some other element $y \in E_{i-1}$. Since $d^2 = 0$, we have the infamous result:

Proposition 1.52. *In a chain complex (E, d) , every exact element is closed.*

The converse (that every closed element is exact) is in general not true, and actually those closed elements that are not exact carry important informations on the problem. That is why mathematicians have defined the following central notion in modern mathematics:

Definition 1.53. *Let (E, d) be a chain complex (of vector spaces). We define its cohomology as the graded vector space $H^\bullet = (H^i)_{i \in \mathbb{Z}}$, where for each $i \in \mathbb{Z}$, the space H^i is called the i -th cohomology group of E and is defined as the quotient:*

$$H^i = \frac{\operatorname{Ker}(d_i)}{\operatorname{Im}(d_{i-1})}$$

We say that the chain complex is exact – equivalently, that it is a resolution – if $H^i = 0$ for every $i \in \mathbb{Z}$.

Remark 1.54. While in our context, the cohomology groups H^i are vector spaces, the word ‘group’ is widely used because the notion of cohomology applies to much more general objects than complexes of vector spaces. In any case, a vector space can be seen as an abelian group, with respect to the vector addition.

Elements of H^i are *equivalence classes* of vectors of E_i . For every element $x \in \text{Ker}(d_i) \subset E_i$, we write $[x]$ the corresponding equivalence class in H^i . We call $[x]$ the *cohomology class* of x . It has the following meaning: in cohomology, x is identified with every other closed element $x' \in E_i$ that can be written as follows:

$$x' = x + d_{i-1}y$$

for some $y \in E_{i-1}$. In such a case we say that x and x' are *cohomologous* and we write $[x] = [x']$. Therefore, any closed element x whose cohomology class is zero, i.e. such that $[x] = 0 \in H^i$, is exact. To every cohomology class $\theta \in H^i$, there exist an infinite number of representatives, i.e. those closed elements $x \in E_i$ such that $[x] = \theta$, because $x + dy$ would be another valid representative. A priori, there is no better choice of representative, except in certain cases (as we may see later).

The cohomology of the de Rham complex is called the *de Rham cohomology*. We write the cohomology groups of the de Rham complex as $H_{dR}^i(\mathbb{R}^n)$. Lemma 1.39 has shown that closed 1-forms are exact. That is to say, that $H_{dR}^1(\mathbb{R}^n) = 0$. This is actually much more general:

Proposition 1.55. *The de Rham cohomology of \mathbb{R}^n satisfies:*

$$H_{dR}^i(\mathbb{R}^n) \simeq \begin{cases} \mathbb{R} & \text{if } i = 0 \\ 0 & \text{otherwise} \end{cases}$$

Proof. This is a consequence of Poincaré's Lemma, which states that the de Rham cohomology on every star-shaped open set (of a smooth manifold) is trivial (except for the 0-th cohomology group). See Theorem 15.11 in [Lee, 2003]. \square

What kind of objects span the 0-th group of de Rham cohomology $H_{dR}^0(\mathbb{R}^n)$? We have the following situation:

$$0 \longrightarrow \Omega^0(\mathbb{R}^n) \xrightarrow{d} \Omega^1(\mathbb{R}^n) \xrightarrow{d} \dots$$

Then $H_{dR}^0(\mathbb{R}^n) = \text{Ker}(d_0)$. Since d_0 is the morphism associating, to every function f , its differential, one deduces that $df = 0$ if and only if f is a constant function. Then $H_{dR}^0(\mathbb{R}^n) = \{\text{constant functions on } \mathbb{R}^n\}$, which is indeed a one-dimensional space. Another simple example sits at the other side of the chain complex: we know that $\Omega^n(\mathbb{R}^n)$ is one-dimensional and spanned by the standard volume form $\omega = dx^1 \wedge \dots \wedge dx^n$. Since $d(\Omega^n(\mathbb{R}^n)) = 0$, Proposition (1.55) tells us that there should be a differential $n-1$ -form ν such that $\omega = d\nu$. There are several actually: for example $x_1 dx^2 \wedge \dots \wedge dx^n$ or, more generally, those of the form $(-1)^{k-1} x_k dx^1 \wedge \dots \wedge dx^{k-1} \wedge dx^{k+1} \wedge \dots \wedge dx^n$.

Let us now apply all this machinery to Maxwell equations. They are equations that the electric field \vec{E} and the magnetic field \vec{B} should satisfy. Recall what they are (in *three-dimensional space*):

$$\text{div}(\vec{E}) = \rho \tag{1.38}$$

$$\text{div}(\vec{B}) = 0 \tag{1.39}$$

$$\overrightarrow{\text{curl}}(\vec{E}) + \frac{\partial \vec{B}}{\partial t} = 0 \tag{1.40}$$

$$\overrightarrow{\text{curl}}(\vec{B}) - \frac{\partial \vec{E}}{\partial t} = \vec{j} \tag{1.41}$$

We have used the *rationalized Planck units*, where:

$$c = 4\pi G = \hbar = \varepsilon_0 = k_B = 1$$

Although \vec{E} and \vec{B} are usually considered as vector fields, the discussion in Example 1.51 has shown that using the musical isomorphisms allow us to adopt a much more synthesized perspective. However, from the knowledge we have of the differences between the respective behavior of the electric and the magnetic field, we expect that they would not carry the same degree as differential forms. Let us be more specific.

Let M be Minkowski space, i.e. \mathbb{R}^4 equipped with a metric $\eta_{\mu\nu}$ of signature $(3, 1)$ – the indices ranging from 0 to 3, corresponding to the coordinates t, x, y and z . In other words: $\eta_{00} = -1$, and $\eta_{ii} = +1$ for $1 \leq i \leq 3$. The other components of the metric vanish. The volume form would then be $\omega = dt \wedge dx \wedge dy \wedge dz$. The order is important here because if we had taken t to be the fourth coordinate, then the corresponding volume form $dx \wedge dy \wedge dz \wedge dt$ would be minus ω . This would have repercussions on the definition of the Hodge star operator. The electric and magnetic fields are 1-forms and 2-forms on \mathbb{R}^4 , respectively:

$$E = E_x dx + E_y dy + E_z dz \quad (1.42)$$

and

$$B = B_x dy \wedge dz + B_y dz \wedge dx + B_z dx \wedge dy \quad (1.43)$$

So, in particular, $E = \vec{E}^\flat$ and $B = \star(\vec{B}^\flat \wedge dt)$. We define the *field strength* as the following differential 2-form on M :

$$F = B + E \wedge dt$$

In particular, this 2-form decomposes on the canonical frame $dx^\mu \wedge dx^\nu$ of the vector bundle $\wedge^2 T^*M$, as:

$$F = \frac{1}{2} F_{\mu\nu} dx^\mu \wedge dx^\nu, \quad \text{where} \quad F_{\mu\nu} = \begin{pmatrix} 0 & -E_x & -E_y & -E_z \\ E_x & 0 & B_z & -B_y \\ E_y & -B_z & 0 & B_x \\ E_z & B_y & -B_x & 0 \end{pmatrix} \quad (1.44)$$

The current density \vec{j} can be merged with the charge density ρ into a 4-vector $\vec{J} = \rho \frac{\partial}{\partial t} + \vec{j}$. Using the musical isomorphism \sharp , we transform this 4-vector into a differential 1-form J called the *current*:

$$J = -\rho dt + j_x dx + j_y dy + j_z dz$$

This allows us to have synthesized Maxwell equations:

Proposition 1.56. Geometric Maxwell equations Equations (1.39) and (1.40) are equivalent to the Bianchi identity:

$$dF = 0 \quad (1.45)$$

whereas Equations (1.38) and (1.41) are equivalent to:

$$\star d \star F = J \quad (1.46)$$

Proof. Equation (1.45) contains two terms:

$$dF = dB + dE \wedge dt \quad (1.47)$$

because $d^2t = 0$. Let us focus on the first term dB , using Proposition (1.45), and deleting the terms containing $dx \wedge dx$, $dy \wedge dy$ or $dz \wedge dz$:

$$\begin{aligned} dB &= \partial_x B_x dx \wedge dy \wedge dz + \partial_y B_y dy \wedge dz \wedge dx + \partial_z B_z dz \wedge dx \wedge dy \\ &\quad + \partial_t B_x dt \wedge dy \wedge dz + \partial_t B_y dt \wedge dz \wedge dx + \partial_t B_z dt \wedge dx \wedge dy \\ &= \operatorname{div}(\vec{B}) dx \wedge dy \wedge dz + \partial_t B_x dy \wedge dz \wedge dx + \partial_t B_y dz \wedge dx \wedge dt + \partial_t B_z dx \wedge dy \wedge dt \end{aligned}$$

On the other hand, the second term of Equation (1.47) can be written as:

$$\begin{aligned} dE \wedge dt &= \partial_y E_x dy \wedge dx \wedge dt + \partial_z E_x dz \wedge dx \wedge dt \\ &\quad + \partial_x E_y dx \wedge dy \wedge dt + \partial_z E_y dz \wedge dy \wedge dt \\ &\quad + \partial_x E_z dx \wedge dz \wedge dt + \partial_y E_z dy \wedge dz \wedge dt \\ &= (\partial_x E_y - \partial_y E_x) dx \wedge dy \wedge dt + (\partial_y E_z - \partial_z E_y) dy \wedge dz \wedge dt \\ &\quad + (\partial_z E_x - \partial_x E_z) dz \wedge dx \wedge dt \end{aligned}$$

Writing $dB + dE \wedge dt = 0$, one obtains the following identity:

$$\begin{aligned} 0 &= \operatorname{div}(\vec{B}) dx \wedge dy \wedge dz + (\partial_t B_x + \partial_y E_z - \partial_z E_y) dy \wedge dz \wedge dt \\ &\quad + (\partial_t B_y + \partial_z E_x - \partial_x E_z) dz \wedge dx \wedge dt + (\partial_t B_z + \partial_x E_y - \partial_y E_x) dx \wedge dy \wedge dt \end{aligned}$$

Thus, each term in parenthesis is equal to zero, and we obtain Equations (1.39) and (1.40). \square

Exercise 1.57. Using the fact that the volume form is $\omega = dt \wedge dx \wedge dy \wedge dz$ in our convention for Minkowski space, show that:

$$(\star F)_{\mu\nu} = \begin{pmatrix} 0 & B_x & B_y & B_z \\ -B_x & 0 & E_z & -E_y \\ -B_y & -E_z & 0 & E_x \\ -B_z & E_y & -E_x & 0 \end{pmatrix}$$

and prove Equation (1.46). Beware of the timelike direction dt that satisfies $\langle dt, dt \rangle = -1$ in Minkowski space.

The Bianchi identity (1.45) implies that the field strength is a closed 2-form. We have seen in Proposition 1.55 that the de Rham cohomology is vanishing, except for zero forms. Then, it means that F is an exact form, i.e. there exists a differential 1-form A such that:

$$F = dA \tag{1.48}$$

Using the musical isomorphism \sharp on the 1-form:

$$A = A_\mu dx^\mu = -V dt + A_x dx + A_y dy + A_z dz \tag{1.49}$$

it gives a vector field on M that is written $A^\sharp = V \frac{\partial}{\partial t} + \vec{A}$, where V is the *scalar potential* and \vec{A} is the *vector potential*. To keep this analogy in mind, we often call the differential 1-form A a potential for F . Obviously, in the case where the 2nd group of de Rham cohomology is not zero (this does not happen in \mathbb{R}^n but could happen on other smooth manifolds), it may not be possible to find a vector potential for F . That is why Equation (1.45) is a topological condition. The physical information is contained in Equation (1.46): it is a necessary condition to the existence of a potential A for F .

Exercise 1.58. Check that Equation (1.48), with the potential $A = -V dt + A_x dx + A_y dy + A_z dz$ is equivalent to the two equations:

$$\vec{E} = -\overrightarrow{\text{grad}}(V) - \frac{\partial \vec{A}}{\partial t} \quad (1.50)$$

$$\vec{B} = \overrightarrow{\text{curl}}(\vec{A}) \quad (1.51)$$

Where $\overrightarrow{\text{curl}}$ and $\overrightarrow{\text{grad}}$ are considered to be the usual operators in \mathbb{R}^3 .

Reinjecting Equation (1.48) in Equation (1.46), one obtains the following identity:

$$\star d \star dA = J \quad (1.52)$$

We will see later that $\star d \star d$ is (in Minkowski space) the d'Alembertian operator $\square = \frac{\partial^2}{\partial t^2} - \Delta$, so that one may show that Equation (1.52) is equivalent to the following two equations:

$$\Delta V + \frac{\partial}{\partial t} \text{div}(\vec{A}) = -\rho \quad (1.53)$$

$$\overrightarrow{\square} \vec{A} + \overrightarrow{\text{grad}} \left(\text{div}(\vec{A}) + \frac{\partial V}{\partial t} \right) = \vec{j} \quad (1.54)$$

Under the assumption that \vec{E} and \vec{B} are related to \vec{A} and V through Equations (1.50) and (1.51), Equations (1.53) and (1.54) are equivalent to Equations (1.38) and (1.41). Hence we see that the geometric Maxwell equations are equivalent to the classical Maxwell equations. The Bianchi identity is a topological condition, automatically satisfied in M (but not on every manifold), so that the existence of A depends on the possibility of solving Equation (1.52).

The fact that the choice of potential A is fixed, up to an exact 1-form df – because $d(A + df) = dA = F$ – implies that one can make a specific choice for A that possibly simplifies Equations (1.53) and (1.54). When we make such a choice, we say that we fix the *gauge*. Let us choose the *Lorenz gauge*⁵, defined by the condition:

$$\partial_\mu A^\mu = 0 \quad (1.55)$$

The notation A^μ symbolizes the components of the vector field A^\flat , so $A^\mu = \eta^{\mu\nu} A_\nu$: $A^0 = V$, and $A^i = A_i$ for $1 \leq i \leq 3$. Then, Equation (1.55) translates as:

$$\frac{\partial V}{\partial t} + \text{div}(\vec{A}) = 0$$

In this gauge, Equations (1.53) and (1.54) become:

$$\begin{aligned} \square V &= \rho \\ \overrightarrow{\square} \vec{A} &= \vec{j} \end{aligned}$$

Fixing a gauge allows to obtain differential equations that may be easier to solve. There are several gauges in electromagnetism: the *Coulomb gauge*, where $\text{div}(\vec{A}) = 0$; the *Weyl* – or *temporal* – *gauge*, where $V = 0$. Electromagnetism is one of the simplest gauge theories. Its straightforward generalization is the Yang-Mills theory, whose study is postponed to a later chapter.

⁵From the Danish physicist Ludvig Lorenz, not to be confused with the Dutch physicist Hendrik Lorentz, to whom we attribute the *Lorentz transformations* in the theory of relativity, nor with the American physicist Edward Lorenz, who gave his name to the attractor looking like a butterfly in dynamical systems.

2 Differential calculus on smooth manifolds

In this Chapter, we will introduce the notion of smooth manifold and, relying on the mathematical background of the first chapter, develop the machinery needed to study action functionals and turn to more involved topics. The material presented in Chapter 1 will be central to the present chapter, because we will soon understand that a smooth manifold is *locally* like \mathbb{R}^n . It means that at least locally, in a neighborhood of a point, we should think of a smooth manifold as a n -dimensional vector space. The tangent bundle and the cotangent bundle on a smooth manifold, although defined globally, are thus always locally trivial. Differential forms on a manifold are thus always locally exact (because de Rham cohomology on \mathbb{R}^n is almost trivial). Integration of differential forms, though, needs considering the global structure of the manifold. That is why it is often used to probe the topological structure of the manifold, e.g. in topological field theories.

2.1 Smooth manifolds

We emphasize in this presentation the role of functions on manifolds. There is indeed a deep relationship between a manifold, and the algebra of functions on this manifold. One should consider that defining a smooth manifold M from its topology and additional properties satisfied by the open sets is actually equivalent to characterizing what are smooth functions on this manifold M . This point of view illustrates the equivalence between the geometrico-analytic point of view, and the algebraic point of view:

$$\begin{array}{ccc} \text{Geometry} & & \text{Algebra} \\ M & \longleftrightarrow & \mathcal{C}^\infty(M) \end{array}$$

A smooth manifold is a particular case of a topological manifold which, in turn, is defined as follows:

Definition 2.1. *A topological manifold of dimension n is a topological space M (i.e. a set equipped with a topology of open subsets), that is:*

1. Hausdorff, i.e. *points can be separated by neighborhoods: for every pair of points $x, y \in M$, there are disjoint open subsets $U, V \subset M$ such that $x \in U$ and $y \in V$;*
2. second-countable, i.e. *there exists a countable basis for the topology of M ;*
3. locally euclidean, i.e. *every point of M has a neighborhood that is homeomorphic to an open subset of \mathbb{R}^n .*

The first property is a minimal assumption to avoid pathological cases that are not fit to do analysis. The second property means that the topology is generated by a countable family of open sets. This axiom is desirable as in (non necessarily Hausdorff) second-countable spaces, compactness, sequential compactness, and countable compactness are all equivalent properties. In non-compact Hausdorff spaces, second-countability can then be interpreted as a weaker version of countable compactness:

Proposition 2.2. *Let M be a locally euclidean Hausdorff topological space, then M is second-countable if and only if M is paracompact – i.e. every open cover of M has a locally finite open refinement – and has countably many connected components.*

Proof. See Proposition 2.24 and Exercise 2.15 in [Lee, 2003]. □

This consequence is crucially needed to define partitions of unity, which are central to define integration on smooth manifolds and metrics on a manifold. The last property of Definition 2.1 means, more precisely, that for every point $x \in M$, there exists an open neighborhood U of x and an open subset $\tilde{U} \subset \mathbb{R}^n$, together with a homeomorphism $\varphi : U \rightarrow \tilde{U}$ from U onto its image. We call the pair (U, φ) a *chart* or *coordinate chart* on M . At the cost of translating the image of the map φ in \mathbb{R}^n , one can always send x to $0 \in \mathbb{R}^n$. We then say that the chart is *centered at x* ; every chart can be made centered at x by subtracting the vector $\varphi(x)$. Denoting by x^1, \dots, x^n the standard coordinates centered at 0 on \mathbb{R}^n , we often define by abuse of notation the composite functions $x^i \circ \varphi$ with the same letters x^i . We then call the continuous functions x^1, \dots, x^n *local coordinates at x* . We define an *atlas* for M to be a collection \mathcal{A} of charts whose domains cover M . Let us now give three pathological examples illustrating why we need the three assumptions in Definition 2.1.

Example 2.3. The ‘line with two origins’ is second-countable and locally euclidean, but not Hausdorff. It is obtained as the quotient of the union of the two horizontal lines $\{(x, y) \in \mathbb{R}^2 \mid y = 1\}$ and $\{(x, y) \in \mathbb{R}^2 \mid y = -1\}$ (with their respective subspace topology) under the following relation: $(x, 1) \sim (x, -1)$, whenever $x \neq 0$. Due to this very particular choice of quotient, the two origins cannot be separated by neighborhoods.

Example 2.4. The ‘long line’ is Hausdorff and locally Euclidean but not second-countable. It consists of segments $[0, 1[$ glued one after the other, but uncountably many times (contrary to the real line). The ‘long ray’ is the cartesian product $L = \omega_1 \times [0, 1[$ equipped with the order topology that arises from the lexicographical order on L . The long line is obtained by putting together a long ray in each direction (positive and negative).

Example 2.5. The ‘figure eight’ is Hausdorff and second-countable but not locally Euclidean at the origin.

Example 2.6. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^k$ be a continuous function. The *graph* of f is the subset of $\mathbb{R}^n \times \mathbb{R}^k$:

$$\Gamma(f) = \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^k \mid y = f(x)\}$$

Equipped with the subspace topology, it is a topological manifold. Indeed, denoting the projection on the first factor $pr_1 : (x, y) \mapsto x$, we set $\varphi = pr_1|_{\Gamma(f)}$. Then φ is a continuous surjective map that has a continuous inverse: $\varphi^{-1}(x) = (x, f(x))$. Then it is a homeomorphism, and $(\Gamma(f), \varphi)$ is a global coordinate chart, turning $\Gamma(f)$ into a topological manifold of dimension n .

The fact that topological manifolds of dimension n are locally homeomorphic to \mathbb{R}^n implies that we may be able to do differential calculus on it. For instance, given a continuous function $f : M \rightarrow \mathbb{R}$ and a chart (U, φ) on M , one could consider the composition $f \circ \varphi^{-1} : \tilde{U} \rightarrow \mathbb{R}$, which is a real-valued function whose domain is an open subset \tilde{U} of a Euclidean space. Then it would make sense to say that f is *smooth* if and only if for every chart (U, φ) on M , $f \circ \varphi^{-1}$ is infinitely differentiable. However, this definition is not stable when passing from one open set U to another open set V , for the following reason: let (U, φ) and (V, ψ) be two charts whose underlying open sets U and V are overlapping; then, the *transition map* $\varphi \circ \psi^{-1}$ is a homeomorphism from $\psi(U \cap V)$ to $\varphi(U \cap V)$, both open subsets of \mathbb{R}^n . However, this map is not necessarily smooth, and this has the following consequence when we write f over the intersection $U \cap V$:

$$f \circ \psi^{-1} = f \circ \varphi^{-1} \circ (\varphi \circ \psi^{-1})$$

Then, even if $f \circ \varphi^{-1}$ and $f \circ \psi^{-1}$ are both differentiable, it does not imply that the function $\varphi \circ \psi^{-1}$ is, which is a bit problematic regarding the derivation rule of composite functions: $\partial_k(g \circ h) = \sum_{i=1}^n \partial_k(h^i) \partial_i(g)$ that one should expect in differential calculus.

To solve this issue, one should restrict the choice of coordinate charts adapted to the topological space M and pick up only a sub-family of those, that are ‘compatible’, i.e. such that the

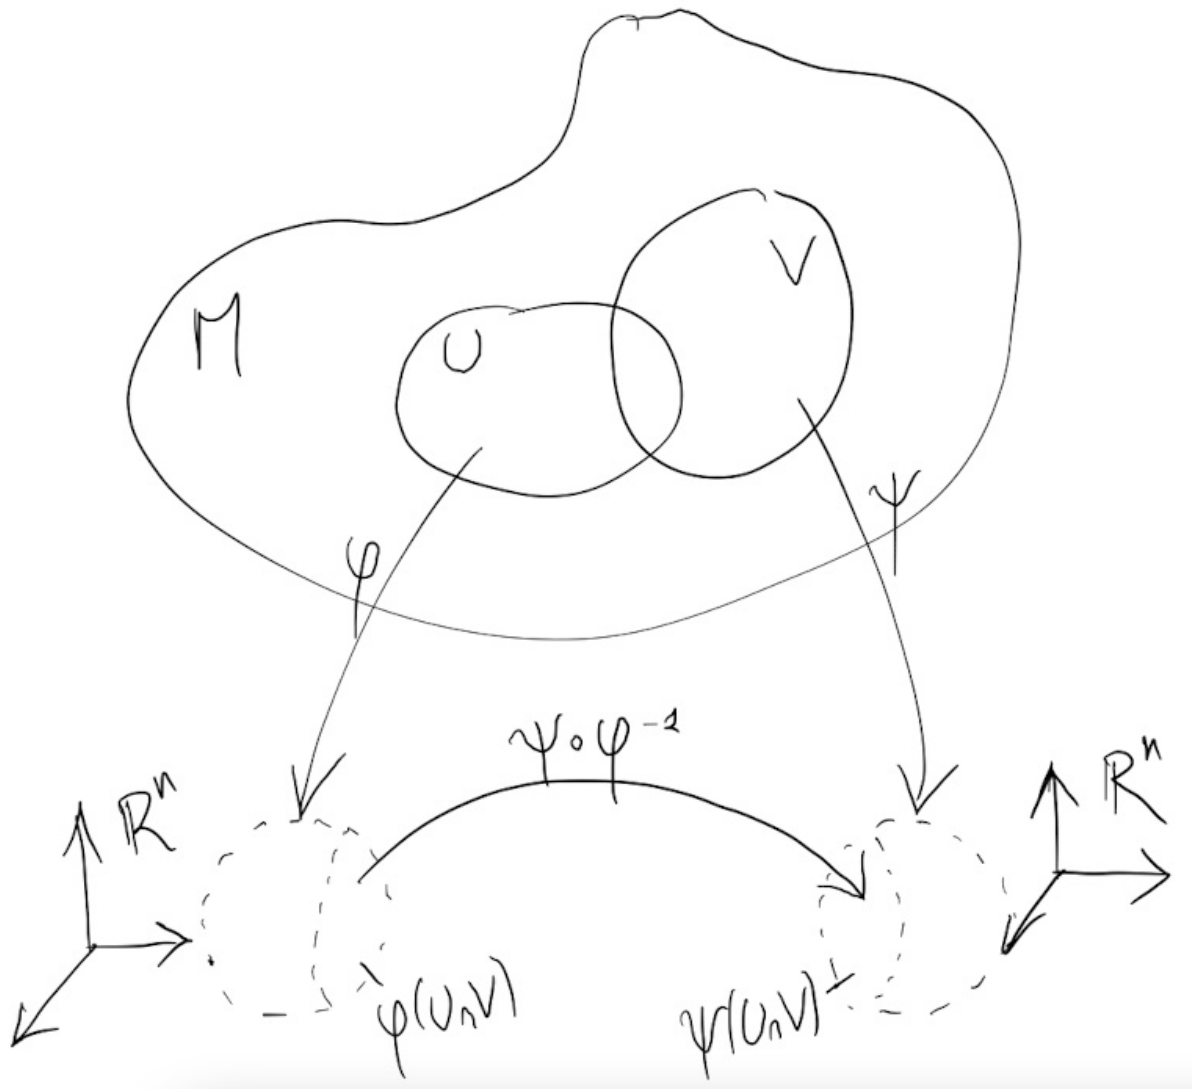


Figure 8: Two overlapping charts (U, φ) and (V, ψ) are smoothly compatible if the map $\psi \circ \varphi^{-1} : \varphi(U \cap V) \rightarrow \psi(U \cap V)$ is a diffeomorphism. A smooth atlas is a collection of smoothly compatible charts covering M .

transition functions between two charts of that family are smooth. More precisely, two charts (U, φ) and (V, ψ) are said to be *smoothly compatible* if either $U \cap V = \emptyset$ or the transition map $\varphi \circ \psi^{-1} : \psi(U \cap V) \rightarrow \varphi(U \cap V)$ is a diffeomorphism, i.e. a smooth homeomorphism from $\psi(U \cap V)$ to $\varphi(U \cap V)$, whose inverse is smooth as well. An atlas \mathcal{A} is called a *smooth atlas* if any two charts in \mathcal{A} are smoothly compatible with each other. Obviously a given (topological) atlas on M can give rise to several smooth atlases if, for instance, two families of charts covering M are smoothly compatible within the families, but not between them. Given a smooth atlas \mathcal{A} on M , one says that a chart is *smoothly compatible* with the atlas, if this chart is smoothly compatible with every chart comprised in \mathcal{A} . The union of all compatible charts to a given smooth atlas \mathcal{A} then defines a smooth atlas that is said *maximal*: it is not contained in any strictly larger smooth atlas. Such a smooth atlas is always very large since it contains every possible choice of smoothly compatible charts on the topological manifold M . Alternatively, one can work with equivalence classes of smooth atlases: two smooth atlases \mathcal{A} and \mathcal{A}' are considered equivalent if every chart of \mathcal{A} is smoothly compatible with \mathcal{A}' . That allows working

on a manifold with a single smooth atlas, consisting of only a few and practical charts, with the implicit understanding that many other charts and differentiable atlases are equally legitimate. Then, maximal smooth atlases are distinguished representents of their respective equivalence classes of compatible smooth atlases. Lemma 1.10 in [Lee, 2003] provides some understanding of the relationship between maximal smooth atlases and equivalence classes of smooth atlases:

Lemma 2.7. *Let M be a topological manifold of dimension n .*

1. *Every smooth atlas for M is contained in a unique maximal smooth atlas.*
2. *Two smooth atlases are equivalent if and only if their union is a smooth atlas.*

In particular, this shows that there may exist non-equivalent maximal smooth atlases for a given topological manifold M . Then, we can now define the central definition of this subsection:

Definition 2.8. *A smooth structure on a topological n dimensional manifold M is a maximal smooth atlas \mathcal{A} . A smooth manifold of dimension n is a pair (M, \mathcal{A}) – often only written M , omitting \mathcal{A} – where M is a topological manifold of dimension n and \mathcal{A} is a smooth structure on M .*

Remark 2.9. The smooth structure is an additional piece of data added to a topological manifold M . Most topological manifolds have uncountably many different smooth structures, but there exist topological manifolds that do not admit any smooth structure.

Example 2.10. The vector space \mathbb{R}^n is a smooth manifold, when equipped with the chart $(\mathbb{R}^n, \text{id}_{\mathbb{R}^n})$: the smooth structure consists of all the charts on \mathbb{R}^n that are compatible with the first one.

Exercise 2.11. Check that the following charts on the 2-sphere are smoothly compatible:

$$\begin{aligned} U_x^+ &= \{(x, y, z) \in \mathbb{S}^2 \mid x > 0\} && \text{(resp. } U_x^- \text{ for } x < 0) \\ U_y^+ &= \{(x, y, z) \in \mathbb{S}^2 \mid y > 0\} && \text{(resp. } U_y^- \text{ for } y < 0) \\ U_z^+ &= \{(x, y, z) \in \mathbb{S}^2 \mid z > 0\} && \text{(resp. } U_z^- \text{ for } z < 0) \end{aligned}$$

and thus induce a smooth structure on \mathbb{S}^2 (the smooth atlas of every chart compatible with the above three charts). This kind of charts can be generalized to every n -sphere and defines the standard smooth structure on the n -sphere.

It turns out that if the dimension of the topological manifold M is higher than or equal to 1, then it has uncountably many distinct smooth structures (see Problem 1.3 in [Lee, 2003]). Thus we would like a notion of equivalence of smooth structures that mimic the topological equivalence of homeomorphic topological spaces: for this reason we introduce the notion of *diffeomorphism*. Let M, N be smooth manifolds, and let $f : M \rightarrow N$ be any map (of sets). We say that f is a *smooth map* if for any $x \in M$, there exist smooth charts (U, φ) containing x and (V, ψ) containing $f(x)$ such that $f(U) \subset V$ and the composite map $\psi \circ f \circ \varphi^{-1}$ is smooth in the usual sense (i.e. infinitely differentiable) from $\varphi(U)$ to $\psi(V)$. The smooth map f is a *diffeomorphism* if it is bijective, and its inverse f^{-1} is smooth as well. The coordinate map φ of a smooth chart (U, φ) is a diffeomorphism onto its image $\varphi(U)$. While homeomorphisms define an equivalence relation between topological manifolds, diffeomorphisms define an equivalence relation between smooth manifolds. This relation allows to probe inequivalent smooth structures, for there exist topological manifolds admitting several smooth structures that are not diffeomorphic to one another. Finally, it is always useful to have the local variant of the former notion: $f : M \rightarrow N$ is called a *local diffeomorphism* if every point $x \in M$ has a neighborhood U such that $f(U)$ is open in N and $f : U \rightarrow f(U)$ is a diffeomorphism (onto its image).

Example 2.12. The euclidean vector space \mathbb{R}^n has a unique smooth structure (up to diffeomorphism) unless $n = 4$, in which case \mathbb{R}^4 admits an uncountable number of non-diffeomorphic smooth structures, and these are called *exotic* \mathbb{R}^4 . See [Lee, 2003, p. 37] for more details on this deep and exciting topic.

Example 2.13. The situation for the spheres is a bit more complicated. The following table shows how many smooth *types*, i.e. smooth-structures up to diffeomorphism, a n -sphere admits:

Dim.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Types	1	1	1	≥ 1	1	1	28	2	8	6	992	1	3	2	16256	2	16	16	523264	24

The n -spheres whose smooth structure is not diffeomorphic to the standard one are called *exotic sphere*. It is not known yet how many types the 4-sphere possesses.

A smooth map $f : M \rightarrow \mathbb{R}$ is called a *smooth function* on M . The set of smooth functions on a smooth manifold M is denoted $\mathcal{C}^\infty(M)$. It is a commutative associative unital \mathbb{R} -algebra. The definition applies locally as well: any open set U of M inherits a smooth structure by restriction of the atlas to U (this can be seen by applying Lemma 1.23 in [Lee, 2003] to U); then we note $\mathcal{C}^\infty(U)$, or $\Omega^0(U)$, the space of functions on the open set $U \subset M$. Not every function in $\mathcal{C}^\infty(U)$ descend from a function in $\mathcal{C}^\infty(M)$: for example $]0, 1[$ is a smooth manifold, whose smooth structure is inherited from its embedding in \mathbb{R} , but there are functions on $]0, 1[$ that *do not* descend from functions on \mathbb{R} , e.g. $f : x \mapsto \frac{1}{x(1-x)}$. Thus we see that $\mathcal{C}^\infty(U)$ is not a subalgebra of $\mathcal{C}^\infty(M)$. Rather, the assignment $U \rightarrow \mathcal{C}^\infty(U)$ which associates to any open set a commutative associative unital \mathbb{R} -algebra is what is called a *sheaf of (commutative associative unital) \mathbb{R} -algebras* over M . There is a deep relationship between smooth manifolds and their algebras of functions. As for finite dimensional vector spaces, where the dual space E^* is an alternative characterization of a given vector space E , we expect some sort of duality between a smooth manifold M and its space of smooth functions $\mathcal{C}^\infty(M)$. There exists such a result in operator algebra:

Theorem 2.14. Gel'fand duality *For every arbitrary unital commutative C^* -algebra A there exists a compact Hausdorff topological space X such that A is equivalent to the algebra of complex-valued continuous functions on X : $A \simeq C(X)$. More precisely, there exists an equivalence of categories between the (opposite) category of unital commutative C^* -algebras and the category of compact Hausdorff topological spaces.*

The idea is not to understand this theorem but to see that for any given algebra of a certain type, there exists a geometric space such that this algebra plays the role of a subalgebra of functions – or operators – on this space. We expect this result to hold as well for smooth manifolds, that is to say: to any commutative, associative algebra with unit over \mathbb{R} with some additional property, one can associate a smooth manifold, in the sens of Definition 2.8. Using a metaphor with physics, the algebra of functions would be considered as ‘physical observables’, and the associated smooth manifold would be what ‘could be observed’ by using these functions. It is thus meaningful that, given a different choice of observables, then what could be observed would change, and thus the associated manifold. There exists such a correspondence in algebraic geometry, between a choice of a commutative ring R , and its associated set of points that we call the *spectrum* of R : it is the set of prime ideals of R and is denoted $\text{Spec}(R)$. Then, the ring R is considered as playing the role of the ring of functions on $\text{Spec}(R)$. Then a *scheme* is a topological space X admitting a covering by open sets U_i , such that each U_i is the spectrum of a given ring R_i .

We can define a smooth manifold using the same kind of ideas. Let us start from a commutative, associative algebra with unit over \mathbb{R} denoted \mathcal{C} which will play the role of the algebra of smooth functions $\mathcal{C}^\infty(M)$ on the manifold M yet to define. Drawing an analogy from finite dimensional vector spaces, for which the dual of the dual of E is E (this is not true anymore in infinite dimension), we define M – also denoted $|\mathcal{C}|$ – to be the ‘dual’ of \mathcal{C} , i.e. the set of all \mathbb{R} -algebra homomorphisms to \mathbb{R} :

$$|\mathcal{C}| = \{x : \mathcal{C} \longrightarrow \mathbb{R}, f \longmapsto x(f) \text{ is an } \mathbb{R}\text{-algebra homomorphism}\}$$

To this set we can associate an algebra of ‘physical observables’ $\tilde{\mathcal{C}}$, i.e. the \mathbb{R} -algebra of objects $\tilde{f} : |\mathcal{C}| \longrightarrow \mathbb{R}$ associated to some $f \in \mathcal{C}$ via the formula $\tilde{f}(x) = x(f)$. It turns out that \mathcal{C} is surjective onto $\tilde{\mathcal{C}}$, but not injective because there may be some non-trivial element $f \in \mathcal{C}$ which satisfies $x(f) = 0$ for every x . Since we want the elements of \mathcal{C} to be in one-to-one correspondence with the physical observables on $M = |\mathcal{C}|$, we require \mathcal{C} to satisfy the additional assumption that:

$$\bigcap_{x \in |\mathcal{C}|} \text{Ker}(x) = 0$$

This condition is equivalent to saying that every element of \mathcal{C} ‘observe’ at least something – for if $x(f) = 0$ for every $x \in |\mathcal{C}|$, the element f could not be used as a physical observable. Then, under this assumption, one can show that \mathcal{C} becomes canonically isomorphic to the algebra of ‘observables’ $\tilde{\mathcal{C}}$. Equipping the set $|\mathcal{C}|$ with the weakest topology for which all such functions are continuous turns $M = |\mathcal{C}|$ into a Hausdorff topological space. Then, \mathcal{C} can be identified through its isomorphism with $\tilde{\mathcal{C}}$ as a subalgebra of the algebra of continuous functions on M .

At this point, one would expect that the algebra \mathcal{C} represents smooth functions on M . However this claim is still far from reality. A naive postulate would be to additionally require that \mathcal{C} be locally isomorphic to $\mathcal{C}^\infty(\mathbb{R}^n)$ i.e., by assuming that there exists an open cover of M with a family of open sets U_i , $M = \bigcup_i U_i$, such that the restriction of \mathcal{C} to each U_i is isomorphic to $\mathcal{C}^\infty(\mathbb{R}^n)$ as a \mathbb{R} -algebra. This would be the algebraic way of saying that the manifold M is locally like \mathbb{R}^n . However, this is mathematically not sufficient or does not define a smooth manifold as we understand it. The precise condition that one should require on \mathcal{C} is much more subtle and very close to mathematical notions that are commonly used in algebraic geometry. Since it is not the topic of the current course, I refer to [Nestruev, 2003] for precise statements:

Theorem 2.15. *There is an equivalence of categories between the category of smooth manifolds and the category of complete geometric commutative associative unital \mathbb{R} -algebras.*

$$\begin{array}{ccc} \textit{Geometry} & & \textit{Algebra} \\ M & \longrightarrow & \mathcal{C}^\infty(M) \\ |\mathcal{C}| & \longleftarrow & \mathcal{C} \end{array}$$

What should be remembered from this discussion, is that there exists a canonical one-to-one correspondence between smooth manifolds and commutative associative unital \mathbb{R} -algebras satisfying additional specific conditions mirroring the topological and differential properties characterizing smooth manifolds. This bijective correspondence will be used frequently when we study gauge theories and constraint surfaces, and ultimately will be the fundamental characterization of graded manifolds in graded geometry.

2.2 Vector bundles, pushforwards, pullbacks

In this section we are interested in smooth maps, and there associated pushforwards and pullbacks. Most notions that we have seen in Chapter 1 will be understood as the ‘local’ versions of

the objects presented in the present section. We have seen that over \mathbb{R}^n a vector bundle is always trivial. This property will only be observed locally for vector bundles over smooth manifolds:

Definition 2.16. A vector bundle of rank k over M is a topological space E together with a surjective continuous map $\pi : E \rightarrow M$, satisfying the two following conditions:

1. for every $x \in M$, the preimage $\pi^{-1}(x) \subset E$ is a k -dimensional vector space, called the fiber of E at x and denoted E_x ;
2. for each $x \in M$, there exists a neighborhood U of x in M and a homeomorphism $\Phi_U : \pi^{-1}(U) \rightarrow U \times \mathbb{R}^k$ (called a local trivialization of E over U), making the following triangle commutative:

$$\begin{array}{ccc}
 \pi^{-1}(U) & \xrightarrow{\Phi_U} & U \times \mathbb{R}^k \\
 \pi \searrow & & \swarrow pr_1 \\
 & U &
 \end{array}$$

where $pr_1 : U \times \mathbb{R}^k \rightarrow U$ is the projection on the first variable; and such that for every $y \in U$, the restriction of Φ_U to E_y is a linear isomorphism from E_y to $\{y\} \times \mathbb{R}^k \simeq \mathbb{R}^k$.

The space E is called the total space of the bundle, M is called its base, and π is called its projection. If E is a smooth manifold, π is a smooth map, and the local trivializations can be chosen to be diffeomorphisms, then E is said to be a smooth vector bundle. If there exists a local trivialization over all of M (called a global trivialization of E), then E is said to be a trivial bundle. In this case, E itself is homeomorphic (resp. diffeomorphic if E is smooth) to the product space $M \times \mathbb{R}^k$.

Every point x of M admits a tangent space $T_x M$, whose definition is straightforward since it does not depend on the neighboring points of x : the tangent space to M at a given point x is the vector space of linear morphisms that are derivations at x , i.e. all the maps $X_x : \mathcal{C}^\infty(M) \rightarrow \mathbb{R}$ satisfying Equations (1.2) and (1.3). The *tangent bundle* of the smooth manifold M is the disjoint union of the tangent spaces at each point:

$$TM = \bigsqcup_{x \in M} T_x M$$

It can be equipped with a natural topology and a natural smooth structure, making it into a rank n smooth vector bundle over M (see Lemma 4.1 in [Lee, 2003]). Similarly, the *cotangent bundle* is the disjoint union of the cotangent spaces at each point, i.e. the spaces dual to the tangent spaces: $T_x^* M = (T_x M)^*$. It can be showned that it is a smooth vector bundle of rank n (see Proposition 6.5 in [Lee, 2003]).

A *local section* of a vector bundle E over an open set $U \subset M$ is a continuous map $\sigma : U \rightarrow E$ such that:

$$\pi \circ \sigma = \text{id}_U$$

A *global section* is a local section defined over the whole manifold, i.e. such that $U = M$. When E is a smooth vector bundle and σ is a smooth map, we say it is a *smooth section*. Vector fields and differential 1-forms are smooth sections of the vector bundles TM and T^*M , respectively.




Figure 9: A (smooth) vector bundle is locally trivial, i.e. in the neighborhood of every point, there is an open set U , over which the pre-image $\pi^{-1}(U)$ is homeomorphic (resp. diffeomorphic) to $U \times \mathbb{R}^k$.

Some are defined only locally, while other are defined globally. The space of vector fields on an open set U is noted $\mathfrak{X}(U)$ while the space of differential 1-forms on the same open set is denoted $\Omega^1(U)$. A vector bundle always admits a smooth global section: the *zero section*, that has the particularity that it sends every point $x \in M$ to the zero vector in the fiber E_x . A set of k local sections $\sigma_1, \dots, \sigma_k$ of E over U is called a *local frame* of E over U if for every $x \in U$, the vectors $\sigma_1(x), \dots, \sigma_k(x)$ form a basis of the fiber E_x . It is called a *global frame* if $U = M$, and it is called *smooth* if the sections σ_i are smooth sections of the smooth vector bundle E .

Proposition 2.17. *A smooth vector bundle is trivial if and only if it admits a smooth global frame.*

Remark 2.18. Unless explicitly said, in the following we will always assume that vector bundles and their sections are smooth.

The space of smooth local sections of E over U is denoted $\Gamma_U(E)$ or $\Gamma(U, E)$; it is an infinite dimensional \mathbb{R} -vector space but a $\mathcal{C}^\infty(U)$ -module. If there exists a smooth local frame on U – this occurs U is an open set trivializing E , i.e. satisfying the second item of 2.16 – then one observes that the frame plays the role of independent generators of $\Gamma_U(E)$, with respect to the action of $\mathcal{C}^\infty(U)$. One can always find such a frame in the neighborhood of every point, turning the assignment $U \rightarrow \Gamma(U, E)$ in what is called a *locally free and finitely generated sheaf* (it is

actually what is called a \mathcal{C}^∞ -module, because $\Gamma(U, E)$ is a $\mathcal{C}^\infty(U)$ -module for every U). Pushing the idea further, $\mathcal{C}^\infty(U)$ can be seen as the space of local sections over U of the trivial bundle $M \times \mathbb{R}$. In the same manner that a smooth manifold can be defined by its algebra of functions, a smooth vector bundle can be defined through its space of sections. This fact is a central tenet of the general duality between geometry and algebra. The category of real vector bundles on M is equivalent to the category of locally free and finitely generated sheaves of \mathcal{C}^∞ -modules on M . This is the well-known Serre-Swan theorem which, in modern language, can be expressed as:

Theorem 2.19. Serre-Swan *There is an equivalence of categories between smooth vector bundles of finite rank over a smooth manifold M and finitely generated projective (equivalently: locally free) \mathcal{C}^∞ -modules over M .*

$$\begin{array}{ccc} \text{Geometry} & & \text{Algebra} \\ E & \longleftrightarrow & \Gamma(-, E) \end{array}$$

Proof. It is explained in Chapter 11 of [Nestruev, 2003]. □

Smooth sections of the vector bundle $\bigwedge^m T^*M = \bigsqcup_{x \in M} \bigwedge^m T_x^*M$ are called differential m -forms. They can be either locally defined or globally defined. The de Rham differential acts on these differential forms via Equation (1.31), and it induces the same notion of closedness and exactness of differential forms. For any open set U , the m -th de Rham cohomology group is:

$$H_{dR}^m(U) = \frac{\text{Ker}(d : \Omega^m(U) \longrightarrow \Omega^{m+1}(U))}{\text{Im}(d : \Omega^{m-1}(U) \longrightarrow \Omega^m(U))}$$

where we understand that $\Omega^{-1}(U) = \Omega^{n+1}(U) = 0$. Since a smooth manifold is locally Euclidean, it means that in the neighborhood of every point, the cohomology groups are trivial except at level 0 (see Proposition 1.55), because a small enough open set is homeomorphic to \mathbb{R}^n . However, globally, the de Rham cohomology of a smooth manifold has no reason to be trivial. On the contrary, it is often not trivial because it contains information on the topological structure of the manifold, as the following examples show:

Example 2.20. The de Rham cohomology of the n -sphere \mathbb{S}^n satisfies:

$$H_{dR}^i(\mathbb{S}^n) \simeq \begin{cases} \mathbb{R} & \text{if } i = 0 \text{ or } i = n \\ 0 & \text{otherwise} \end{cases}$$

Example 2.21. The de Rham cohomology of the n -torus \mathbb{T}^n satisfies:

$$H_{dR}^i(\mathbb{T}^n) \simeq \mathbb{R}^{\binom{n}{i}}$$

In order to define local frames of the tangent and cotangent bundles, one needs to introduce the notion of pushforwards, and pull backs. First, let us define the following important notion:

Definition 2.22. *A morphism of vector bundles, or bundle map, between smooth vector bundles E (over M) and E' (over N) is a pair (ψ, ϕ) of smooth maps $\psi : M \longrightarrow N$ and $\phi : E \longrightarrow E'$, making the following square commutative:*

$$\begin{array}{ccc} E & \xrightarrow{\phi} & E' \\ \pi \downarrow & & \downarrow \pi' \\ M & \xrightarrow{\psi} & N \end{array}$$

and such that the restriction to the fibers $\phi_p : E_p \longrightarrow E'_{\psi(p)}$ is a linear morphism of vector spaces.

When $N = M$ and $\psi = \text{id}_M$, the above diagram reduces to a triangle:

$$\begin{array}{ccc} E & \xrightarrow{\phi} & E' \\ \pi \searrow & & \swarrow \pi' \\ & M & \end{array}$$

Since both E and E' are smooth vector bundles over M , for every smooth section σ the composite $\phi \circ \sigma$ defines a smooth section of E' . Then, fiberwise linearity of ϕ implies that, for every smooth section σ of E , and every function $f \in C^\infty(M)$, one has:

$$\phi(f(x)\sigma(x)) = f(x)\phi(\sigma(x))$$

Forgetting about the point x , this equation reads: $\phi \circ (f\sigma) = f(\phi \circ \sigma)$. Thus, vector bundle morphisms over the same base manifold are morphisms of the corresponding sheaves of sections that are C^∞ -linear. This is actually an alternative characterization of vector bundle morphisms over a smooth manifold. This is a consequence of the Serre-Swan theorem:

Proposition 2.23. *Let E and E' be smooth vector bundles over a smooth manifold M . Then a map of sheaves $\Phi : \Gamma(-, E) \longrightarrow \Gamma(-, E')$ is linear over $C^\infty(U)$ for every open set U if and only if there exists a smooth bundle map $\phi : E \longrightarrow E'$ over M such that $\Phi(\sigma) = \phi \circ \sigma$ for all smooth section σ .*

Let us provide an example of such a vector bundle morphism, that will become central in the following parts of the course:

Definition 2.24. *Let M be a smooth manifold. A Lie algebroid over M is a smooth vector bundle A , together with:*

1. *a Lie algebra structure $[\cdot, \cdot]_A : \Gamma(A) \otimes \Gamma(A) \longrightarrow \Gamma(A)$ on the space of sections,*
2. *and a vector bundle morphism $\rho : A \longrightarrow TM$ called the anchor,*

such that the following Leibniz rule holds:

$$[a, fb]_A = f[a, b]_A + \rho(a)(f)b \tag{2.1}$$

for every $a, b \in \Gamma(A)$ and $f \in C^\infty(M)$.

A Lie algebroid is a generalization of the tangent bundle, since Equation (4.86) is resembling Equation (1.12). Indeed, the tangent bundle is a particular case of a Lie algebroid, where the anchor is the identity map. Lie algebroids also generalize Lie algebras since a Lie algebra is a Lie algebroid over a point. As Lie algebras are infinitesimal counterparts of Lie groups, Lie algebroids are infinitesimal counterparts of *Lie groupoids*. These objects are widely used in mathematical physics nowadays.

Example 2.25. The space of endomorphisms of \mathbb{R}^n is denoted $\text{End}(\mathbb{R}^n)$ or $\mathfrak{gl}_n(\mathbb{R})$. By [Example 1.15](#), it is a finite dimensional Lie algebra, with respect to the commutator of endomorphisms $[M, N] = M \circ N - N \circ M$. This Lie algebra additionally defines an infinitesimal Lie algebra action on \mathbb{R}^n via the following Lie algebra homomorphism:

$$\begin{aligned} \bar{\rho} : \mathfrak{gl}_n(\mathbb{R}) &\longrightarrow \mathfrak{X}(\mathbb{R}^n) \\ M &\longrightarrow X_M : (x, f) \longmapsto \left. \frac{d}{dt} \right|_{t=0} f(x \cdot \exp(tM)) \end{aligned}$$

where, on the right-hand side, the group element acts from the right. On the basis $(E_{i,j})_{1 \leq i,j \leq n}$ of $\mathfrak{gl}_n(\mathbb{R})$ this homomorphism then reads at the point x :

$$\bar{\rho}(E_{i,j})_x = X_{E_{i,j},x} = x^i \frac{\partial}{\partial x^j}$$

where the x^i are the coordinates of the point x . These data are sufficient to define a Lie algebroid over \mathbb{R}^n via the following data: $A = \mathbb{R}^n \times \mathfrak{gl}_n(\mathbb{R})$ (it is a trivial vector bundle); $[\cdot, \cdot]_A$ is defined on the constant sections as the bracket on $\mathfrak{gl}_n(\mathbb{R})$ and then it is generalized to every smooth sections by the Leibniz rule ([4.86](#)); the anchor map is defined on the frame of constant sections $(E_{i,j})_{1 \leq i,j \leq n}$ of A by:

$$\rho(E_{i,j}) = x^i \frac{\partial}{\partial x^j}$$

Then the infinitesimal action of $\mathfrak{gl}_n(\mathbb{R})$ on \mathbb{R}^n straightforwardly translates in the data contained in a Lie algebroid. More generally, the action of a Lie algebra \mathfrak{g} on a manifold M can be encoded in what is called an *action Lie algebroid* $M \times \mathfrak{g}$.

Remark 2.26. One could have defined the infinitesimal action of $\mathfrak{gl}_n(\mathbb{R})$ on \mathbb{R}^n as a left action, but in that case we need to add a minus sign to have a Lie algebra homomorphism:

$$\begin{aligned} \mathfrak{gl}_n(\mathbb{R}) &\longrightarrow \mathfrak{X}(\mathbb{R}^n) \\ M &\longrightarrow X_M : (x, f) \longmapsto - \left. \frac{d}{dt} \right|_{t=0} f(\exp(tM) \cdot x) \end{aligned} \quad (2.2)$$

and the basis vectors $E_{i,j}$ are sent to $-x^j \frac{\partial}{\partial x^i}$, which is not very practical. If the minus sign had not been present, we would have a Lie algebra *anti-homomorphism* $\mathfrak{g} \longrightarrow \mathfrak{X}(M)$. The choice of a minus sign or, more conveniently, a right action, comes from the following facts (that we summarize very sketchily): to any smooth manifold M , one can associate its set of diffeomorphisms, denoted $\text{Diff}(M)$. This space can be equipped with an infinite dimensional Lie group structure whose local charts are modeled over the infinite vector space $\mathfrak{X}(M)$.

Then, the left invariant vector fields over this Lie group form a Lie algebra $\mathfrak{diff}(M)$, in bijection with the space of vector fields on M . However, the choice of Lie bracket on $\mathfrak{X}(M)$, as defined in [Equation \(1.10\)](#), corresponds to *minus* the Lie bracket on $\mathfrak{diff}(M)$, and we write $\mathfrak{X}(M) \simeq \mathfrak{diff}(M)^{op}$. This can be explained by the fact that the diffeomorphisms act on M from the left, and thus the induced linear map between the Lie algebras of $\text{Diff}(M)$ and $\mathfrak{X}(M)$ is an anti-homomorphism (because the vector field associated to a given diffeomorphism is obtained without involving the minus sign appearing in [Equation \(2.2\)](#)). More generally, a left action of a Lie group G on the manifold M is equivalent to a group homomorphism $G \longrightarrow \text{Diff}(M)$. This group homomorphism induces in turn a Lie algebra homomorphism $\mathfrak{g} \longrightarrow \mathfrak{diff}(M)$, and thus a Lie algebra anti-homomorphism from \mathfrak{g} to $\mathfrak{X}(M)$. However, a right-action of a Lie group G on a manifold M is equivalent to a Lie group homomorphism $G \longrightarrow \text{Diff}(M)^{op}$, where $\text{Diff}(M)^{op}$ is the Lie group modeled on $\text{Diff}(M)$ but with multiplication from the right. Then in that case, the Lie algebra homomorphism $\mathfrak{g} \longrightarrow \mathfrak{diff}(M)^{op}$ is equivalent to a Lie algebra homomorphism $\mathfrak{g} \longrightarrow \mathfrak{X}(M)$. For more details on these questions see [Section 3.3](#) of these [lecture notes](#).

Exercise 2.27. By using the Jacobi identity on $\Gamma(A)$ and Equation (4.86), show that the anchor map is a *Lie algebra homomorphism* from $\Gamma(A)$ to $\mathfrak{X}(M)$. That is to say, it satisfies the following equation:

$$\rho([a, b]_A) = [\rho(a), \rho(b)] \quad (2.3)$$

for every smooth sections $a, b \in \Gamma(A)$.

Here is another important example of a vector bundle morphism:

Definition 2.28. Let M, N be smooth manifolds. For every smooth map $F : M \rightarrow N$ we associate a vector bundle morphism $F_* : TM \rightarrow TN$ called the *pushforward*, defined on each fiber $T_x M$ as:

$$F_*(X_x)(f) = X_x(f \circ F)$$

for every $f \in C^\infty(N)$ and $X_x \in T_x M$.

The pushforward is a vector bundle morphism sending tangent vectors on M to tangent vectors on N :

$$\begin{array}{ccc} TM & \xrightarrow{F_*} & TN \\ \pi \downarrow & & \downarrow \pi' \\ M & \xrightarrow{F} & N \end{array}$$

Given a point $x \in M$ (resp. $y \in N$) and a trivializing neighborhood (U, φ) centered at x (resp. (V, ψ) centered at y), then the matrix of the linear morphism $F_* : T_x M \rightarrow T_{F(x)} N$ at x is the Jacobian of the smooth map $\psi \circ F \circ \varphi^{-1} : \tilde{U} \rightarrow \tilde{V}$ at $\varphi(x)$, where $\tilde{U} = \varphi(U)$ is an open subset of \mathbb{R}^n centered at 0 (and respectively for \tilde{V}):

$$F_*|_x = \left(\frac{\partial(\psi \circ F \circ \varphi^{-1})^j}{\partial x^i}(\varphi(x)) \right)_{1 \leq i, j \leq n} \quad (2.4)$$

In the above formula, the coordinates x^i in the denominator denote the standard coordinates on \tilde{U} . Notice that the numerator $(\psi \circ F \circ \varphi^{-1})^j$ can alternatively be written $\psi^j \circ F \circ \varphi^{-1}$, where ψ^j is a smooth function on V and denotes the j -th component of ψ with respect to the standard coordinates on \tilde{V} .

The pushforward F_* is then the best linear approximation of F at the point x . The rank of the Jacobian matrix at x characterizes this smooth map at this point and is called the *rank* of F at x . If the rank of F is constant for every point of the smooth manifold M then we say that F has *constant rank* and denote it by $\text{rk}(F)$. We have the following conventions:

1. if $\text{rk}(F) = \dim(M)$ at every point (i.e. F_* is injective everywhere), then F is called an *immersion*;
2. if $\text{rk}(F) = \dim(N)$ at every point (i.e. F_* is surjective everywhere), then F is called a *submersion*.

In both cases, these properties are partly independent from the fact that F being injective, surjective or bijective. For instance, when $\dim(M) = \dim(N)$, F is a (local) diffeomorphism if and only if it is an immersion or a submersion. However, F needs not be a global diffeomorphism: for that it should be either injective or surjective.

Remark 2.29. The pushforward admits several other notations: dF because it is the differential of the map F (so that when $N = \mathbb{R}$, we retrieve the usual differential of functions), TF to symbolize that it is a map between tangent spaces, etc.

Lemma 2.30. *Given two smooth functions $F : M \rightarrow N$ and $G : N \rightarrow P$, the pushforward of the composite $G \circ F : M \rightarrow P$ preserves the order:*

$$(G \circ F)_* = G_* \circ F_*$$

Be aware that although tangent vectors always behave well under pushforwards, it may not be the case for vector fields, i.e. sections of the tangent bundle. This phenomenon actually exists for every vector bundle morphism, so we will study this problem in the general setting. Let E (resp. E') be a smooth vector bundle over the smooth manifold M (resp. N), and let (ψ, ϕ) be a vector bundle morphism from E to E' . Let σ be a smooth section of E , then under the action of ϕ it becomes a map $\sigma^{(\psi, \phi)} : \text{Im}(\psi) \rightarrow E'$ defined on the subset $\text{Im}(\psi) \subset N$ by:

$$\sigma^{(\psi, \phi)}(\psi(x)) = \phi(\sigma(x))$$

There may be several obstructions to the fact that $\sigma^{(\psi, \phi)}$ forms a smooth section of E' . This can be seen in several situations: 1) local sections should be defined on open sets, but if the smooth map ψ is not open (i.e. if $\psi(U)$ is not necessarily open while U is open) then $\text{Im}(\psi)$ may not be even open in N , so that the map $\sigma^{(\psi, \phi)}$ could not be qualified as a local section of E' ; 2) if ψ is not injective then ϕ can send two conflicting informations to the same point of N : take $x, y \in M$ such that $z = \psi(x) = \psi(y)$, but then for any choice of smooth section σ , how would be defined $\sigma^{(\psi, \phi)}(z)$? As $\phi(\sigma(x))$ or as $\phi(\sigma(y))$?

These problems can be explicitly solved if one introduces the notion of *pullback bundle*, that would be introduced as an intermediary bundle between E and E' . Given a smooth map $\psi : M \rightarrow N$, and a vector bundle E' over N , one defines the *pullback bundle of E' along ψ* , denoted $\psi^!E'$, as the vector bundle over M such that the fiber over the point $x \in M$ is $(\psi^!E')_x = E'_{\psi(x)}$. Thus, as a set, the pullback bundle is the disjoint union $\psi^!E' = \bigsqcup_{x \in M} E'_{\psi(x)}$ and the projection map is denoted $\psi^!\pi' : E'_{\psi(x)} \mapsto x$. Notice that the fact that we have a disjoint union (and not a mere union) is crucial so that the fibers associated to the pre-image of the same point stay disjoint in $\psi^!E'$. Under this convention, the vector bundle morphism (ψ, ϕ) induces a vector bundle morphism (id_M, ϕ) covering the identity of M :

$$\begin{array}{ccc} E & \xrightarrow{\phi} & \psi^!E' \\ & \searrow \pi & \swarrow \psi^!\pi' \\ & & M \end{array}$$

In that context, any local smooth section σ of E induces a local smooth section $\sigma^{(\text{id}_M, \phi)}$ of $\psi^!E'$. Indeed, since $\psi^!E'$ is a vector bundle over M , if σ is defined over an open set U then $\sigma^{(\text{id}_M, \phi)}$ stays defined over the same open set U . Moreover, the possible lack of injectivity of ψ is now solved: the images of two different fibers of E through ϕ are sent to different fibers of $\psi^!E'$ so they cannot be confounded. Thus, even though $\psi(x) = \psi(y)$, the element $\sigma^{(\text{id}_M, \phi)}(x)$ is a vector of the fiber of $\psi^!E'$ over x , while $\sigma^{(\text{id}_M, \phi)}(y) \in (\psi^!E')_y$. To conclude, the vector bundle morphism $(\text{id}_M, \phi) : E \rightarrow \psi^!E'$ sends smooth sections to smooth sections.

Now, notice that any smooth section τ of E' over some open set $U \subset N$ defines a map $\psi^!\tau : \psi^{-1}(U) \rightarrow \psi^!E'$ by the following identity:

$$(\psi^!\tau)_x = \tau_{\psi(x)}$$

This assignment is well defined, and it is easy to see that it is additionally smooth, hence $\psi^! \tau$ is a smooth section of $\psi^! E'$ over $\psi^{-1}(U)$. Then, we say that a smooth section $\sigma \in \Gamma(E)$ over some open set $U \subset M$ and a smooth section $\tau \in \Gamma(E')$ over some open set $V \subset N$ containing $\psi(U)$ are (ψ, ϕ) -related if we have the following identity over U :

$$\sigma^{(\text{id}_M, \phi)} = \psi^! \tau$$

In that case, we can consider that the image of the smooth section σ through (ψ, ϕ) is τ . Obviously, if ψ is not surjective, σ can be related to many sections τ (that could for instance differ outside $\text{Im}(\psi)$). Moreover, not every smooth section of E is related to a smooth section of E' .

Example 2.31. Let $M = N = \mathbb{R}$ and let $E = E' = \mathbb{R}^2$. Let $\psi(x) = x^2$ and let $\phi(x, y) = (x, xy)$ (the latter is indeed linear on the fibers). Be aware that although $M = N$ and $E = E'$, the fact that ψ is not the identity implies that not all smooth sections of E are (ψ, ϕ) related to smooth sections of E' . Let $\sigma : x \mapsto (x, \sin(x))$ a smooth section of E ; determine what is the section $\sigma^{(\text{id}_M, \phi)} \in \Gamma(\psi^! E')$. Then, find a global smooth section τ of E' which is (ψ, ϕ) -related to σ . Find a smooth section σ' of E such that there exist *no* global smooth section of E' that would be (ψ, ϕ) -related to σ' .

To summarize we have the following situation (this diagram should not be understood as a commutative diagram, but as a metaphor, even though the square on the left is commutative):

$$\begin{array}{ccccc} E & \xrightarrow{\phi} & \psi^! E' & \xleftarrow{\psi^!} & E' \\ \pi \downarrow & & \psi^! \pi' \downarrow & & \downarrow \pi' \\ M & \xrightarrow{\text{id}_M} & M & \xrightarrow{\psi} & N \end{array}$$

Using this construction, we understand that, given a smooth map $F : M \rightarrow N$, a vector field X on M is F -related to a vector field Y on N , if:

$$F_* X = F^! Y$$

The notion of pullback bundle, as the name indicates, allows to make sense of so-called *pullbacks*:

Definition 2.32. Let M, N be smooth manifolds. For every smooth map $F : M \rightarrow N$ we associate the pullback $F^* : F^! T^* N \rightarrow T^* M$, defined on each fiber $T_{F(x)}^* N$ as:

$$F^*(\xi_{F(x)})(X_x) = \xi_{F(x)}(F_*(X_x)) \tag{2.5}$$

for every $\xi_{F(x)} \in T_{F(x)}^* N$ and $X_x \in T_x M$.

The pullback is a vector bundle morphism covering the identity of M :

$$\begin{array}{ccc} F^! T^* N & \xrightarrow{F^*} & T^* M \\ & \searrow F^! \pi' & \swarrow \pi \\ & M & \end{array}$$

Differential 1-forms on N can be pullbacked on M via Equation (2.5) but, contrary to vector fields that do not behave well under pushforwards, differential forms actually behave very well under pullbacks. For every covector field $\xi \in \Omega^1(N)$, the pullback of ξ is the unique section $F^*\xi$ of T^*M defined at x as in Equation (2.5):

$$(F^*\xi)_x = F^*(\xi_{F(x)})$$

Note that there is no ambiguity in the definition of $F^*(\xi)$, contrary to the case of the pushforward of vector fields. This section is smooth because the function $(F^*\xi)(X) : x \mapsto \xi_{F(x)}(F_*(X_x))$ is a smooth function of x (it can be seen from the fact that $x \mapsto \xi_{F(x)}$ is a smooth section of $F^!T^*N$, while $x \mapsto F_*X_x$ is a smooth section of $F^!TN$), thus it satisfies criterion 2. of Scholie 1.22.

Thus, the pullback can be extended to a smooth map $F^* : \Omega^1(N) \longrightarrow \Omega^1(M)$. We can also extend F^* to smooth functions, for if $f \in \mathcal{C}^\infty(N)$, we define, for every $x \in M$:

$$F^*(f)(x) = f(F(x))$$

More generally, for every differential m -form η on N ($m \geq 1$), one defines the pullback of η to M from its action on m vector fields $X_1, \dots, X_m \in \mathfrak{X}(M)$:

$$F^*(\eta)(X_1, \dots, X_m) = F^!\eta(F_*X_1, \dots, F_*X_m)$$

where $F^!\eta \in \Gamma(F^! \wedge^m T^*N)$ is the pullback section of $F^! \wedge^m T^*N$ associated to η , i.e. the smooth map associating to every point $x \in M$ the covector $\eta_{F(x)}$. Using this result, one can extend the pullback as a graded commutative algebra morphism $F^* : \Omega^\bullet(N) \longrightarrow \Omega^\bullet(M)$ from the following identity:

$$F^*(\eta \wedge \mu) = F^*(\eta) \wedge F^*(\mu) \tag{2.6}$$

for any m -form η and p -form μ on N . For a proof of this statement see Lemma 12.10 in [Lee, 2003]. Then, the pullback somehow defines a dual version of a smooth map:

Geometry		Algebra
M	\longleftarrow	$\mathcal{C}^\infty(M)$
$F : M \longrightarrow N$	\longleftarrow	$F^* : \Omega^\bullet(N) \longrightarrow \Omega^\bullet(M)$

In this correspondence the pullback is actually characterized by the following algebraic property:

Proposition 2.33. *The pullback $F^* : \Omega^\bullet(N) \longrightarrow \Omega^\bullet(M)$ of the smooth map $F : M \longrightarrow N$ is a morphism of differential graded commutative algebras from $(\Omega^\bullet(N), d_N)$ to $(\Omega^\bullet(M), d_M)$. In particular, it commutes with the respective de Rham differentials d_M on M and d_N on N :*

$$d_M \circ F^* = F^* \circ d_N$$

Proof. The fact that F^* is a morphism of graded commutative algebra is transparent in Equation (2.6). For $m = 0$, let $f \in \mathcal{C}^\infty(N)$ and let X be a vector field on M . Then, one has:

$$F^*(d_N f)(X) = F^!d_N f(F_*X) = F_*X(f) = X(f \circ F) = X(F^*f) = d_M F^*(f)(X)$$

where the third term is an explicitation of the second, as the action of the section $F_*X \in \Gamma(F^!TN)$ on f is understood to be the expected one: $x \mapsto X_x(f \circ F(x))$. Now let $m \geq 1$, let $\eta \in \Omega^m(N)$ be a differential m -form on N , and let X_1, \dots, X_{m+1} be m vector fields on

M . One can easily check that on their respective pullback bundles, $F^!d_N\eta = d_NF^!\eta$ and $[F_*X_i, F_*X_j] = F_*[X_i, X_j]$. From this, we deduce:

$$\begin{aligned}
(F^*(d_N\eta))(X_1, \dots, X_{m+1}) &= F^!d_N\eta(F_*X_1, \dots, F_*X_{m+1}) = d_NF^!\eta(F_*X_1, \dots, F_*X_{m+1}) \\
&= \sum_{i=1}^{m+1} (-1)^{i-1} F_*X_i(F^!\eta(F_*X_1, \dots, \widehat{F_*X_i}, \dots, F_*X_{m+1})) \\
&\quad + \sum_{1 \leq i < j \leq m+1} (-1)^{i+j-1} F^!\eta(\underbrace{[F_*X_i, F_*X_j]}_{= F_*[X_i, X_j]}, F_*X_1, \dots, \widehat{F_*X_i}, \dots, \widehat{F_*X_j}, \dots, F_*X_{m+1}) \\
&= \sum_{i=1}^{m+1} (-1)^{i-1} X_i(F^*(\eta)(X_1, \dots, \widehat{X_i}, \dots, X_{m+1})) \\
&\quad + \sum_{1 \leq i < j \leq m+1} (-1)^{i+j-1} F^*(\eta)([X_i, X_j], X_1, \dots, \widehat{X_i}, \dots, \widehat{X_j}, \dots, X_{m+1}) \\
&= d_MF^*(\eta)
\end{aligned}$$

□

Since pullback goes the other way around compared to F , contrary to pushforwards, composition of pullbacks is not the pullback of the composite maps:

Lemma 2.34. *Given two smooth functions $F : M \rightarrow N$ and $G : N \rightarrow P$, the pullback of the composite $G \circ F : M \rightarrow P$ reverts the order:*

$$(G \circ F)^* = F^* \circ G^*$$

Remark 2.35. The correspondence between geometry and algebra can be further exploited to describe Lie algebroid morphisms. Without entering into much details, a morphism of Lie algebroids $\phi : A \rightarrow B$ is a vector bundle morphism that is additionally a Lie algebra morphism on sections, and which is compatible with the anchor map. This complicated condition can be equivalently stated as the following: a Lie algebroid morphism is a morphism of differential commutative graded algebra $\Phi : (\Omega^\bullet(B), d_B) \rightarrow (\Omega^\bullet(A), d_A)$, where $(\Omega^\bullet(A), d_A)$ is the so-called *Lie algebroid cohomology*. This one-to-one correspondence was originally found by Vaintrob [Vaintrob, 1997], and is still valid for higher Lie algebroids.

Pushforwards and pullbacks allow to define smooth local frames on the tangent and cotangent bundle. Let $x \in M$ and let (U, φ) be a trivializing chart of the tangent bundle (and then, by duality of the fiber, of the cotangent bundle as well) centered at x . We denote by x^1, \dots, x^n the standard coordinates on \tilde{U} centered at 0 (because $\varphi(x) = 0$), and by abuse of notation they also denote the composite function $x^i \circ \varphi$. Then one can define:

1. a local smooth frame of TM over U from the constant vector fields $\frac{\partial}{\partial x^i}$ on $\tilde{U} = \varphi(U)$, via the push-forward of $\varphi^{-1} : \tilde{U} \rightarrow U$. This time the pushforward is well defined because $\varphi : U \rightarrow \tilde{U}$ is a diffeomorphism. For brevity, we denote the induced local smooth frame on U by the same notation $\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^n}$ and we call it the *coordinate frame*;
2. a local smooth frame of T^*M over U from the constant covector fields dx^i on $\tilde{U} = \varphi(U)$ via the pull back of φ . We denote this local smooth frame on U by the same notation dx^1, \dots, dx^n and we call it the *coordinate coframe*.

Both frames are well-defined because they are constant sections. In order to differentiate the frame on $U \subset M$ and the one on $\tilde{U} \subset \mathbb{R}^n$, we write $\frac{\partial}{\partial x^i} \Big|_y$ (resp. $dx^i|_y$) to indicate the former,

and $\frac{\partial}{\partial x^i} \Big|_{\varphi(y)}$ (resp. $dx^i|_{\varphi(y)}$) to indicate the latter, for every $y \in U$. A manifold that admits a smooth global frame for the tangent bundle is said *parallelizable*. The only spheres that are parallelizable are \mathbb{S}^1 , \mathbb{S}^3 and \mathbb{S}^7 .

Example 2.36. For every $1 \leq m \leq n$, the exterior algebra of the cotangent space at each point defines a smooth vector bundle:

$$\wedge^m T^* \mathbb{R}^n = \bigsqcup_{x \in \mathbb{R}^n} \wedge^m T_x^* \mathbb{R}^n$$

Then a smooth local frame consists of the sections $dx^{i_1} \wedge \dots \wedge dx^{i_m}$ for $1 \leq i_1 < \dots < i_m \leq n$. It is not a global frame because the coordinates functions x^i are only defined locally.

Now let us understand how the coordinate functions of vector fields and of differential forms transform under a change of local coordinates. Assume that there exists another compatible chart (V, ψ) centered at x so that $\tilde{V} = \psi(V)$ and x^1, \dots, x^n are the standard coordinates on \tilde{V} centered at 0. Then, under the change of coordinates $\psi \circ \varphi^{-1} : \tilde{U} \rightarrow \tilde{V}$, the constant sections $\frac{\partial}{\partial x^i}$ transform as:

$$\frac{\partial}{\partial x^i} \Big|_y = (\varphi^{-1})_* \frac{\partial}{\partial x^i} \Big|_{\varphi(y)} \quad (2.7)$$

$$\begin{aligned} &= (\psi^{-1})_* \circ (\psi \circ \varphi^{-1})_* \frac{\partial}{\partial x^i} \Big|_{\varphi(y)} \\ &= (\psi^{-1})_* \left(\frac{\partial(\psi \circ \varphi^{-1})^j}{\partial x^i}(\varphi(y)) \frac{\partial}{\partial x'^j} \Big|_{\psi(y)} \right) \\ &= \frac{\partial(\psi^j \circ \varphi^{-1})}{\partial x^i}(\varphi(y)) \frac{\partial}{\partial x'^j} \Big|_y \end{aligned} \quad (2.8)$$

for every $y \in U \cap V$. We pass from the first line to the second line by using Lemma 2.30, and from the second line to the third by Equation (2.4). Then the push-forward $(\psi^{-1})_*$ is linear so that we obtain the fourth line.

Remark 2.37. Usually the term $\frac{\partial(\psi^j \circ \varphi^{-1})}{\partial x^i}(\varphi(y))$ is denoted $\frac{\partial x'^j}{\partial x^i}(y)$, because it is transparent and for practical purposes. We will pick up this convention from then on.

A vector field $X \in \mathfrak{X}(U \cap V)$ decomposes as $X^i \frac{\partial}{\partial x^i}$ with respect to the coordinate functions x^i of φ , and $X'^j \frac{\partial}{\partial x'^j}$ with respect to the coordinate functions x'^i of ψ . Then Equations (2.7)-(2.8) show that:

$$X'_y{}^j = \frac{\partial x'^j}{\partial x^i}(y) X_y^i$$

We observe that under a change of coordinates $x^i \mapsto x'^i$, the coordinate functions of the vector field transform in the opposite way than the constant sections $\frac{\partial}{\partial x^i}$:

$$\begin{array}{ccc} \frac{\partial}{\partial x^i} \Big|_y & \longrightarrow & \frac{\partial}{\partial x'^i} \Big|_y = \frac{\partial x^j}{\partial x'^i}(y) \frac{\partial}{\partial x^j} \Big|_y \\ X_y^i & \longrightarrow & X_y'^i = \frac{\partial x'^j}{\partial x^i}(y) X_y^i \end{array}$$

The first line has been obtained from Equations (2.7)-(2.8) by inverting the Jacobian matrix $\frac{\partial x'^j}{\partial x^i}$. Since the coordinate functions of vector fields transform in the opposite way than the way

in which the canonical frame of the tangent bundle transforms, we say that these coordinates are *contravariant*. Changes of coordinates impact also the way differential forms transform, since for any covector field $\xi = \xi_i dx^i = \xi'_j dx'^j$, by using Equations (2.7)-(2.8), one has:

$$\xi_{y,i} = \xi_y \left(\frac{\partial}{\partial x^i} \Big|_y \right) = \frac{\partial x'^j}{\partial x^i}(y) \xi_y \left(\frac{\partial}{\partial x'^j} \Big|_y \right) = \frac{\partial x'^j}{\partial x^i}(y) \xi'_{y,j}$$

Here, we consider that the coordinates x^i and x^j are those on $U \cap V$. Thus, we observe that the coordinate functions of differential 1-forms transform in the same way as the constant sections $\frac{\partial}{\partial x^i}$:

$$\begin{array}{ccc} \frac{\partial}{\partial x^i} \Big|_y & \longrightarrow & \frac{\partial}{\partial x'^i} \Big|_y = \frac{\partial x^j}{\partial x'^i}(y) \frac{\partial}{\partial x^j} \Big|_y \\ \xi_{y,i} & \longrightarrow & \xi'_{y,i} = \frac{\partial x^j}{\partial x'^i}(y) \xi_{y,j} \end{array}$$

Since the coordinate functions of differential forms transform in the same way as the way in which the coordinate frame of the tangent bundle transforms, we say that these coordinates are *covariant*. In general the position of the indices indicates when it is a covariant (at the bottom) or a contravariant (at the top) coordinate. The names ‘contravariant’ and ‘covariant’ come from the fact that the pushforward functor, assigning to any smooth manifold its tangent bundle and to any smooth function its pushforward, is a covariant functor, while the pullback functor, assigning to any smooth manifold its algebra of functions and to any smooth function between manifolds its pullback, is a contravariant functor.

2.3 Submanifolds in differential geometry

The notion of pullback and pushforward allows to define various kinds of subspaces in a smooth manifold, that can be additionally equipped with a distinguished smooth structure that turn them into *submanifolds*. Since it is a very subtle topic, I strongly advise the reader to refer to [Lee, 2003] and to [Lee, 2009] to get a much more clear understanding of the notions discussed in the present section. Also, I refer to [Karshon et al., 2022] in order to deepen one’s understanding of the differences between the various exposed notions and their relationship with diffeological manifolds. There are three main kinds of submanifold objects:

$$\{\text{embedded submanifolds}\} \subset \{\text{weakly embedded submanifolds}\} \subset \{\text{immersed submanifolds}\}$$

An *immersed submanifold* of a smooth manifold M is a subset S , equipped with a smooth structure (i.e. a topology composed of smoothly compatible charts) such that the inclusion $\iota : S \longrightarrow M$ is a smooth map (with respect to the respective smooth structures on S and M) and an immersion. It does not mean that the topology on S is the subspace topology, and in general it will not be! There may exist various non-diffeomorphic smooth structures on the subset S such that the inclusion ι is an immersion. A famous example of an immersed manifold is the figure eight:

Example 2.38. Let $\gamma :]-\pi, \pi[\longrightarrow \mathbb{R}^2$ be the smooth map defined as:

$$\gamma(t) = (\sin(2t), \sin(t))$$

The image of γ , denoted S , is the locus of points (x, y) defined by $x^2 = 4y^2(1 - y^2)$. This subset can be equipped with a topology of open sets defined as follows: a subset $U \subset \text{Im}(\gamma)$ is open if

and only if $\gamma^{-1}(U)$ is open in the topology of N . This implies in particular that any subset of the form $(\sin(2t), \sin(t))$ for $t \in]-\epsilon, \epsilon[$ is an open set of S . The map γ^{-1} then turn these open sets into smooth open charts, that are smoothly compatible by construction. Then, S is a smooth manifold, but its smooth structure does not descend from the smooth structure on M for the following reason: in the subspace topology, a neighborhood of 0 in S has the shape of a cross, and is not homeomorphic to any region of euclidean space, while in the manifold topology, there exist neighborhoods of the origin that are homeomorphic to an open one-dimensional segment. The inclusion map $\iota : S \rightarrow M$ being an immersion since $\dot{\gamma}(t)$ never vanishes, the subset S equipped with its smooth manifold structure is an immersed manifold of M .




Figure 10: The image of the path γ is a subset of \mathbb{R}^2 that has the shape of a ‘eight’. In particular it is not simply connected, and at the origin it looks like a crossroad (as a set), although as a topological space, an open neighborhood of the origin is an open set of dimension 1, of the form $\gamma(]-\epsilon, \epsilon[)$.

Example 2.39. Another example consists of any irrational curve on the 2-torus: pick up an irrational number $\alpha \in \mathbb{R} \setminus \mathbb{Q}$ and define S to be the subset of \mathbb{T}^2 induced by the slope of slope α (it can be understood visually by using the well-known identification between \mathbb{T}^2 and a square). The topology of the irrational curve is so that connected open line segments are open, while the subspace topology does not allow this because the irrational curve is dense in \mathbb{T}^2 . Thus, the submanifold smooth structure does not descend from the ambient smooth structure so the submanifold is not embedded but merely immersed.

Let us now turn to *embedded* (or *regular*) *submanifolds* of a smooth manifold M : they are subsets S such that the subspace topology of M defines a canonical smooth structure on S (see

Theorem 8.2 in [Lee, 2003]). This property is a consequence of the fact that these submanifolds are modeled locally on the standard embedding of \mathbb{R}^k into \mathbb{R}^n . More precisely, let \tilde{U} be an open subset of \mathbb{R}^n , then a k -slice of \tilde{U} is any subset of the form:

$$\{(x^1, \dots, x^k, x^{k+1}, \dots, x^n) \in \tilde{U} \mid x^{k+1} = c^{k+1}, \dots, x^n = c^n\}$$

for some constants c^{k+1}, \dots, c^n . Clearly any k -slice is homeomorphic to an open subset of \mathbb{R}^k . Let M be a smooth manifold, and let (U, φ) be a smooth chart on M . If S is a subset of U such that $\varphi(S)$ is a k -slice of $\tilde{U} = \varphi(U)$, then we say simply that S is a k -slice of U . A subset $S \subset M$ is called a k -dimensional embedded submanifold of M if for each point $x \in S$, there exists a smooth chart (U, φ) for M such that $x \in U$, and $U \cap S$ is a k -slice of U . Equivalently, S is a k -dimensional embedded manifold in M if every point $x \in S$ is in the domain of a coordinate chart (U, φ) such that:

$$\varphi(U \cap S) = \varphi(U) \cap \{\mathbb{R}^k \times 0\} \quad (2.9)$$

The definition of embedded submanifolds is a local one, so that we can summarize it under the following Lemma:

Lemma 2.40. *Let M be a smooth manifold and let S be a subset of M . Suppose that for some k , every point $x \in S$ has a neighborhood $U \subset M$ such that $U \cap S$ is a k -dimensional embedded submanifold of U . Then S is a k -dimensional embedded submanifold of M .*

Example 2.41. The figure eight (Figure 10) is not an embedded submanifold because, although any point of the figure eight outside the origin belong to a 1-dimensional slice, there is no open set $U \subset \mathbb{R}^2$ containing the origin such that the intersection of U and the figure eight is an embedded submanifold of U (i.e. a one-dimensional slice of U).

Example 2.42. Let M and N be smooth manifolds of dimensions n and k , respectively, and let $F : M \rightarrow N$ be a smooth map. Let us call the *graph* of F the following subset of $\mathbb{R}^k \times \mathbb{R}^n$:

$$\text{Gr}(F) = \{(y, x) \in \mathbb{R}^k \times \mathbb{R}^n \mid y = F(x)\}$$

Indeed, Lemma 8.6 in [Lee, 2003] shows that locally this graph is embedded. Hence, by Lemma 2.40, it is an embedded submanifold.

A nice characterization of embedded submanifolds is obtained through the observation that the slice property carried by embedded submanifolds is equivalent to being locally the level set of a submersion:

Proposition 2.43. Constant rank level set theorem *Let M and N be smooth manifolds, and let $F : M \rightarrow N$ be a smooth map with constant rank equal to k . Each level set of F is a closed embedded submanifold of codimension k in M . In particular, a subset S of M is an embedded submanifold of M of codimension k if and only if every point $x \in S$ has a neighborhood U in M such that $U \cap S$ is a level set of a submersion $U \rightarrow \mathbb{R}^k$.*

Proof. See Chapter 8 in [Lee, 2003]. □

Remark 2.44. We say that a submanifold S of M is *closed* if its complement is open in M . The level set of a submersion is a closed embedded submanifold. From this we conclude that any embedded submanifold is locally closed in M .

A straightforward and very useful corollary of Proposition 2.43 relies on the following notions. If $F : M \rightarrow N$ is a smooth map, a point $x \in M$ is said to be a *regular point* of F if the push-forward $F_* : T_x M \rightarrow T_{F(x)} N$ is surjective; it is a *critical point* otherwise. A point $y \in N$ is said to be a *regular value* of F if every point of the level set $F^{-1}(y)$ is a regular point, and a *critical value* otherwise. Finally, a level set $F^{-1}(y)$ is called a *regular level set* if y is a regular value; in other words, a regular level set is a level set consisting entirely of regular points. Then, one has:

Theorem 2.45. Regular level set theorem *Every regular level set of a smooth map is a closed embedded submanifold whose codimension is equal to the dimension of the range.*

Proof. This is Corollary 8.10 in in [Lee, 2003]. □




Figure 11: The image of S through the map $\varphi : U \rightarrow \mathbb{R}^n$ is an open subset of $\mathbb{R}^k \times \{0\}$. Thus the preimage is a closed embedded submanifold of U .

Example 2.46. An alternative argument to show that the figure eight is not an embedded submanifold is that the figure eight is the zero level set of the smooth function:

$$F : \mathbb{R}^2 \longrightarrow \mathbb{R}$$

$$(x, y) \longmapsto x^2 - 4y^2(1 - y^2)$$

This function does not satisfy the latter part of Proposition (2.43) because at $(0, 0) \in S$ it is not a submersion.

Example 2.47. Let \mathbb{R}^n be the *configuration space*, with the corresponding coordinates q^i . We call the cotangent bundle $P = T^*\mathbb{R}^n$ the *phase space*, with coordinates q^i and p_i (the latter are linear forms on the fibers). Then a *constraint* is a smooth function $\phi : P \rightarrow \mathbb{R}$. The 0-level locus of a set of constraints ϕ_1, \dots, ϕ_r is a subset Σ of P that we call the *constraint surface*. Usually, physicists assume that the constraints satisfy a so-called *regularity condition* that often take the form that for each point $x \in \Sigma$ there exists an open neighborhood U such that only r' constraints $\phi_{i_1}, \dots, \phi_{i_{r'}}$ are functionally independent over U , making $\Sigma \cap U$ an embedded submanifold of U of codimension r' (as a level set of the constant rank smooth map $(\phi_{i_1}, \dots, \phi_{i_{r'}}) : P \rightarrow \mathbb{R}^{r'}$). Then by Proposition 2.43, the constraint surface is an embedded submanifold of dimension $2n - r'$. For more details, see Chapters 1 and 2 of [Henneaux and Teitelboim, 1992].

Associated to immersed submanifolds and embedded submanifolds, there exist corresponding notions of maps: injective immersions and smooth embeddings. An *injective immersion* between

two smooth manifolds S and M is an injective smooth map $F : S \hookrightarrow M$ that is additionally an immersion, i.e. such that the pushforward $F_* : TS \rightarrow TM$ is injective (we can consider that F_* takes values in TM because F is injective). In particular, an injective smooth map is an immersion if and only if it has constant rank. A *topological embedding* $F : S \hookrightarrow M$ is a continuous map that is a homeomorphism onto its image, where the topology on the image $F(S)$ is the subspace topology induced from the smooth atlas on M . A *smooth embedding* is a topological embedding that is smooth and of constant rank (then it is automatically an immersion). Obviously not every injective immersion is a smooth embedding (not even on its image), however here are two cases where it happens:

1. S is compact
2. F is proper (i.e. $F^{-1}(K)$ is compact if and only if $K \subset M$ is compact)

because in both cases the map $F : S \hookrightarrow M$ is closed (see Proposition 7.4 in [Lee, 2003]). Moreover, immersions locally behave as smooth embeddings, but not globally (hence justifying that the figure eight is the image of an immersion and not an embedding). See Lemma 8.18 in [Lee, 2003] for more details.

Proposition 2.48. *Immersed submanifolds are precisely the images of injective immersions and embedded submanifolds are precisely the images of smooth embeddings.*

Proof. See Chapter 8 in [Lee, 2003]. □

Exercise 2.49. Define the following three open subsets of \mathbb{R} :

$$A_- =] - \infty, -\frac{\pi}{2}[, \quad A_0 =] - \frac{\pi}{2}, +\frac{\pi}{2}[, \quad A_+ =] + \frac{\pi}{2}, +\infty[$$

Denote by A their disjoint union so, in particular, $A = \mathbb{R} - \{-\frac{\pi}{2}, \frac{\pi}{2}\}$. Let $f : A \rightarrow \mathbb{R}^2$ be the smooth map defined on each subset as follows:

$$f|_{A_-} = \left(-e^{x - \frac{1}{x + \frac{\pi}{2}}}, e^{x - \frac{1}{x + \frac{\pi}{2}}} \right), \quad f|_{A_0} = (\tan(x), \tan(x)), \quad f|_{A_+} = \left(e^{-x + \frac{1}{x - \frac{\pi}{2}}}, -e^{-x + \frac{1}{x - \frac{\pi}{2}}} \right)$$

Prove that f is an injective immersion (with respect to the standard smooth structures on A and \mathbb{R}^2). Draw a conclusion about the image of f , and determine the tangent space of $\text{Im}(f)$ at the point $(0, 0)$.


Now let us study in more details the difference between immersed and embedded submanifolds. Notice that if one had chosen another parametrization for the figure eight in Example 2.38, we would have inherited a totally different topology. Another, alternative smooth map defining the figure eight (as a set) can be chosen to be :

$$\eta(t) = (-\sin(2t), \sin(t))$$

and the path corresponding to η would be the symmetric image of that with respect to γ with respect to the vertical axis (see Figure 10). The open sets would not be the same either, because for example the image of $\gamma(] - \epsilon, \epsilon[)$, although a connected open set with respect to the topology induced by γ , would not be open in the topology induced by η , for its preimage would consist of two disjoint intervals, and the origin $t = 0$ (closed point). Then it seems that, although the subset S is uniquely defined as the level sets of the points (x, y) satisfying the equation $x^2 = 4y^2(1 - y^2)$, it admits several – a priori non-equivalent – smooth structures. The above argument shows that the map $\eta^{-1} \circ \gamma :] - \pi, \pi[\rightarrow] - \pi, \pi[$ is not a smooth map, not even


a continuous one. However, if one would have a diffeomorphism ψ from $] - \pi, \pi[$ such that $\eta = \gamma \circ \psi$, we would certainly conclude that the two smooth structures on S can be considered as ‘equivalent’. This is not the case, but this equivalence property is worth extending to every immersed submanifolds.

Definition 2.50. *Immersed submanifolds $N_1 \xrightarrow{\varphi_1} M$ and $N_2 \xrightarrow{\varphi_2} M$ are called equivalent when there exists a diffeomorphism $\psi : N_1 \rightarrow N_2$ making the following diagram commutative:*



This is an equivalence relation on the set of immersed submanifolds of M , and thus each equivalence class has a unique representative (S, \mathcal{A}, ι) where S is a subset of M with a given smooth structure \mathcal{A} such that the inclusion ι is an immersion. We emphasized the presence of the maximal atlas \mathcal{A} because it will turn out to be central in the discussion. For example, as seen above, there are two possible atlases for the figure eight to be an immersed submanifold, which are not equivalent because the map $\eta \circ \gamma^{-1}$ is not continuous. Thus, the figure eight admits several non-equivalent smooth structures making it an immersed submanifold. More generally now assume that there exist two injective immersions $N_1 \xrightarrow{\varphi_1} M$ and $N_2 \xrightarrow{\varphi_2} M$ whose images coincide $S = \text{Im}(\varphi_1) = \text{Im}(\varphi_2)$. What is the condition on φ_1 and φ_2 for N_1 and N_2 to be equivalent? Obviously, if φ_1 and φ_2 are smooth embeddings (i.e. if S is an embedded submanifold), then a diffeomorphism between N_1 and N_2 satisfying the commutative triangle is $\varphi_2^{-1} \circ \varphi_1$. This solution work for smooth embeddings because they have the following property:

Definition 2.51. *Let N and M be smooth manifolds. A smooth map $F : N \rightarrow M$ will be called smoothly universal if for any smooth manifold N' and any smooth map $H : N' \rightarrow M$ such that $H(N') \subset F(N)$, there exists a smooth map $G : N' \rightarrow N$ making the following triangle commutative:*



Remark 2.52. This kind of maps is also called *initial morphism* elsewhere [Kolar et al., 1993]. The name comes from the fact that such maps are initial in the categorical sense [Karshon et al., 2022].

Since smooth embeddings are homeomorphisms onto their image, the following consequence holds: the topology on an embedded submanifold $S \subset M$ is unique and is the subspace topology of M on S . Then, the smooth structure on M induces a unique smooth structure on S . In other words, embedded submanifolds carry a unique smooth structure inherited from that of their ambient manifold, such that the inclusion map is a smooth embedding (this is the content of Theorem 8.2 in [Lee, 2003]). This is not the case for immersed submanifolds, because the

underlying set of an immersed submanifold S may admit different smooth structures making the inclusion map $\iota : S \rightarrow M$ an immersion. Indeed, not every injective immersion is smoothly universal as the following discussion shows: given N, N', F and G as in the Definition (and assuming that F is an injective immersion), one naive idea would be to use F^{-1} to lift the smooth map G to H . However, the smoothness of the map H then crucially depends on the smooth structure of N and the properties of the smooth map F or, said differently, if $F(N)$ is an immersed or an embedded manifold. In the latter case, one can always define the map H as the composite $F^{-1} \circ G$, which is a smooth map because F is a diffeomorphism onto its image. However, when $F(N)$ is an immersed submanifold, although well-defined the map $H = F^{-1} \circ G$ needs not be a smooth map with respect to the smooth structure on N . For example, although the two paths γ and η define the same subset S – the figure eight – in \mathbb{R}^2 , the map $\eta \circ \gamma^{-1}$ is not continuous. This is the content of Scholie 1.31-33 in [Warner, 1983], which contain a nice discussion on this topic. In particular Theorem 1.32 states that the lift G is smooth if and only if it is continuous, thus showing that not having the smoothly universal property has tremendous consequences. This justifies the following definition which, for Molino [Molino, 1988], goes back to Pradines [Pradines, 1985]:


Definition 2.53. *An injective immersion $N \xrightarrow{\iota} M$ that has the smoothly universal property is called a weak embedding. The image of such a map in M is called a weakly embedded submanifold.*

Remark 2.54. The authors in [Karshon et al., 2022] notice that a map having the smoothly universal property is necessarily injective. If the inclusion of a subset $S \subset M$ has this property, then they call S a *diffeological submanifold*. In that case the smooth structure on S is unique. Weakly embedded submanifolds are precisely those diffeological submanifolds such that the inclusion map is an immersion, as can be seen from the following example: the cusp $S = \{(x, y) \in \mathbb{R}^2 \mid x^2 = y^3\}$ is a diffeological submanifold of \mathbb{R}^2 but not a weakly embedded submanifold. See [Karshon et al., 2022] for details.

We deduce from the above discussion that smooth embeddings are weak embeddings, while injective immersions need not be. Hence the following sequence of inclusions:

$$\{\text{smooth embeddings}\} \subset \{\text{weak embeddings}\} \subset \{\text{injective immersions}\}$$

As for Proposition 2.48, *weakly embedded submanifolds* correspond to the images of weak embeddings (some authors call them *regularly immersed* submanifolds). By construction, they are immersed submanifolds, but need not be embedded submanifolds. Weak embeddings possess a universal property making the smooth structure of a weakly embedded manifold unique, up to the equivalence given in Definition 2.50. More precisely, assume that a weakly embedded submanifold S of M is obtained via a weak embedding $\varphi_1 : N_1 \rightarrow M$ – which is a smooth map with respect to a maximal smooth atlas \mathcal{A}_1 , and assume moreover that S admits another weak embedding $\varphi_2 : N_2 \rightarrow M$ with respect to a smooth structure \mathcal{A}_2 on N_2 . Then, by the smoothly universal property of weak embeddings, both maps $\varphi_1^{-1} \circ \varphi_2 : (N_2, \mathcal{A}_2) \rightarrow (N_1, \mathcal{A}_1)$ and $\varphi_2^{-1} \circ \varphi_1 : (N_1, \mathcal{A}_1) \rightarrow (N_2, \mathcal{A}_2)$ are smooth. Being injective and inverse to one another, they define a diffeomorphism between (N_1, \mathcal{A}_1) and (N_2, \mathcal{A}_2) , thus showing that the two smooth structures are equivalent in the sense of Definition (2.50):



Weakly embedded submanifolds can then be considered as those submanifolds that have the right amount of regularity so that they carry only one possible smooth structure making the inclusion map an immersion. We will now give more details on their geometric properties and explain why their smooth structure – although uniquely defined by that of M – is not necessarily induced by the subspace topology. Given a subset S of a smooth manifold M and a point $x \in S$, we denote by $C_x(S)$ the (smooth) path connected component of x in S , i.e. the set of points that are reachable from x by smooth curves contained entirely in S (here, a smooth curve is a smooth map $\gamma : \mathbb{R} \rightarrow M$). Then a weakly embedded submanifold is characterized by the following property, that mimick Equation (2.9), but only at the level of path connected components:

Proposition 2.55. *Let M be a smooth manifold and let S be a weakly embedded submanifold of M of dimension k . Then for every $x \in S$, there exists a coordinate chart (U, φ) centered at x such that:*

$$\varphi(C_x(U \cap S)) = \varphi(U) \cap \{\mathbb{R}^k \times 0\} \tag{2.10}$$

Proof. See Propositions 3.19 and 3.20 in [Lee, 2009]. □




Figure 12: Assume that S is a weakly embedded submanifold of M that is additionally dense in M . Then, although $S \cap U$ consists of an infinite number of disjoint ‘plaques’, the path component of $x \in S$ in U is connected (by definition). Then its image through $\varphi : U \rightarrow \mathbb{R}^n$ is an open subset of $\mathbb{R}^k \times \{0\}$.

Once again, we see why the figure eight is not even a weakly embedded submanifold: at the origin, the path-connected component forms a cross shaped set, which does not have the

slice property of Equation (2.10). An example of a weakly embedded manifold which is not an embedded manifold is any leaf of the Kronecker foliation of the torus:

Example 2.56. Let $\mathbb{T} = \mathbb{R}^2 / \mathbb{Z}^2$ be the torus and let $X = \frac{\partial}{\partial x} + \alpha \frac{\partial}{\partial y}$ be a vector field on \mathbb{T} such that $\alpha \in \mathbb{R} - \mathbb{Q}$. Then the integral curve of X through any point (x, y) is a dense subset of \mathbb{T} , that is a weakly embedded submanifold. Indeed, any open neighborhood U of $(x, y) \in \mathbb{T}$ intersects infinitely many times the leaf through (x, y) (because it is dense). However, the set of points of U which are path-connected to (x, y) satisfy Equation (2.10) when U is taken to be sufficiently small. More generally it has been shown by Štefan in 1974 that leaves of (possibly singular) foliations are weakly embedded submanifolds, see [Miyamoto, 2023] for a thorough discussion on this question.

Let us conclude this section by a rather useful result, which is a variation of Proposition 2.55 for immersed submanifolds. Although immersed submanifolds do not admit the local structure of embedded or weakly embedded submanifold as a level set of a constant rank smooth map, there exist local distinguished coordinates characterizing open sets of immersed submanifolds:

Proposition 2.57. *Let $N \xrightarrow{F} M$ be an immersed submanifold of M and let $x \in N$. Then there exists a connected open neighborhood V of x in N and a coordinate chart (U, φ) centered at $F(x)$ such that:*

$$\varphi(U \cap F(V)) = \varphi(U) \cap (\mathbb{R}^k \times \{0\}) \quad (2.11)$$

Proof. The proof can be found in the discussion on page 131 of [Lee, 2009] and complemented by Proposition 1.35 of [Warner, 1983]. \square

Condition (2.11) emphasizes that the image in M of some connected open neighborhood of every point of S is embedded in M . Notice the difference with Lemma 2.40 which characterize embedded submanifolds. Also notice the difference between Equation (2.11) and the one for weakly embedded submanifolds (2.10) and for embedded submanifolds (2.9). We see that in each case the condition is stronger and stronger as we climb the hierarchy of submanifolds:

$$\{\text{embedded submanifolds}\} \subset \{\text{weakly embedded submanifolds}\} \subset \{\text{immersed submanifolds}\}$$

While conditions (2.9) and (2.10) are necessary and sufficient conditions to define embedded and weakly embedded submanifolds (see Propositions 3.19 and 3.20 in [Lee, 2009]), condition (2.11) does not characterize immersed submanifolds, as it is a particular case of the so-called *rank theorem* (see theorem 7.12 in [Lee, 2003]). However if the function F is injective and has constant rank then it implies that it is an immersion and then that the image is an immersed submanifold.

2.4 Distributions and foliations

Submanifolds possess their own tangent bundles, but it is often useful to see them as sub-bundles of the tangent bundle of M . That is why we benefit from the fact that every submanifold – be it immersed, weakly embedded or embedded – is the image of an immersion, to identify the tangent space to a submanifold $S \subset M$ at $x \in S$ with the image of the tangent space $T_x S$ as the image in $T_x M$ of the pushforward of the inclusion map ι_* – or the pushforward of the map $F : N \hookrightarrow M$ defining S :

$$T_{F(x)} S = F_*(T_x N)$$

Then, we often identify the tangent bundle of S (in M) with the subbundle of TM whose base is restricted to S and whose fiber is $T_x S$ at any point $x \in S$. This subbundle satisfies the following

nice characterization:

$$T_x S \subset \left\{ X_x \in T_x M \mid X_x(f) = 0 \text{ whenever } f \in C^\infty(M) \text{ and } f|_S \equiv 0 \right\}$$

where the inclusion is an equality (at least) when S is an embedded submanifold (see Proposition 8.5 in [Lee, 2003] for a demonstration). However, by Proposition 2.57, one observes that there exists an open neighborhood V of x in S such that:

$$T_x S = \left\{ X_x \in T_x M \mid X_x(f) = 0 \text{ whenever } f \in C^\infty(M) \text{ and } f|_{F(V)} \equiv 0 \right\}$$

This equality can be explained by the fact that $F(V)$ is an embedded submanifold of M , and that $T_x S = T_x F(V)$. We can go further to conclude that when S is a *closed* embedded submanifold generated by a set of smooth functions, the following useful Lemma holds:

Lemma 2.58. *Let $f_1, \dots, f_k \in C^\infty(M)$ be a set of functionally independent functions and let $S = \bigcap_{i=1}^k f_i^{-1}(0)$ be the closed embedded submanifold obtained as the intersection of the zero-level sets of each such function. Then the tangent space of S at a point $x \in S$ is:*

$$T_x S = \left\{ X_x \in T_x M \mid X_x(f_i) = 0 \text{ for all } 1 \leq i \leq k \right\}$$

Proof. The ideal of smooth functions vanishing on S is generated by f_1, \dots, f_k . □

Example 2.59. Although the origin in the figure eight is located at the crossroad of two one-dimensional paths, the tangent space at the origin of the figure eight is considered to be one dimensional, since it is the pushforward of $T] - \pi, \pi[$ through γ_* .

Example 2.60. The inclusion may not hold for weakly embedded submanifolds, as the example for the Kronecker foliation shows: since every leaf is dense in \mathbb{T} , the only function f that vanish on the leaf passing through x is the zero function, and hence every tangent vector at x satisfies $X_x(f) = 0$. The tangent space to the leaf is one dimensional, hence strictly included into $T_x \mathbb{T}$. However, if one had replaced the condition $f|_S \equiv 0$ by $f|_{C_x(S)} \equiv 0$ – for any smooth function $f \in C^\infty(U)$, for any arbitrary open neighborhood U of x – then the inclusion would have been an equality for weakly embedded submanifolds, but still not for immersed submanifolds (think of the figure eight).

The tangent bundle of a submanifold $S \subset M$ is a subbundle of the tangent bundle of M , when the base is restricted to S : $TS \subset T_S M$. However, assume now that we have a subbundle of TM defined over the entire manifold M . Then we expect that, under some circumstances, there may exist a family of ‘parallel’ submanifolds whose tangent bundles are precisely these subbundle.

Definition 2.61. *A (smooth) distribution on M is a smooth assignment⁶, to every point $x \in M$, of a vector subspace D_x of the tangent space $T_x M$. We say that the distribution D is regular if the function $x \mapsto \dim(D_x)$ is constant over M – in that case D forms a subbundle of TM – and it is said singular (or generalized) otherwise. We say that the distribution is involutive if the sheaf of smooth sections of D is stable under the Lie bracket of vector fields:*

$$\forall X, Y \in \Gamma(D) \quad [X, Y] \in \Gamma(D)$$

An integral manifold of D is an immersed submanifold S such that $T_x S = D_x$ for every $x \in S$. A distribution D is said integrable if through each point of M passes an integral manifold of D .

⁶Here, smooth means that for every tangent vector $X_x \in D_x$, it is always possible to find a locally defined vector field X such that for every point y in a neighborhood of x , $X_y \in D_y$.

Remark 2.62. Sometimes people define integral manifolds to be those submanifolds that satisfy the following inclusion $T_x S \subset D_x$, and then define a *maximal integral manifold of D* to be an integral manifold that is maximal with respect to inclusion; in particular, which satisfies the equality $T_x S = D_x$. On the other hand, an *invariant manifold of D* would be an immersed submanifold S such that $D_x \subset T_x S$ for every $x \in S$. The name ‘invariant’ comes from the fact that S is invariant under the action of the flows of sections of D .

Remark 2.63. Notice that the function $x \mapsto \dim(D_x)$, as a map from a topological space into the integers, is lower semi-continuous, and thus, the rank of the distribution D is locally constant and, in a vicinity of any given any point x , it can only be higher than or equal to that of D_x .

Example 2.64. There exist non-smooth integrable distributions. Let $M = \mathbb{R}^2$ and let D be the distribution defined as follows:

$$D_{(x,y)} = \begin{cases} \{0\} & \text{if } x \neq 0 \\ \langle \partial_y \rangle & \text{if } x = 0 \end{cases}$$

The corresponding integral manifolds are the points (x, y) when $x \neq 0$ and the vertical axis. The distribution is not smooth because there is no way of extending – as a smooth section of D – a non-trivial tangent vector defined at the origin $(0, 0)$ to a small neighborhood because the distribution outside the vertical axis is trivial. Although the distribution D is integrable, we do not consider it forms a *singular foliation* because it does not satisfy the axioms that we will soon present.

An integrable regular (resp. singular) distribution corresponds to what is commonly known as a *regular* (resp. *singular*) *foliation*. We do not want to enter the wide area of foliation theory for now, so we stick to the regular case and to regular distributions. The following definition should certainly be sufficient to understand the basic idea: a *foliation atlas of codimension p* on M (where $0 \leq p \leq n$) is an atlas made of charts call *foliation charts* and that are such that:

1. the image of the domain of any foliation chart (U, φ) through φ decomposes as a product of connected open sets $\varphi(U) = \tilde{U}' \times \tilde{U}'' \subset \mathbb{R}^{n-p} \times \mathbb{R}^p$
2. the transition function between two foliation charts (U, φ) and (V, ψ) is of the form:

$$\psi \circ \varphi^{-1}(a, b) = (g(a, b), h(b)) \in \mathbb{R}^{n-p} \times \mathbb{R}^p \quad (2.12)$$

where $g : \mathbb{R}^n \longrightarrow \mathbb{R}^{n-p}$ and $h : \mathbb{R}^p \longrightarrow \mathbb{R}^p$ are smooth maps.

Thus the domain U of each foliation chart (U, φ) is partitioned into the connected components of the submanifolds $\varphi^{-1}(\mathbb{R}^{n-p} \times y)$, $y \in \mathbb{R}^p$, called *plaques*. Being the connected components of the level sets of a smooth map of constant rank, plaques are connected embedded submanifolds of M of dimension $n - p$. The change-of-charts diffeomorphisms defined in Equation (2.12) preserve the plaques. Then, the union of plaques which overlap amalgamate into an immersed (in fact, weakly embedded) submanifold of M . Those submanifolds which are maximal with respect to inclusion are called *leaves*. More precisely (and we do not have time nor space to detail it here), two points $x, y \in M$ lie on the same leaf if there exists a sequence of foliation charts U_1, \dots, U_k and a sequence of points $x = x_0, x_1, \dots, x_k = y$ such that x_{i-1} and x_i lie on the same plaque in U_i . This defines an equivalence relation, so that the leaves of a foliated atlas of codimension p forms a partition of M by disjoint connected immersed submanifolds of dimension $n - p$. This observation justifies the following abstract definition (although one should stick to the idea that a foliation is a partition of M into leaves):

Definition 2.65. A foliation of codimension p on a smooth manifold M is a choice of maximal foliation atlas on M of codimension p .

Remark 2.66. The immersed submanifolds are actually weakly embedded [Stefan, 1974]. The charts defined in the definition are called *foliated charts* and there exists an atlas for M made of foliated charts, that are additionally compatible to one another, in the sense that the transition maps $\psi \circ \varphi^{-1}$ send slices to slices and preserve their transversal. We call such an atlas a *foliated atlas*. See [Candel and Conlon, 2000] and [Moerdijk and Mrcun, 2003] for details on foliations.




Figure 13: Three different foliations of the torus: the vertical one, the horizontal one, and the last one being characterized by the slope α . When α is an irrational real number, each leaf is dense in the torus: this is the Kronecker foliation. Picture taken from [Lee, 2003].

The relationship between distributions and foliations is that maximal (with respect to inclusion) connected integral manifolds of an integrable regular distribution form the leaves of a regular foliations:

Proposition 2.67. *Let D be an integrable regular distribution on a smooth manifold M . The collection of all maximal connected integral manifolds of D forms a foliation of M .*

This statement justifies the name ‘integrable’, since the regular distribution is thus integrable to a regular foliation, such that the leaves are the maximal connected integral manifolds of the distribution. However, this proposition does not tell us under which circumstances a regular distribution D is integrable. First observe that the tangent spaces to the leaves of a regular foliation define an involutive distribution. Thus an integrable distribution is necessarily involutive. The converse is actually also true, and this is the celebrated theorem of Frobenius (although he was not the first to state it):

Theorem 2.68. Frobenius Theorem *A regular distribution D on a smooth manifold is integrable (to a regular foliation) if and only if it is involutive.*

Proof. For more details on this subject, see Chapter 19 in [Lee, 2003], or Chapter 11 in [Lee, 2009], or [Candel and Conlon, 2000] and [Moerdijk and Mrcun, 2003]. \square

Example 2.69. Let ϕ_1, \dots, ϕ_r be a set of constraints on a phase space $T^*\mathbb{R}^n$, satisfying the regularity condition of Example 2.47: the constraint surface Σ is then a $2n - r'$ -dimensional embedded submanifold of $T^*\mathbb{R}^n$. The vector fields $X_i = \{\phi_i, \cdot\}$ generate a distribution on $T^*\mathbb{R}^n$ that is regular of rank r' on the constraint surface. We say that the constraints are *first-class* if the canonical Poisson bracket on $T^*\mathbb{R}^n$ of two such constraints vanishes on the constraint surface, i.e. if we have:

$$\{\phi_i, \phi_j\} = C_{ij}^k \phi_k$$

where the C_{ij}^k are smooth functions on $T^*\mathbb{R}^n$. If otherwise, we say that they are *second-class*. Then, a set of first-class constraints define an involutive, and then integrable, distribution on Σ . The leaves of the induced foliation are immersed (in fact, weakly embedded) submanifolds in Σ (and thus in $T^*\mathbb{R}^n$) of dimension r' , and correspond to the gauge equivalent physical configurations.

Exercise 2.70. By using the Jacobi identity satisfied by the Poisson bracket, compute $[X_i, X_j]$ and show that the distribution generated by the X_i is involutive (at least) on Σ .

Now what happens when the distribution is not involutive? It means that there exist (smooth) sections X, Y of D such that their Lie bracket $[X, Y]$ is not a section of D anymore. In particular, there is a point x such that the tangent vector $[X, Y]_x$ does not belong to D_x . Taking the successive brackets of (smooth) sections of D , and evaluating them at the point x thus may generate a subspace at x that is way bigger than D_x . We set $\text{Lie}(\Gamma(D))_x$ to be the distribution corresponding to the Lie algebra generated by $\Gamma(D)$ under the successive action of the Lie bracket of vector fields on smooth sections of D :

$$\text{Lie}(\Gamma(D))_x = D_x + \text{Span}\left([X_1, X_2]_x, [[X_1, X_2], X_3]_x, [[[X_1, X_2], X_3], X_4]_x, \dots \mid X_i \in \Gamma(D)\right)$$

Notice that it may not be a regular distribution, although interesting things happen when it is:

Definition 2.71. Hormander's condition *Let D be a distribution. We say that D is bracket generating at x if:*

$$\text{Lie}(\Gamma(D))_x = T_x M$$

We say that D is maximally non-integrable if D is bracket generating at every point.

The latter notion comes from the fact that Theorem 2.68 can be reformulated as the statement that a distribution is integrable if and only if $\text{Lie}(\Gamma(D))_x = D_x$. Then, obviously, if at some point $\text{Lie}(\Gamma(D))_x$ is strictly bigger than D_x , the distribution will not be integrable. Consequently, the situation where $\text{Lie}(\Gamma(D))_x = T_x M$ at every point can legitimately be considered as the worst case scenario where D is non-integrable in the worst possible way. However, maximally non-integrable distribution have a nice property: from the fact that if a distribution D is bracket generating at a given point x , every point in a small neighborhood of x can be reached through a so-called 'horizontal' path. A *horizontal path* is a path $\gamma : [0, 1] \rightarrow M$ that is:

1. absolutely continuous on every local coordinate chart, and
2. such that $\dot{\gamma}(t) \in D_{\gamma(t)}$ almost everywhere.

The notion of absolutely continuous paths is often met in the field of control theory under the following form: assume that X_1, \dots, X_m are smooth sections of D that are defined in a neighborhood of $\gamma([0, 1])$, where γ is up to now only a continuous path. Then it is said absolutely continuous if there exist m absolutely continuous functions $u_i \in L^1([0, 1])$ such that the following equation holds almost everywhere:

$$\dot{\gamma}(t) = \sum_{i=1}^m u_i(t) X_{i, \gamma(t)}$$

The functions u_1, \dots, u_m are called the controls of γ with respect to the vector fields X_1, \dots, X_m . When the distribution is induced by a physical system, and that $D_x \neq T_x M$, we say that the system is *non-holonomic* – joining two points may not be possible if one restricts itself to horizontal paths only – while if $D_x = T_x M$, we say that the system is *holonomic* – one could always join one point of the state space to any other through horizontal paths. Then we have the infamous following result that answer the problem for non-holonomic systems:

Theorem 2.72. Chow-Rashevskii theorem Let M be a smooth manifold and let D be a smooth distribution that is bracket generating at a given point $x \in M$. Then, there exists a neighborhood of x on which every point can be joined from x by an horizontal path.

Corollary 2.73. If D is maximally non-integrable, every two points of the manifold M can be joined through a horizontal path.

Proof. See Section 3.2 of [Agrachev et al., 2019]. □

Exercise 2.74. Check that the distribution D of rank 2 on \mathbb{R}^3 generated by the vector fields $X = \frac{\partial}{\partial x^1} + x^2 \frac{\partial}{\partial x^3}$ and $Y = \frac{\partial}{\partial x^2}$ is maximally non-integrable.




Figure 14: Although the distribution D defined in Exercise 2.74 does not contain the vertical tangent vector $\frac{\partial}{\partial x^3}$, we can however reach the point $(0, 0, 1)$ from $(0, 0, 0)$ through a sequence of paths whose tangent vectors are in D at each point.

The corollary give some more insight on the denomination *maximally non-integrable*: such a distribution does not have ‘leaves’ *per se*, and on the contrary, every two points of the manifold

can be joined though an absolutely continuous path almost everywhere tangent to the distribution. We will use these notions to explain how Constantin Carathéodory defined a geometric approach to thermodynamics, and how he deduced the existence of a function called the entropy. The following discussion is mainly inspired by Chapter 22 of [Bamberg and Sternberg, 1988a].

In thermodynamics, we distinguish between two kinds of physical systems: *closed systems* are those that are spatially bounded and that allow heat transfer with the exterior but no matter transfer of any kind, while *open systems* are those physical systems allowing both heat and matter transfers. Although open systems are those that are found in nature, we will restrict ourselves to closed ones, which are a very practical modelization. To every closed thermodynamical system is associated a thermodynamical state space, consisting of all its equilibrium states. Although we may assume that it is a smooth manifold (possibly with boundary), it turns out that it is often a vector space or a half space. Usually, it admits three types of coordinates: some empirical temperature θ or several, depending on the number of reservoir; some intensive variables corresponding to generalized force such as the pressure P or a magnetic intensity; and extensive variables measuring variations of volume V or of magnetization, etc. [Zemansky, 1966]. It turns out that the existence of equations of states – such as the one relating the internal energy to the thermodynamic variables, see Scholie 2.75 – implies that the intensive variables can be made dependent on the (then independent) temperatures and on the extensive variables. We will adopt the convention that paths in the state manifold correspond to reversible thermodynamic processes.

There are mainly two kinds of thermodynamic transformations: those in which we apply some *work* W to the physical system, and those in which there is a *heat transfer* Q between the system and the exterior. The corresponding infinitesimal thermodynamic transformations are denoted δW and δQ , respectively. They are differential one-forms which, when integrated over a reversible thermodynamic process represented by a path γ , gives the total amount of work and of heat that has been exchanged:

$$Q_\gamma = \int_0^1 \delta Q(\dot{\gamma}(t)) dt \quad \text{and} \quad W_\gamma = \int_0^1 \delta W(\dot{\gamma}(t)) dt$$

The symbol δ not only symbolizes that the objects δW and δQ correspond to infinitesimal transformations, but also that they are *not* exact one-forms. More precisely, the quantity of work and of heat that is applied to or retrieved from the system depends on the way we apply or retrieve it (it is process dependent). One family of such processes is fundamental in thermodynamics for its usefulness: an *adiabatic process* is a thermodynamic process for which there is no heat transfer with the exterior, i.e. for which $Q_\gamma = 0$. Adiabaticity is a property that is central in Carathéodory's reformulations of the first and the second principle of thermodynamics [Sears, 1966]:

Scholie 2.75. Carathéodory's first principle of Thermodynamics *For a closed thermodynamic system, in all adiabatic reversible thermodynamic processes between an initial state and a final state, the work does not depend on the path chosen. In particular, this implies that there exists a well-defined function U called the internal energy such that its infinitesimal variations satisfies:*

$$dU = \delta Q + \delta W$$

Proof. The proof that the second statement is a consequence of the first can be found in [Sears, 1963]. \square

The integration of an exact one form over a path γ joining two points x and y in a simply

connected space only depends on the ends points, and not on the path chosen:

$$U(y) - U(x) = \int_{\gamma} dU = \int_0^1 \dot{\gamma}(t)(U) dt$$

This makes the internal energy a *state function*, i.e. a function on the state space whose variations only depend on the initial state and the final state of the system – which is not the case for work and heat transfer. There may be different kinds of work δW and one of the most used is the one consisting of increasing or decreasing the volume of a given volume of gas, so that:

$$\delta W = -pdV + \nu_1 d\mu_1 + \dots$$

The (certainly non-exact) differential one-form $\delta Q = dU - \delta W$ then corresponds to the infinitesimal heat production or absorption. The kernel of the differential one form $\alpha = \delta Q$ defines a distribution $D = \text{Ker}(\alpha)$ such that at every point $x \in M$, $D_x = \text{Ker}(\alpha_x) \subset T_x M$, and that for the sake of the presentation we will assume to be regular. This distribution has rank $n - 1$ and then the question is: is it integrable or maximally non-integrable? More precisely, given the equivalence between involutivity and integrability for regular distributions, do we have $\text{Lie}(\Gamma(D))_x = D_x$ or, on the contrary, do we have $\text{Lie}(\Gamma(D))_x = T_x M$? There exists obviously a middle ground: at some point the distribution may be bracket generating while at others it may not, but we will see that this situation is not met in our context.

If the distribution $D = \text{Ker}(\alpha)$ is maximally non-integrable, it means that every two points of the state space can be joined through a horizontal path, i.e. through a succession of reversible adiabatic transformations. On the contrary, if the distribution D is integrable, then we can deduce some properties of the differential one form δQ . Indeed, one can show that an alternative form of Frobenius theorem states that the graded ideal $\mathcal{I}_{\alpha}^{\bullet} = \bigoplus_{1 \leq m \leq n} I_{\alpha}^m$ of the graded commutative algebra $\Omega^{\bullet}(M)$ generated by α – i.e. $I_{\alpha}^1 = \text{Span}(\alpha)$ and $I_{\alpha}^m = \text{Span}(\eta_1 \wedge \dots \wedge \eta_{m-1} \wedge \alpha \mid \eta_i \in \Omega^1(M))$ – is actually a *differential graded ideal*, i.e. it is stable under the de Rham differential:

$$d\mathcal{I}_{\alpha}^{\bullet} \subset \mathcal{I}_{\alpha}^{\bullet}$$

For details on this statement, see for example Theorem 1.3.8 and Exercise 1.3.12 in [Candel and Conlon, 2000]. So, in particular, since $d\alpha \in I_{\alpha}^2$, there exists a one form η such that, at least locally, $d\alpha = \eta \wedge \alpha$. One can then show that this identity holds if and only if there exist two smooth functions $f, g \in C^{\infty}$ such that $\alpha = fdg$. This observation leads to Carathéodory's (partial) reformulation of the second principle of thermodynamics:

Scholie 2.76. Carathéodory's second principle of Thermodynamics, *a.k.a* adiabatic inaccessibility *Given a closed system, in every neighborhood of any state x there are states inaccessible from x through adiabatic (reversible) processes. In particular, this implies that there exists two smooth functions – T called the temperature and S called the entropy – such that the differential form δQ takes the following form:*

$$\delta Q = TdS$$

Proof. If, in the vicinity of each equilibrium state, there are other states which are *not* reachable through adiabatic reversible transformations, then the distribution $D = \text{Ker}(\alpha)$ is *not* maximally non-integrable and, the assumption of non-accessibility holding at every point, we deduce that it is integrable. But then, by the above discussion on integrable distribution of rank $n - 1$, we deduce that $\alpha = \delta Q$ can be written as fdg or, for the sake of consistency with traditional notations, $\delta Q = TdS$. The fact that this equality holds globally and not only locally comes from the fact that the thermodynamic state space is often a vector space, on which the cotangent bundle is trivial. \square

Remark 2.77. Actually, Carathéodory's principle is not a faithful second principle, because it says nothing about the conditions under which the entropy increases. That is why it is necessary to supplement it with Planck's principle, stating that adiabatic isochoric processes always increase the internal energy of a closed system, hence corresponding to an increase of entropy. The only way that the entropy of a closed system can decrease is when heat is transferred from the system to the exterior. See e.g. [Sears, 1966] and [Zemansky, 1966] for a discussion on the relationships between non-equivalent statements of the second principle.

2.5 Orientation of smooth manifolds and integration of differential forms

Now we have enough material to define integration of differential forms on smooth manifolds. Theoretically one can integrate any differential k -forms, but this relies on advanced mathematics so we would rather only concentrate on integrating differential n -forms. This is consistent with what theoretical physicists mostly do in their everyday life. We would proceed as usual: integration on a manifold M would first be defined locally, because we know how to integrate differential n -forms in \mathbb{R}^n , and then using a partition of unity we can sum up all the local contributions to obtain an integral over M . A necessary condition to integrate is to have an *orientable* manifold. In this section we assume that the dimension of manifolds and vector spaces are greater than or equal to 1.

Given a n -dimensional vector space E , we say that two ordered basis e_1, \dots, e_n and e'_1, \dots, e'_n are *consistently oriented* if the transition matrix from one to the other has positive determinant. This relation is an equivalence relation. Since $\mathbb{R} - \{0\}$ has two disjoint connected components, there are only two equivalence classes of consistently oriented ordered bases: either the determinant of the transition matrix is positive and we stay in the equivalence class, or it is negative, and we change class. We call an *orientation* on E either of those equivalence classes of those consistently oriented ordered bases. There is no absolute choice of orientation on a vector space (except maybe for \mathbb{R}^n), there are only relative choices: once we have chosen an ordered basis e_1, \dots, e_n , it is a convention to say that every other consistently oriented ordered basis is *positively oriented*. A basis that is obtained from e_1, \dots, e_n through a transition matrix with negative determinant is said *negatively oriented*.

Example 2.78. The vector space \mathbb{R}^n admits the following standard basis $e_i = (0, \dots, 0, 1, 0, \dots, 0)$ where 1 is located at the i -th position. We say that the orientation defined by this basis is the standard orientation of \mathbb{R}^n .

Lemma 2.79. *Let E be a vector space of dimension $n \geq 1$, and suppose ω is a nonzero element of $\Lambda^n(E^*)$. The set of ordered bases e_1, \dots, e_n such that $\omega(e_1, \dots, e_n)$ has the same sign is an orientation for E .*

Proof. Let e_1, \dots, e_n and e'_1, \dots, e'_n be two basis of E and let B the transition matrix from the former to the latter: $e'_i = B_i^j e_j$. Then:

$$\omega(e'_1, \dots, e'_n) = \det(B)\omega(e_1, \dots, e_n)$$

so that e_1, \dots, e_n and e'_1, \dots, e'_n are consistently oriented if and only if $\omega(e'_1, \dots, e'_n)$ and $\omega(e_1, \dots, e_n)$ have the same sign. \square

Thus, choosing an orientation of a vector space E amounts to choosing an element ω of $\Lambda^n E^*$. Once this choice is made, we say that ω is a positively oriented n -covector. For example, if the ordered basis of E is given by e_1, \dots, e_n , the n -covector $\omega = e^1 \wedge \dots \wedge e^n$ is positively oriented. For any real scalar $\lambda > 0$, $\lambda\omega$ is another positively oriented n -covector, while for any real scalar

$\mu < 0$, $\mu\omega$ is said to be a negatively oriented n -covector. This plays some role in the definition of the Hodge star operator. Indeed, it depends on a choice of orientation of E because the volume form $\omega = e^1 \wedge e^2 \wedge \dots \wedge e^n$ is given by the choice of an ordered basis e_1, e_2, \dots, e_n . If one had taken the ordered basis $e_2, e_1, e_3, \dots, e_n$ instead – with reverse orientation, then – the associated volume form positively oriented with respect to the orientation defined by $e_2, e_1, e_3, \dots, e_n$ would be $\omega' = e^2 \wedge e^1 \wedge e^3 \wedge \dots \wedge e^n = -\omega$, so that the Hodge star operator \star' associated with ω' would be the opposite to the one associated with ω : $\star' = -\star$.

Since a smooth manifold is locally euclidean, we can define an orientation locally, at the level of the tangent bundle. A *pointwise orientation* on M is the assignment, to every point x , of an orientation of the fiber $T_x M$. It is always possible to equip a smooth manifold with such a pointwise orientation, but the difficulty comes from having this orientation varying smoothly over the manifold. A local smooth frame X_1, \dots, X_n of the tangent bundle over some open set U is said *positively oriented* if, for every $x \in U$, the orientation of the basis $X_{1,x}, \dots, X_{n,x}$ coincides with the orientation of $T_x M$. A pointwise orientation is said *smooth* if every point of M is in the domain of an oriented local smooth frame. Given two smooth manifolds M and N of the same dimension that admit smooth pointwise orientations, we say that a local diffeomorphism $F : M \rightarrow N$ is *orientation-preserving* if, for every $x \in M$, F_* takes oriented bases of $T_x M$ to oriented bases of $T_{F(x)} N$, and *orientation-reversing* if it takes (positively) oriented bases of $T_x M$ to negatively oriented bases of $T_{F(x)} N$.

We want to study how the existence of a smooth pointwise orientation translates at the level of charts and transition functions. Let M be a smooth manifold equipped with a (non necessarily smooth) pointwise orientation. Any smooth chart (U, φ) whose coordinate frame $\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^n}$ is positively (resp. negatively) oriented is called a positively (resp. negatively) *oriented chart*. Any smooth oriented chart (U, φ) on M always induce another chart $(U, \bar{\varphi})$ with reverse orientation. Indeed let $\bar{\varphi}$ be the composite of φ and a reflectional symmetry (which is a smooth map from \mathbb{R}^n to \mathbb{R}^n), then $(U, \bar{\varphi})$ has reverse orientation compared to (U, φ) . Obviously, there exist choices of pointwise orientations such that some charts are neither positively nor negatively oriented. However we will see that for smooth pointwise orientations, the situation is really nice. We say that two smooth oriented charts (U, φ) and (V, ψ) are *consistently oriented* if the transition map $\psi \circ \varphi^{-1}$ has positive Jacobian determinant, i.e. if it is orientation preserving. An *oriented atlas* on M is a smooth atlas for which all smooth charts are consistently oriented. M is *orientable* if it admits an oriented atlas, and an *orientation of M* is a choice of a maximal oriented atlas. The following proposition shows that the existence of a smooth pointwise orientation on M is equivalent to an orientation on M . Thus, oriented atlases form a subclass of smooth atlases, where the transition functions are not only diffeomorphisms, but also orientation preserving. The relationship between the two notions of orientability is actually very simple:

Proposition 2.80. *Let M be a positive ($n > 1$) smooth manifold equipped with a pointwise orientation. Then it is smooth if and only if it induces an orientation on M .*

Proof. First, notice that a smooth pointwise orientation on M implies that there exists an open cover of positively oriented charts. This can be seen as follows: let $x \in M$ and let X_1, \dots, X_n be an oriented frame defined on an open neighborhood of x . One can assume that this neighborhood is a smooth chart (U, φ) . Then, the induced coordinate frame is either positively oriented, or negatively oriented, but in that latter case the smooth chart $(U, \bar{\varphi})$ is positively oriented. Then we can find an open cover of positively oriented charts.

Assume that the chosen pointwise orientation on M is smooth and pick up such an open cover of oriented charts. Then, by Equation 2.2, on overlapping oriented charts, the transition functions are orientation-preserving. This implies that the open cover of positively oriented charts is an oriented atlas, providing M with an orientation. Conversely, every orientation

makes the pointwise orientation smooth because the coordinate frames are oriented frames, and two such frames define the same orientations on the fibers since the transition functions between oriented charts are orientation-preserving by hypothesis. \square

Thus, being smoothly pointwise orientable is equivalent to being orientable. If a smooth manifold is orientable, there are essentially two possible choices of orientations. Pick up a tangent space and attribute an orientation to this vector space (here we make a choice between two orientations). Then, by Proposition 2.80, the respective orientations of the other fibers of the tangent bundle will be automatically determined by this first choice. This can be seen from the fact that transition functions from one oriented chart to another are orientation preserving. Non-orientable manifolds are precisely those manifolds for which there are always at least one transition function that is not orientation preserving, whatever the choice of smooth chart we make. For a zero dimensional manifold, i.e. a point $\{*\}$, an orientation is a choice of map $\{*\} \mapsto \{\pm 1\}$. We know at least one evident situation where a smooth manifold is orientable:

Proposition 2.81. *Every parallelizable smooth manifold is orientable.*

Example 2.82. Every Lie group is parallelizable, hence is orientable.

Example 2.83. Spheres, planes and tori are orientable.

Example 2.84. The *Mobius bundle* is the vector bundle E over S^1 whose total space is defined as a quotient of \mathbb{R}^2 by the following relation:

$$(x, y) \sim (x + n, (-1)^n y)$$

The *Mobius band* is the subset $M \subset E$ of the Mobius bundle that is the image, under the above quotient map, of the set $\{(x, y) \in \mathbb{R}^2 \mid |y| \leq 1\}$. It is a smooth 1-manifold with boundary, which is non orientable.

The most important point with orientations is that we can characterize it through differential forms. A *volume form* on M is a global *nowhere vanishing* smooth section of the vector bundle $\wedge^n T^*M$. We usually denote such a section by the letter ω . Over a local smooth chart U , with respect to a coordinate coframe, the volume form decomposes as $\omega = f dx^1 \wedge \dots \wedge dx^n$, for some smooth function $f \in C^\infty(U)$. The existence of volume forms is tightly connected to that of orientations:

Proposition 2.85. *Let M be a smooth manifold of dimension $n \geq 1$. Any nowhere vanishing differential n -form $\omega \in \Omega^n(M)$ determines a unique orientation of M for which the n -covector $\omega(x) \in \wedge^n T^*M$ is positively oriented for every $x \in M$. Conversely, if M is given an orientation, there is a smooth nowhere vanishing differential n -form on M that is positively oriented at each point.*

Proof. Assume that there exists such a volume form ω , so by Lemma 2.79, the evaluation of the volume form ω at a point x induces an orientation of the tangent space $T_x M$, that is considered to be positively oriented. Now let us check that there exists an oriented smooth atlas for M . Let (U, φ) and (V, ψ) be two intersecting oriented charts. Let us denote by x^1, \dots, x^n and x'^1, \dots, x'^n the coordinates respectively associated to the maps φ and ψ . Then, on U the volume form reads $\omega = f dx^1 \wedge \dots \wedge dx^n$, while on V it reads $\omega = g dx'^1 \wedge \dots \wedge dx'^n$, for two nowhere vanishing functions $f \in C^\infty(U)$ and $g \in C^\infty(V)$. Over the intersection $U \cap V$, using the transformation laws found in Equations (2.7)-(2.8), one obtains that:

$$dx'^i \Big|_y = \frac{\partial x'^i}{\partial x^j}(\varphi(y)) dx^j \Big|_y \tag{2.13}$$

Then, we obtain that $g dx'^1 \wedge \dots \wedge dx'^n = g \text{Jac}(\psi \circ \varphi^{-1}) dx^1 \wedge \dots \wedge dx^n$, where Jac symbolizes the Jacobian determinant. Then, we have:

$$f(y) = g(y) \text{Jac}(\psi \circ \varphi^{-1})(\varphi(y))$$

for every $y \in U \cap V$. The sign of the Jacobian determinant is determined by the sign of the function $\frac{f}{g}$ which is nowhere vanishing over $U \cap V$.

Now, either f and g have the same sign, and then (U, φ) and (V, ψ) are consistently oriented, or they do not have the same sign. However, in that case, one may define another chart $(V, \bar{\psi})$ by changing a sign in the definition of ψ , e.g. $\psi(y) \mapsto \bar{\psi}(y) = (-\psi^1(y), \psi^2(y), \dots, \psi^n(y))$. This is possible because reflectional symmetries with respect to hyperplanes are diffeomorphisms of \mathbb{R}^n . We label the corresponding new coordinates as \bar{x}^i , and in particular $\bar{x}^1 = -x^1$ whereas for $2 \leq i \leq n$, $\bar{x}^i = x^i$. Then the volume form decomposes in this new coordinate coframe as $\omega = -g d\bar{x}^1 \wedge \dots \wedge d\bar{x}^n$, and then, the new Jacobian determinant reads: $\det(\bar{\psi} \circ \varphi^{-1})(\varphi(y)) = -\frac{f(y)}{g(y)}$ which is now a positive fraction for every $y \in U \cap V$. Thus, the oriented chart $(V, \bar{\psi})$ is consistently oriented with (U, φ) . This proves the first statement. For the converse statement – that any orientable smooth manifold admits a volume form – see Proposition 13.4 in [Lee, 2003]. \square

Since the vector bundle $\wedge^n T^*M$ has rank 1, any other *nowhere vanishing* smooth section $f\omega$, where $f \in \mathcal{C}^\infty(M)$, is a volume form as well. Since there are two disjoint connected components in $\mathbb{R} - \{0\}$, there are two equivalence classes of sections of the line bundle $\wedge^n T^*M$: those that are positively related to ω , and those that are negatively related to ω . Moreover, those volume forms that are negatively related with ω are still positively related among themselves. Thus, as the proof of Proposition 2.85 shows, picking up any other representent of the equivalence class of ω – i.e. of the form $f\omega$ for some strictly positive function – defines the same orientation on M as ω . Actually, the oriented atlas associated to ω is obtained as the collection of all smooth charts for which the standard volume form on \mathbb{R}^n (induced from the standard oriented basis presented in Example 2.78) pulls back to a positive multiple of ω . That is why some authors define an orientation on M as a choice of an equivalence class of positively related volume forms (see e.g. [Baez and Muniain, 1994, p. 84]):

Corollary 2.86. *There is a one-to-one correspondence between orientations on M and equivalence classes of positively related globally defined volume forms.*

Remark 2.87. There are homological and cohomological characterization of orientability. For example, a smooth manifold is orientable if and only if the first Stiefel-Whitney characteristic class is 0.

Orientability is necessary to define integration on smooth manifolds. Since a smooth manifold is locally euclidean, let us first define integration over \mathbb{R}^n , before generalizing to any smooth manifold using pullbacks. A subset of \mathbb{R}^n is a *domain of integration* if its boundary has n -dimensional measure 0. We usually define integration in \mathbb{R}^n defining first the integral of bounded continuous functions on ‘rectangles’, i.e. products of closed intervals. Then, every continuous function can be locally approximated by such functions, and every domain of integration can be covered by rectangles (given by the closure of open sets inherited from the subspace topology of \mathbb{R}^n on D), so that in the end we can define the integral of bounded continuous functions on any domain of integration. Then, a choice of domain of integration D defines a linear form on the space of bounded continuous functions on D :

$$\begin{aligned} \int_D : \mathcal{C}_b^0(D) &\longrightarrow \mathbb{R} \\ f &\longmapsto \int_D f dx^1 \dots dx^n \end{aligned}$$

where the notation $dx^1 \dots dx^n$ is purely abstract and needs not appear. It only reminds the reader that we integrate the function over a subset of \mathbb{R}^n . It is sometimes noted $d\mu$ to symbolize the Lebesgue measure. You can find more details on this construction in Appendix A of [Lee, 2003].

This definition straightforwardly generalizes to differential n -forms. Let $U \subset \mathbb{R}^n$ be an open set and let ω be a differential n -form compactly supported on some compact set $K \subset U$:

$$K = \text{supp}(\omega) = \overline{\{x \in M \mid \omega(x) \neq 0\}}$$

Lemma 14.1 in [Lee, 2003] shows that there always exists a domain of integration D such that $K \subset D \subset U$. Then, assuming that this differential form can be written as $\omega = f dx^1 \wedge \dots \wedge dx^n$ over its support K , the *integral of ω over U* is given by:

$$\int_U \omega = \int_D f dx^1 \dots dx^n$$

The notation here is very convenient: it is as if we had ‘erased’ the wedges. Notice that the above definition does not depend on the choice of domain of integration $K \subset D \subset U$.

Lemma 2.88. *Suppose U, V are open sets of \mathbb{R}^n and that $F : U \rightarrow V$ is a diffeomorphism. Let ω be a compactly supported differential n -form on V . Then:*

$$\int_V \omega = \begin{cases} \int_U F^* \omega & \text{if } F \text{ is orientation-preserving} \\ - \int_U F^* \omega & \text{if } F \text{ is orientation-reversing} \end{cases}$$

This lemma provides another formulation of the fact that a differential n -form can be written equivalently in two sets of coordinates x^1, \dots, x^n and x'^1, \dots, x'^n , defined over overlapping open sets U and V , and related through a diffeomorphism $F : U \rightarrow V$, e.g. such that $F(x^1, \dots, x^n) = (x'^1, \dots, x'^n)$. Then if one writes the differential form on V as $\omega = g dx'^1 \wedge \dots \wedge dx'^n$, Equation (2.13) implies:

$$\begin{aligned} F^* \omega &= F^*(g dx'^1 \wedge \dots \wedge dx'^n) \\ &= (F^*g) F^*(dx'^1) \wedge \dots \wedge F^*(dx'^n) \\ &= g \circ F \cdot \text{Jac}(F) dx^1 \wedge \dots \wedge dx^n \end{aligned}$$

where \cdot indicates the multiplication of two smooth functions on U . Thus, we obtain the infamous formula for a change of coordinates under integration:

$$\int_V g(x'^1, \dots, x'^n) dx'^1 \dots dx'^n = \int_U g \circ F(x^1, \dots, x^n) \text{Jac}(F)(x^1, \dots, x^n) dx^1 \dots dx^n$$

We have now enough material to define integration on manifolds. Let M be a n -dimensional oriented smooth manifold. First, let ω be a differential n -form compactly supported in the domain of a single oriented smooth chart (U, φ) . Then, we define the *integral of ω over U* as the following object:

$$\int_U \omega = \int_{\varphi(U)} (\varphi^{-1})^* \omega \tag{2.14}$$

The right-hand side is an integral over an open subset $\tilde{U} = \varphi(U)$ of \mathbb{R}^n . It is well defined because $(\varphi^{-1})^* \omega$ is a compactly supported differential n -forms on this open set. Lemma 2.88

implies that the integral of ω over any other choice of oriented smooth chart (V, ψ) containing its compact support would have given the same result:

$$\int_U \omega = \int_V \omega$$

Now that we have defined integration over compact support, we can extend integration over the whole manifold M by using a notion that is specific to smooth manifolds:

Definition 2.89. Let $\mathcal{U} = (U_\alpha)_{\alpha \in A}$ be any open cover of M (indexed over some set A). A partition of unity subordinate to \mathcal{U} is a collection of continuous functions $(\psi_\alpha : M \rightarrow \mathbb{R})_{\alpha \in A}$, with the following properties:

1. $0 \leq \psi_\alpha(x) \leq 1$ for all $\alpha \in A$ and all $x \in M$;
2. $\text{supp}(\psi_\alpha) \subset U_\alpha$;
3. for every $x \in M$, there is only a finite number of ψ_α such that $x \in \text{supp}(\psi_\alpha)$;
4. $\sum_{\alpha \in A} \psi_\alpha(x) = 1$ for all $x \in M$.

The third condition is equivalent to saying that the set of supports $\{\text{supp}(\psi_\alpha)\}_{\alpha \in A}$ is locally finite. Because of this condition, the sum in the last item has only finitely many nonzero terms in the neighborhood of each point, so there is no issue of convergence. When the functions ψ_α are smooth, then we say that they form a *smooth* partition of unity. The importance of partition of unity is central in differential geometry, as the following theorem shows:

Theorem 2.90. Any open cover \mathcal{U} of a smooth manifold admits a smooth partition of unity subordinate to \mathcal{U} .

Proof. The main point is that the fact that a topological manifold is Hausdorff and second-countable implies that it is paracompact (and that it has countably many connected components), which is the crucial property needed to demonstrate the result. However, the proof is long and subtle, so we refer to Chapter 2 of [Lee, 2003]. \square

Remark 2.91. While Theorem 2.90 show that smooth partitions of unity subordinate to any open cover of a smooth manifold exist, it is no longer the case for analytic manifolds. Indeed, the proof relies on the existence of smooth bump functions on $[-1, 1]$. Unfortunately, those bump functions are not analytic because of the so-called identity theorem, which is then an obstruction to the existence of analytic partitions of unity subordinate to any open cover of an analytic manifold.

To integrate over a (connected) smooth manifold, one needs an orientation. The latter is needed to integration over the entire manifold in order to ensure that local contributions, as defined by Equation (2.14), do not artefactually cancel one another because of a change in open chart, as shown by the change of sign in Lemma 2.88. Let ω be a compactly supported differential n -form on a connected oriented smooth manifold M . Then there exists a finite open cover $\{(U_i, \varphi_i)\}$ of oriented charts for $\text{supp}(\omega)$, and a partition of unity $\{\psi_i\}$ subordinated to this open cover. We define the integral of ω as follows:

$$\int_M \omega = \sum_i \int_{U_i} \psi_i \omega \tag{2.15}$$

For each i , the n -form is compactly supported in U_i , so that the integral on the right-hand side is obtained through Equation (2.14). There are finitely many non-zero integrals on the right

because the open cover of $\text{supp}(\omega)$ is finite. It turns out that Equation (2.15) neither depends on the choice of finite cover, nor on the choice of partition of unity (see Lemma 14.5 in [Lee, 2003]). The disconnected case requires to define an orientation on each connected component. In the following proposition are listed several properties of this integral:

Proposition 2.92. *Let M and N be oriented smooth n -dimensional manifolds, and let ω, η be compactly supported differential n -forms on M . Then:*

1. **Linearity:** for every $a, b \in \mathbb{R}$,

$$\int_M a\omega + b\eta = a \int_M \omega + b \int_M \eta$$

2. **Positivity:** if ω is positively oriented, then $\int_M \omega > 0$;

3. **Orientability:** If $F : N \rightarrow M$ is a diffeomorphism, then:

$$\int_M \omega = \begin{cases} \int_N F^* \omega & \text{if } F \text{ is orientation-preserving} \\ - \int_N F^* \omega & \text{if } F \text{ is orientation-reversing} \end{cases}$$

Remark 2.93. Equation (2.15) is still valid for non-compactly supported differential forms, but in that case the integral is improper since the sum on the right may not converge. On a compact manifold, the integral is defined for every differential n -form.

We conclude this section by briefly discussing two important results relying on integration of differential forms. We note $\Omega_c^p(M)$ the compactly supported differential p -forms on M . The de Rham differential applies to compactly supported differential forms and induces a cohomology, denoted $H_c^m(M)$. However, this cohomology is different than the de Rham cohomology, as the following observation shows:

Proposition 2.94. *The compactly supported de Rham cohomology of \mathbb{R}^n satisfies:*

$$H_c^i(\mathbb{R}^n) \simeq \begin{cases} \mathbb{R} & \text{if } i = n \\ 0 & \text{otherwise} \end{cases}$$

One notices that the m -th compactly supported cohomology group of compact support is isomorphic to the $n - m$ -th de Rham cohomology group. Does this extend to general manifolds? For every $0 \leq m \leq n$, notice that the integral defined in Equation (2.15) induces a linear morphism:

$$\begin{aligned} \mathcal{PD} : \Omega^m(M) &\longrightarrow \Omega_c^{n-m}(M)^* \\ \eta &\longmapsto \mathcal{PD}(\eta) : \mu \longmapsto \int_M \eta \wedge \mu \end{aligned}$$

Following de Rham, we call $n - m$ -currents the elements of $\Omega_c^{n-m}(M)^*$; they are related to the notion of distribution. The de Rham differential on compactly supported $n - m$ -forms induces a degree -1 differential $d' : \Omega_c^\bullet(M)^* \rightarrow \Omega_c^{\bullet-1}(M)^*$ defined by $d'\Phi(\mu) = \Phi((-1)^{|\Phi|} d\mu)$. Then, one can show that \mathcal{PD} commutes with the differentials: $\mathcal{PD}(d\eta)(\mu) = d'\mathcal{PD}(\eta)(\mu)$. This result implies that \mathcal{PD} induces a map at the cohomology level, which turns out to be an isomorphism:

Theorem 2.95. Poincaré duality *Let M be a smooth orientable manifold, then:*

$$H_{dR}^m(M) \simeq H_c^{n-m}(M)^*$$

Proof. See Exercise 16.6 in [Lee, 2003]. □

We conclude this section by mentioning Stokes' theorem. This result relies on the notion of *manifold with boundary*. We will not enter into the details of this notion, because it would take too much time, but many informations can be found in [Baez and Muniain, 1994] and [Lee, 2003]. The main idea is that a manifold with boundary is locally homeomorphic to the euclidean upper half-plane:

$$\mathbb{H}^n = \{(x^1, \dots, x^n) \in \mathbb{R}^n \mid x^n \geq 0\}$$

It means that a chart on a manifold with boundary M is either homeomorphic to an open subset (with respect to the subspace topology) of the interior of \mathbb{H}^n , or to an open subset of \mathbb{H}^n which intersects the boundary. The boundary of M is the set ∂M of points of M which are sent to the boundary $\partial\mathbb{H}^n$ of the upper half plane through the coordinate maps. By construction, the boundary of M is a closed embedded submanifold of M (see Exercise 8.5 in [Lee, 2003]). If the manifold M has an orientation, there is a distinguished orientation on its boundary ∂M . In that case, one can define integration of differential n -forms on M and at the same time define integration of differential $n - 1$ -forms on the boundary ∂M . Stokes' theorem is a statement about the relationship between those two integrals:

Theorem 2.96. Stokes' theorem *Let M be a smooth, oriented n -dimensional manifold with boundary, and let ω be a compactly supported smooth $n - 1$ -form on M . Then:*

$$\int_M d\omega = \int_{\partial M} \omega$$

The meaning of the integral on the right hand side and of this theorem are discussed in details in [Baez and Muniain, 1994] and [Lee, 2003].

3 Poisson geometry

Poisson geometry draws on the work of mathematicians in the 1960s-1970s striving to formalize Hamiltonian mechanics⁷. Recall that in Hamilton's formulation of classical mechanics, a physical system is characterized by a set of positions q^i and conjugate momenta p_i (where $1 \leq i \leq n$) defining a point in a *phase space* $P = \mathbb{R}^{2n}$, and the evolution of the system is governed by a function $H(q, p)$ called the *Hamiltonian*, so that Hamilton's equations are:

$$\dot{q}^i = \frac{\partial H}{\partial p_i} \quad \text{and} \quad \dot{p}_i = -\frac{\partial H}{\partial q^i} \quad (3.1)$$

for every $1 \leq i \leq n$. In this context, the classical Poisson bracket is a skew-symmetric differential operator $\{.,.\} : \mathcal{C}^\infty(P) \times \mathcal{C}^\infty(P) \longrightarrow \mathcal{C}^\infty(P)$ defined as on any two smooth functions $f, g \in \mathcal{C}^\infty(P)$ by:

$$\{f, g\} = \sum_{i=1}^n \frac{\partial f}{\partial q^i} \frac{\partial g}{\partial p_i} - \frac{\partial g}{\partial q^i} \frac{\partial f}{\partial p_i} \quad (3.2)$$

Using the Poisson bracket, Hamilton's equations (3.1) become:

$$\dot{q}^i = \{q^i, H\} \quad \text{and} \quad \dot{p}_i = \{p_i, H\}$$

for every $1 \leq i \leq n$, where q^i and p_i are considered to be the coordinate functions on P . Then, for every solution $\gamma : t \longmapsto (q^1(t), \dots, q^n(t), p_1(t), \dots, p_n(t))$ of the differential equations (3.1), one has:

$$\frac{d(f \circ \gamma)}{dt}(t) = \{f, H\}(\gamma(t))$$

for any smooth function $f \in \mathcal{C}^\infty(P)$. Then, the Hamiltonian defines a vector field $X_H = \{H, .\}$ on P , whose integral curves describe the time evolution of the physical system.

The Poisson bracket is central in Hamilton's description of classical mechanics: Poisson had already noticed that the set of functions which are invariant along the integral curves of X_H – the so-called *constants of motion* – is stable under Poisson bracket. Liouville then showed that the existence of a set of n independent constants of motion commuting under the Poisson bracket allows to integrate Hamilton's equations. This result was then later improved by the infamous *action-angle theorem* which, in the situation where the leaves of the constants of motion are compact, provides a distinguished choice of local coordinates which are such that the Hamiltonian takes a very specific and nice form. The Poisson bracket on \mathbb{R}^{2n} can be generalized to smooth manifolds and the aim of this chapter is to show that there are several deep mathematics that are raised by this new notion.

3.1 Poisson manifolds

Keeping in mind the correspondence between algebra and geometry, we first emphasize that Poisson geometry relies on the notion of *Poisson algebra*. Recall that every associative algebra (A, \cdot) gives rise to a Lie bracket:

$$[a, b] = a \cdot b - b \cdot a \quad (3.3)$$

In particular, because of the associativity, a short computation shows that this Lie bracket is a derivation of the associative product:

$$[a, b \cdot c] = [a, b] \cdot c + b \cdot [a, c] \quad (3.4)$$

⁷This section relies on four main sources: [Dufour and Zung, 2005], [Laurent-Gengoux et al., 2013] and [Vaisman, 1994], [Crainic et al., 2021] as well as these [lectures notes](#).

However, the right hand side of Equation (3.3) is trivial when the associative product is commutative, and hence the Lie bracket vanishes. Then, a non-trivial Lie bracket on such a commutative associative algebra should necessarily form exterior, additional data. A Poisson algebra is precisely such an object, where the commutative associative product is compatible with the Lie bracket so that they satisfy Equation (3.4):

Definition 3.1. A Poisson algebra is a \mathbb{R} -vector space A equipped with two bilinear products \cdot and $\{.,.\}$, such that:

1. (A, \cdot) is a commutative associative algebra;
2. $(A, \{.,.\})$ is a Lie algebra;
3. the Lie bracket is a derivation of the associative product:

$$\{a, b \cdot c\} = \{a, b\} \cdot c + b \cdot \{a, c\}$$

for any elements $a, b, c \in A$.

We call $\{.,.\}$ a Poisson bracket. A morphism of Poisson algebras is a map $\phi : A \longrightarrow B$ that is both a morphism of associative algebras and a morphism of Lie algebras.

Since a Poisson algebra has two main algebraic structures, there are several kinds of ideals and subalgebras: we need to carefully emphasize which product is used in their definitions. We use the denomination *ideal* and *subalgebra* when we refer to these algebraic structure defined with the help of the associative product, and we use the denomination *Lie ideal* and *Lie subalgebra* when we refer to the Poisson bracket. A Poisson ideal (resp. subalgebra) is an ideal (resp. subalgebra) with respect to the associative product *and* to the Poisson bracket. This notions will have a geometric counterpart when we study Poisson manifolds and their submanifolds.

Definition 3.2. A Poisson manifold is a smooth manifold M together with a \mathbb{R} -bilinear Lie bracket $\{.,.\}$ on the commutative associative algebra of smooth functions $\mathcal{C}^\infty(M)$ which makes it a Poisson algebra. The bracket $\{.,.\}$ is called a Poisson structure on M . A Poisson morphism between two Poisson manifolds M and N is a smooth map $\varphi : M \longrightarrow N$ such that the pullback $\varphi^* : \mathcal{C}^\infty(N) \longrightarrow \mathcal{C}^\infty(M)$ is a morphism of Poisson algebras.

Example 3.3. The phase space $P = \mathbb{R}^{2n}$, parametrized with the generalized coordinates q^i and their conjugate momenta p_i , together with the Poisson bracket defined in Equation (3.2). In that particular case, the Poisson bracket actually descends from the canonical symplectic structure $\omega = \sum_{i=1}^n dp_i \wedge dq^i$. Then the so-called canonical transformations correspond to Poisson isomorphisms (which in the present case coincide with *symplectomorphisms*).

Example 3.4. To every finite dimensional Lie algebra $(\mathfrak{g}, [.,.])$ we can associate a *linear Poisson structure* on \mathfrak{g}^* . Elements of \mathfrak{g} can then be seen as linear forms on the dual \mathfrak{g}^* : indeed, every $x \in \mathfrak{g}$ defines a linear map:

$$\begin{aligned} \bar{x} : \mathfrak{g}^* &\longrightarrow \mathbb{R} \\ \xi &\longmapsto \xi(x) \end{aligned}$$

Let (e_1, \dots, e_n) be a basis of \mathfrak{g} . Then by the above assignment they define a system of linear coordinates on \mathfrak{g}^* , denoted $\bar{e}_1, \dots, \bar{e}_n$. Every real analytic function on \mathfrak{g}^* can then be expressed in terms of such coordinates functions, and every smooth function on \mathfrak{g}^* can be differentiated with respect to these coordinates. In particular, the commutators $[e_i, e_j] = C_{ij}^k e_k$ define a

linear function $\eta_{ij} = \overline{[e_i, e_j]} = C_{ij}^k \overline{e_k}$ on \mathfrak{g}^* . These data allow to define a Poisson structure on the dual space \mathfrak{g}^* , called the *linear Poisson structure of \mathfrak{g}^** :

$$\{f, g\} = \sum_{1 \leq i, j \leq n} \eta_{ij} \frac{\partial f}{\partial \overline{e_i}} \frac{\partial g}{\partial \overline{e_j}}$$

for every $f, g \in C^\infty(\mathfrak{g}^*)$. It is the unique Poisson structure on \mathfrak{g}^* that satisfies the following identity:

$$\{\overline{x}, \overline{y}\} = \overline{[x, y]}$$

for every $x, y \in \mathfrak{g}$.

Example 3.5. Example 3.4 applies for example to $\mathfrak{so}(3)$: let e_1, e_2, e_3 its generators, and $[e_i, e_j] = \sum_{k=1}^3 \epsilon_{ijk} e_k$ be the Lie bracket, where ϵ_{ijk} is the Levi-Civita symbol on three elements. Denoting $X = \overline{e_1}$, $Y = \overline{e_2}$ and $Z = \overline{e_3}$, the corresponding linear Poisson structure on $\mathfrak{so}(3)^*$ satisfies:

$$\{X, Y\} = Z, \quad \{Y, Z\} = X, \quad \{Z, X\} = Y.$$

Example 3.6. One can change the former example to the following: instead of $\{X, Y\} = Z$, set $\{X, Y\} = -Z^2 + \frac{1}{4}$, and preserve the other two brackets. This choice defines a non-linear Poisson structure on $\mathfrak{so}(3)^* \simeq \mathbb{R}^3$.

Example 3.7. Example 3.4 extends to Lie algebroids (Definition 2.24): every Lie algebroid structure on a vector bundle A induces a Poisson manifold structure on A^* . In particular this implies that for every smooth manifold, T^*M is a Poisson manifold. If M is the configuration space associated to a given physical system, with local coordinates q^1, \dots, q^n , then the cotangent bundle T^*M is considered to be the associated phase space, admitting fiberwise local coordinates p_1, \dots, p_n , i.e. $p_k(dq^l) = \delta_k^l$. The Poisson bracket on T^*M is then the canonical one, defined in Equation (3.2).

Let us now deduce some properties of a given Poisson bracket $\{.,.\}$ on a Poisson manifold M . Vector fields on M which are derivations of the Poisson bracket are called *Poisson vector fields*. More precisely, such a vector field X satisfies the following identity:

$$X(\{f, g\}) = \{X(f), g\} + \{f, X(g)\} \quad (3.5)$$

for every $f, g \in C^\infty(M)$. Given a Poisson bracket, it is not straightforward to deduce which vector fields are Poisson vector fields. However, it turns out that a subclass of those are easily obtained. Recall from Remark 1.14 that to any element x of a Lie algebra one can associate a derivation, called the adjoint action of x , denoted $\text{ad}_x = [x, -]$. This remark applies to Poisson algebras, since they are particular cases of Lie algebras. In particular, let us study how this materializes in $C^\infty(M)$, when the smooth manifold M is a Poisson manifold.

Definition 3.8. For every $f \in C^\infty(M)$, we call $X_f = \text{ad}_f = \{f, .\}$ the Hamiltonian vector field associated to f . In particular, for any two smooth functions $f, g \in C^\infty(M)$ we have:

$$dg(X_f) = X_f(g) = \{f, g\} \quad (3.6)$$

Lemma 3.9. For any two smooth functions f, g on a Poisson manifold, we have that the Hamiltonian vector fields have the following nice property:

$$[X_f, X_g] = X_{\{f, g\}} \quad (3.7)$$

Proof. This is implied by Equation (3.6), together with the Jacobi identity of the Poisson bracket. \square

In other words the linear map $\mathcal{C}^\infty(M) \longrightarrow \mathfrak{X}(M)$ sending a smooth function to its hamiltonian vector field is a morphism of Lie algebras. Another useful application of Equation (3.6) is in showing that Hamiltonian vector fields are Poisson vector fields because the Jacobi identity for the Poisson bracket can be written as:

$$X_h(\{f, g\}) = \{X_h(f), g\} + \{f, X_h(g)\}$$

for every smooth functions f, g, h . A smooth function of $\mathcal{C}^\infty(M)$ whose hamiltonian vector field is zero is called a *Casimir element*, because it commutes with any other element of the algebra. When M is connected, constant functions on M are always Casimir elements. In Lie theoretic words, the space of Casimir elements corresponds to the center of the Lie algebra $(\mathcal{C}^\infty(M), \{., .\})$. There may then exist many linearly independent such objects.

Exercise 3.10. Show that the Poisson structure on \mathbb{R}^3 as defined in Example 3.6 indeed satisfies Equation (3.7).

Exercise 3.11. Show that the function $C = X^2 + Y^2 + Z^2$ is a Casimir element of the linear Poisson structure on $\mathfrak{so}(3)^*$. Turning to the non-linear structure defined in Example 3.6, check that it admits as a Casimir element (together with constant functions):

$$C = X^2 + Y^2 - \frac{2}{3}Z^3 + \frac{1}{2}Z$$

Given a Poisson bracket, it is not at all evident to deduce which functions are Casimir, and which vector fields are Poisson vector fields (up to hamiltonian vector fields). We will give a partial answer to this question using cohomological techniques. The mathematical machinery set up to describe this so called *Poisson cohomology* will eventually provide another, more geometric point of view on Poisson brackets. We first need to generalize the Lie bracket of vector fields to the whole graded algebra $\mathfrak{X}^\bullet(M) = \bigoplus_{i=1}^n \mathfrak{X}^i(M)$. Recall that the space $\mathfrak{X}^i(M)$ represents the sheaf of smooth sections of the vector bundle $\wedge^i TM$. Every multivector field is locally decomposable because $\wedge^i TM$ admits elements of the form $\partial_{k_1} \wedge \dots \wedge \partial_{k_i}$ (for $1 \leq k_1 < \dots < k_i \leq n$) as local frames. Evaluating an element of $\mathfrak{X}^i(M)$ on i smooth functions – which gives back another smooth function – is then done by using Equation (A.17). Moreover, the pair $(\mathfrak{X}^\bullet(M), \wedge)$ is a graded commutative algebra:

$$\mathfrak{X}^i(M) \wedge \mathfrak{X}^j(M) \subset \mathfrak{X}^{i+j}(M)$$

More precisely, for $P \in \mathfrak{X}^i(M)$ and $Q \in \mathfrak{X}^j(M)$ and $i+j$ smooth functions f_1, \dots, f_{i+j} , one has:

$$P \wedge Q(f_1, \dots, f_{i+j}) = \sum_{\sigma \in Un(i,j)} (-1)^\sigma P(f_{\sigma(1)}, \dots, f_{\sigma(i)}) Q(f_{\sigma(i+1)}, \dots, f_{\sigma(i+j)})$$

where $Un(p, n-p)$ represents the set of $(p, n-p)$ -*unshuffles* (other people call it *shuffles*), i.e. those permutations $\sigma \in S_n$ satisfying the following two unshuffling conditions:

$$\sigma(1) < \sigma(2) < \dots < \sigma(p) \quad \text{and} \quad \sigma(p+1) < \sigma(p+2) < \dots < \sigma(n-1) < \sigma(n)$$

At level 0, i.e. for $\mathfrak{X}^0(M) = \mathcal{C}^\infty(M)$, we understand that wedging with respect to a smooth function f consists in multiplying by this function: $f \wedge P = fP$, for any $P \in \mathfrak{X}^\bullet(M)$.

Example 3.12. Let $P = \partial_x \wedge \partial_y \wedge \partial_z$ be a multivector field on $M = \mathbb{R}^3$. For any three smooth functions f, g, h , we have:

$$P(f, g, h) = \partial_x f \partial_y g \partial_z h - \partial_x f \partial_y h \partial_z g + \circlearrowleft$$

where \circlearrowleft symbolizes circular permutation of the three functions.

While vector fields on a smooth manifold are derivations of smooth functions, multivector fields are multiderivations: $\mathfrak{X}^i(M) \simeq \text{Der}^i(\mathcal{C}^\infty(M))$ (see Lemma 1.2.2 in [Dufour and Zung, 2005]). By multiderivation, we mean the following: for every $P \in \mathfrak{X}^i(M)$ and $f_1, \dots, f_i, g \in \mathcal{C}^\infty(M)$, we have:

$$P(f_1, \dots, f_i g) = P(f_1, \dots, f_i) g + f_i P(f_1, \dots, g)$$

In particular, since P is fully skew-symmetric with respect to permutations of its variables, the derivation property is true for every slot. Multiderivations can be composed: for $P \in \mathfrak{X}^i(M)$ and $Q \in \mathfrak{X}^j(M)$ two multiderivations (here $1 \leq i, j \leq n$), the composite $P \circ Q$ is not a multiderivation, but a priori no more than a multi-operator on $\mathcal{C}^\infty(M)$ (this was already observed for mere vector fields). More precisely, it acts on $i + j - 1$ smooth functions f_1, \dots, f_{i+j-1} as:

$$P \circ Q(f_1, \dots, f_{i+j-1}) = \sum_{\sigma \in Un(j, i-1)} (-1)^\sigma P(Q(f_{\sigma(1)}, \dots, f_{\sigma(j)}), f_{\sigma(j+1)}, \dots, f_{\sigma(i+j-1)}) \quad (3.8)$$

while if Q is a smooth function (whatever P is) then we set $P \circ Q = 0$. Equation (3.8) shows that although P has degree i and Q has degree j , the composite $P \circ Q$ does not respect this graduation because it has $i + j - 1$ arguments. This is why we decide to create a new grading on $\mathfrak{X}^\bullet(M)$, by shifting the original grading by -1 . We denote by $\mathcal{V}^i(M)$ (for $-1 \leq i \leq n-1$) the vector space $\mathfrak{X}^{i+1}(M)$ shifted by a degree -1 :

$$\mathcal{V}^i(M) = \mathfrak{X}^{i+1}(M)$$

In other words, we have the following correspondence:

$$\begin{array}{ccccccc} \mathcal{V}^{-1}(M) & \mathcal{V}^0(M) & \mathcal{V}^1(M) & \dots & \mathcal{V}^{n-1}(M) \\ \parallel & \parallel & \parallel & & \parallel \\ \underbrace{\mathfrak{X}^0(M)}_{\mathcal{C}^\infty(M)} & \mathfrak{X}^1(M) & \mathfrak{X}^2(M) & \dots & \mathfrak{X}^n(M) \end{array}$$

In particular, smooth functions now belong to $\mathcal{V}^{-1}(M) = \mathfrak{X}^0(M)$, vector fields belong to $\mathcal{V}^0(M) = \mathfrak{X}^1(M)$, and multivector fields of degree i belong to $\mathcal{V}^{i-1}(M)$. We label by \bar{P} the degree (with respect to the new convention, in $\mathcal{V}^\bullet(M)$) of the homogeneous element P . In particular, if $P \in \mathfrak{X}^i(M)$, we have $\bar{P} = i - 1$. Given these conventions, we set:

Definition 3.13. *The Schouten-Nijenhuis bracket is the \mathbb{R} -bilinear graded skew-symmetric bracket on $\mathcal{V}^\bullet(M) = \bigoplus_{i=-1}^{n-1} \mathcal{V}^i(M)$ defined by its action on any two homogeneous multivector fields P, Q of degree ≥ 0 :*

$$[P, Q]_{SN} = P \circ Q - (-1)^{\bar{P} \cdot \bar{Q}} Q \circ P \quad (3.9)$$

while, for any function $f \in \mathcal{C}^\infty(M)$:

$$[P, f]_{SN} = P(f, \dots) \quad (3.10)$$

Remark 3.14. Equation (3.10) means that, if for example in local coordinates $P = P^{i_1 \dots i_k} \partial_{x^{i_1}} \wedge \dots \wedge \partial_{x^{i_k}}$ (summation implied), then:

$$[P, f]_{SN} = \sum_{i_1, \dots, i_k} (-1)^{j+1} P^{i_1 \dots i_k} \partial_{x^{i_j}}(f) \partial_{x^{i_1}} \wedge \dots \wedge \widehat{\partial_{x^{i_j}}} \wedge \dots \wedge \partial_{x^{i_k}}$$

where $\widehat{\partial_{x^{i_j}}}$ means that we omit this term in the wedge product of $k - 1$ partial derivatives.

Being *graded skew-symmetric* means that for any two homogeneous multivector fields P, Q , one has:

$$[P, Q]_{SN} = -(-1)^{\overline{P}\overline{Q}}[Q, P]_{SN} \quad (3.11)$$

This implies in particular that, when one considers P, Q as elements of $\mathfrak{X}^\bullet(M)$ of respective degrees i and j , the Schouten-Nijenhuis bracket reads:

$$[P, Q]_{SN} = -(-1)^{(i-1)(j-1)}[Q, P]_{SN}$$

We see that this definition of the bracket – although equivalent to Equation (3.9) – is not so convenient because of the exponents that do not match the degrees of P and Q . For degree reasons, the bracket of two functions is zero because the sum of their degrees is -2 , and the graded vector space $\mathcal{V}^\bullet(M)$ does not possess a vector space of degree -2 .

A more explicit formula for Equation (3.9) when P and Q are decomposable multivector fields, may be the following:

$$[X_1 \wedge \dots \wedge X_i, Y_1 \wedge \dots \wedge Y_j]_{SN} = \sum_{\substack{1 \leq k \leq i \\ 1 \leq l \leq j}} (-1)^{k+l} [X_k, Y_l] \wedge X_1 \wedge \dots \wedge \widehat{X}_k \wedge \dots \wedge X_i \wedge Y_1 \wedge \dots \wedge \widehat{Y}_l \wedge \dots \wedge Y_j \quad (3.12)$$

together with, for Equation (3.10):

$$[X_1 \wedge \dots \wedge X_i, f]_{SN} = \sum_{k=1}^i (-1)^{k+1} X_i(f) X_1 \wedge \dots \wedge \widehat{X}_k \wedge \dots \wedge X_i$$

for every vector fields $X_1, \dots, X_i, Y_1, \dots, Y_j$, and smooth function $f \in \mathcal{C}^\infty(M)$. The latter expression is convenient because we then have:

$$[X, f]_{SN} = X(f) \quad \text{and} \quad [X \wedge Y, f]_{SN} = X(f)Y - Y(f)X$$

Exercise 3.15. Using Equation (A.17), you may check the identity between formula (3.9) and (3.12) on small decomposable multivector fields, such as $P = X$ and $Q = Y_1 \wedge Y_2$.

The Schouten-Nijenhuis bracket has several nice properties. In particular it coincides with the Lie bracket on vector fields when $P, Q \in \mathfrak{X}^1(M)$, since in that case $\overline{P} = \overline{Q} = 0$. It is thus legitimate to wonder whether this bracket generalizes the notion of Lie algebra to that of a *graded Lie algebra* on the graded vector space $\mathcal{V}^\bullet(M)$.

Definition 3.16. A graded Lie algebra is a graded vector space $V = \bigoplus_{i \in \mathbb{Z}} V_i$, equipped with a \mathbb{R} -bilinear operation $[\cdot, \cdot] : V \times V \rightarrow V$ called a graded Lie bracket, and which satisfies the following identities:

$$\begin{aligned} \text{graded skew-symmetry} & \quad [x, y] = -(-1)^{|x||y|}[y, x] \\ \text{graded Jacobi identity} & \quad [x, [y, z]] = [[x, y], z] + (-1)^{|x||y|}[y, [x, z]] \end{aligned}$$

for every $x, y, z \in \mathfrak{V}$. A derivation of degree d of V is an endomorphism $\delta : V^\bullet \rightarrow V^{\bullet+d}$ such that:

$$\delta([x, y]) = [\delta(x), y] + (-1)^{|x|d}[x, \delta(y)]$$

We denote $\text{Der}^d(V)$ the vector space of derivations of degree d of V .

Remark 3.17. Another, more symmetric form of the graded Jacobi identity exists, but it is not very convenient to use:

$$(-1)^{|x||z|}[x, [y, z]] + (-1)^{|y||x|}[y, [z, x]] + (-1)^{|z||y|}[z, [x, y]] = 0$$

Moreover, the graded Jacobi identity appearing in Definition 3.16 shows that the adjoint action of any element of V is a derivation of V : $\text{ad}_x \in \text{Der}^{|x|}(V)$, for every $x \in V$.

The two conditions satisfied by the graded Lie bracket in Definition 3.16 are slight generalizations of what characterizes a Lie algebra because Lie algebras are graded Lie algebras concentrated in degree 0. The idea with grading is very intuitive: for any two homogeneous elements $x, y \in E$, when we swap x and y to form a new term (either in the bracket or by ‘jumping’ over), we add a sign $(-1)^{|x||y|}$ in front of the new term hence created, compared to the classical (non-graded) situation. You can see this phenomenon on the right-hand sides of both equations. The same phenomenon happens for differential forms: $\eta \wedge \mu = (-1)^{|\eta||\mu|} \mu \wedge \eta$, making the wedge product *graded commutative* (and not merely commutative). Here, this has interesting consequences: in a graded Lie algebra, we do not necessarily have $[x, x] = 0$ because the graded Lie bracket is not skew-symmetric anymore when $|x|$ is odd because $[x, x] = -(-1)^{1 \times 1} [x, x]$ hence we cannot conclude on the vanishing of $[x, x]$.

Exercise 3.18. Using the graded Jacobi identity, show that if $x = y$ and $|x|$ is odd, we have:

$$[x, [x, z]] = \frac{1}{2} [[x, x], z]$$

These observations enable us to formulate the following important result:

Proposition 3.19. *The Schouten-Nijenhuis bracket extends the Lie bracket of vector fields to a graded Lie algebra structure on $\mathcal{V}^\bullet(M)$. In particular, the Schouten-Nijenhuis bracket satisfies:*

$$[P, fQ]_{SN} = [P, f]_{SN}Q + f[P, Q]_{SN}$$

for any smooth function f and multivector fields P and Q .

Proof. First of all, one needs to check that the graded vector space $\mathcal{V}^\bullet(M)$ is stable under this bracket. This can be proven on decomposable vector fields, using Equation (3.12). From Equation (3.11), the bracket is obviously graded skew-symmetric. It is just a matter of computation to check with Equation (3.12) that it satisfies the graded Jacobi identity (on decomposable vector fields). See Theorem 1.8.1 in [Dufour and Zung, 2005] for more details. \square

Remark 3.20. Actually, the Schouten-Nijenhuis bracket is the unique extension of the Lie bracket of vector fields to a graded Lie bracket on the space of alternating multivector fields that makes it into a Gerstenhaber algebra.

The Schouten-Nijenhuis bracket on vector fields allows us to characterize Poisson structures in a more geometric flavored approach. Let $(M, \{ \dots \})$ be a Poisson manifold and let x^1, \dots, x^n be local coordinates on M . Then the Poisson bracket between two functions f, g is locally of the form (see Proposition 1.14 in [Fernandes, 2005]):

$$\{f, g\} = \sum_{1 \leq i, j \leq n} \{x^i, x^j\} \frac{\partial f}{\partial x^i} \frac{\partial g}{\partial x^j} \quad (3.13)$$

This Equation is valid locally in the coordinate neighborhood of any point of a smooth manifold, and it is invariant under change of coordinates $x^i \mapsto x'^k$. Indeed, using Equations (2.7)-(2.8), we have that (omitting the sum signs):

$$\{x^i, x^j\} \frac{\partial f}{\partial x^i} \frac{\partial g}{\partial x^j} = \{x^i, x^j\} \frac{\partial x'^k}{\partial x^i} \frac{\partial x'^l}{\partial x^j} \frac{\partial f}{\partial x'^k} \frac{\partial g}{\partial x'^l} = \{x'^k, x'^l\} \frac{\partial f}{\partial x'^k} \frac{\partial g}{\partial x'^l}$$

Then we see from Equation (3.13) that the Poisson bracket can be locally seen as a bivector field $\frac{1}{2} \{x^i, x^j\} \partial_i \wedge \partial_j$ which, when evaluated on two smooth functions f, g , give $\{f, g\}$, as the following short calculation (where we have omitted the sum signs) shows:

$$\frac{1}{2} \{x^i, x^j\} \partial_i \wedge \partial_j (f, g) = \frac{1}{2} \{x^i, x^j\} (\partial_i f \partial_j g - \partial_i g \partial_j f) = \{x^i, x^j\} \frac{\partial f}{\partial x^i} \frac{\partial g}{\partial x^j}$$

We used Equation (A.18) between the first and the second step. We denote π the unique bivector field whose component in local coordinates is $\pi^{ij} = \{x^i, x^j\}$. Thus, the Poisson bracket uniquely defines a bivector field $\pi \in \mathfrak{X}^2(M)$ via the following identity:

$$\pi(f, g) = \{f, g\} \quad (3.14)$$

Exercise 3.21. By applying the Jacobi identity satisfied by the Poisson bracket to the coordinate functions x^i, x^j, x^k , show that the components of the bivector field π satisfies (this is a local expression):

$$\sum_{s=1}^n \pi^{is} \frac{\partial \pi^{jk}}{\partial x^s} + \pi^{js} \frac{\partial \pi^{ki}}{\partial x^s} + \pi^{ks} \frac{\partial \pi^{ij}}{\partial x^s} = 0 \quad (3.15)$$

Obviously, not every bivector field satisfies Equation (3.15). However, those that satisfy it define a Poisson structure on M via Equation (3.14). This translates as the following fundamental fact, due to Lichnerowicz:

Proposition 3.22. *There is a one-to-one correspondence between Poisson structures on a smooth manifold M and bivector fields $\pi \in \mathfrak{X}^2(M)$ such that:*

$$[\pi, \pi]_{SN} = 0 \quad (3.16)$$

Exercise 3.23. Prove that $[\pi, \pi]_{SN} = 0$ is equivalent to Equation (3.15), when evaluated in local coordinates.

Remark 3.24. By the correspondence established by Proposition (3.22), we will now either use the notation $(M, \{., .\})$ or (M, π) (depending on the context) to denote a Poisson manifold.

Example 3.25. The bivector field associated to the canonical Poisson bracket of Example 3.3 is the following:

$$\pi = \sum_{i=1}^n \frac{\partial}{\partial q^i} \wedge \frac{\partial}{\partial p_i} \quad (3.17)$$

So, in particular, if one relabels the coordinates as $x^i = q^i$ and $x^{n+i} = p_i$ for $1 \leq i \leq n$, then $\pi^{ij}(q, p) = 0$ except when $j = i + n$ or $i + n = j$, and in that case we have $\pi^{i(i+n)} = 1$ and $\pi^{(i+n)n} = -1$, so that $\pi = \frac{1}{2} \pi^{ij} \partial_i \wedge \partial_j$.

Example 3.26. On \mathbb{R}^3 one picks the following Poisson bracket:

$$\{f, g\} = x \frac{\partial f}{\partial x} \frac{\partial g}{\partial z} + y \frac{\partial f}{\partial y} \frac{\partial g}{\partial z} - x \frac{\partial g}{\partial x} \frac{\partial f}{\partial z} - y \frac{\partial g}{\partial y} \frac{\partial f}{\partial z}$$

This Poisson bracket corresponds to the following Poisson bivector field:

$$\pi = \left(x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} \right) \wedge \frac{\partial}{\partial z}$$

Exercise 3.27. Check that the Poisson bivector defined in Example 3.25 indeed satisfies Equation (3.14), where the Poisson bracket is that of Equation (3.2).

Exercise 3.28. Given two Poisson structures π_0 and π_1 on a smooth manifold M , show that, if $\pi_t = (1-t)\pi_0 + t\pi_1$ is a Poisson structure for some $t \neq 0, 1$, then it is a Poisson structure for all $t \in \mathbb{R}$. We then call the smooth family $(\pi_t)_t$ a *Poisson pencil*.

Seen from $\mathfrak{X}^\bullet(M)$, bivectors fields have degree 2, while seen from $\mathcal{V}^\bullet(M)$, they have degree 1. Bivector fields satisfying Equation (3.16) are called *Poisson bivector fields* (not to be confused with the Poisson vector fields defined in Equation (3.5)). We now show how a Poisson bivector

field makes $\mathcal{V}^\bullet(M)$ a chain complex. Let $d_\pi : \mathcal{V}^\bullet(M) \longrightarrow \mathcal{V}^{\bullet+1}(M)$ be the unique \mathbb{R} -linear morphism defined on any element $P \in \mathcal{V}^\bullet(M)$ as:

$$d_\pi(P) = [\pi, P]_{SN} \quad (3.18)$$

This operator is well defined, and is indeed of degree 1; it corresponds to the adjoint action of π on the graded Lie algebra $\mathcal{V}^\bullet(M) = \bigoplus_{i=-1}^{n-1} \mathcal{V}^i(M)$. Moreover, the graded Jacobi identity (via Exercise 3.18) together with Equation (3.16) imply that d_π squares to zero:

$$d_\pi^2(P) = [\pi, [\pi, P]_{SN}]_{SN} = \frac{1}{2} [[\pi, \pi]_{SN}, P]_{SN} = 0$$

This operator is often called the *Poisson differential*. These successive facts imply that $(\mathcal{V}^\bullet(M), d_\pi)$ is a chain complex:

$$0 \longrightarrow \underbrace{\mathcal{V}^{-1}(M)}_{\mathcal{C}^\infty(M)} \xrightarrow{d_\pi} \underbrace{\mathcal{V}^1(M)}_{\mathfrak{X}(M)} \xrightarrow{d_\pi} \mathcal{V}^1(M) \xrightarrow{d_\pi} \dots \xrightarrow{d_\pi} \mathcal{V}^{n-1}(M) \longrightarrow 0$$

Notice that the above results can equivalently be expressed with respect to the grading on $\mathfrak{X}^\bullet(M)$. The Poisson differential is still defined from Equation (3.18), at the cost of expressing the Schouten-Nijenhuis with respect to the grading of $\mathfrak{X}^\bullet(M)$, via Equation (3.11). Then, we obtain that $(\mathfrak{X}^\bullet(M), d_\pi)$ is a chain complex concentrated in degrees $0, \dots, n$. This is the natural setup to define a cohomology theory:

Definition 3.29. *Let (M, π) be a Poisson manifold. The cohomology of the chain complex $(\mathfrak{X}(M), d_\pi)$ is called the Poisson cohomology of (M, π) and is denoted, for $0 \leq i \leq n$:*

$$H_\pi^i(M) = \frac{\text{Ker}(d_\pi : \mathfrak{X}^i(M) \longrightarrow \mathfrak{X}^{i+1}(M))}{\text{Im}(d_\pi : \mathfrak{X}^{i-1}(M) \longrightarrow \mathfrak{X}^i(M))}$$

The map d_π is called the Poisson differential.

Notice that Equation (3.18) can be rewritten in a way that shows a huge similarity with de Rham differential (1.31) (and Chevalley-Eilenberg differential as well):

$$\begin{aligned} (-1)^{m-1} d_\pi(P)(f_1, \dots, f_m, f_{m+1}) &= \sum_{i=1}^{m+1} (-1)^{i-1} X_{f_i}(P(f_1, \dots, \widehat{f}_i, \dots, f_{m+1})) \\ &+ \sum_{1 \leq i < j \leq m+1} (-1)^{i+j} P(\{f_i, f_j\}, f_1, \dots, \widehat{f}_i, \dots, \widehat{f}_j, \dots, f_{m+1}) \end{aligned} \quad (3.19)$$

for every $P \in \mathfrak{X}^m(M)$. The sign $(-1)^m$ could have been got rid of if the Poisson differential had been defined following the alternative, although equivalent, convention: $d_\pi(P) = -[P, \pi]_{SN}$. As can be explicitly be seen in Equation (3.19), the Poisson differential carries information on the Poisson structure on M . The next subsection clarifies the relationship between de Rham cohomology and Poisson cohomology.

Exercise 3.30. Show that the map d_π is a derivation of the Schouten-Nijenhuis bracket.

Let us compute the first few cohomology groups. The 0-th Poisson cohomology is given by:

$$H_\pi^0(M) = \text{Ker}(d_\pi : \mathcal{C}^\infty(M) \longrightarrow \mathfrak{X}(M))$$

By definition of the Schouten-Nijenhuis bracket involving functions, we have $d_\pi(f) = [\pi, f]_{SN} = -\pi(f, -) = -\{f, -\}$. Then, the smooth functions that belong to $H_\pi^0(M)$ consists of those that are such that $\{f, g\} = 0$, i.e. they are Casimir elements of the Lie algebra $(\mathcal{C}^\infty(M), \{., .\})$:

$$H_\pi^0(M) = \text{Casimir elements of } (\mathcal{C}^\infty(M), \{., .\})$$

The dimension of $H_\pi^0(M)$ as a vector space is at least 1, because constant functions on M are Casimir elements (assuming M is connected). Going to the next level, we have:

$$H_\pi^1(M) = \frac{\text{Ker}(d_\pi : \mathfrak{X}(M) \longrightarrow \mathfrak{X}^2(M))}{\text{Im}(d_\pi : \mathcal{C}^\infty(M) \longrightarrow \mathfrak{X}(M))}$$

Elements of $\text{Ker}(d_\pi : \mathfrak{X}(M) \longrightarrow \mathfrak{X}^2(M))$ are characterized by the following property: they are vector fields X such that $[\pi, X]_{SN} = 0$. It corresponds to Equation (3.5), hence such vector fields are Poisson vector fields. The space $\text{Im}(d_\pi : \mathcal{C}^\infty(M) \longrightarrow \mathfrak{X}(M))$ consists of Hamiltonian vector fields on M . These are Poisson vector fields of a particular kind: they are somehow "trivial" in the sense that are the easiest Poisson vector fields to find, for they are automatically given as soon as a Poisson structure is defined. The interesting Poisson vector fields are thus those that are not hamiltonian or, more precisely, the classes of Poisson vector fields up to hamiltonian vector fields, which is precisely what the first cohomology group is:

$$H_\pi^1(M) = \text{classes of non-trivial Poisson vector fields}$$

So, in particular, if $H_\pi^1(M) \neq 0$ there are Poisson vector fields which are not Hamiltonian vector fields. Higher Poisson cohomology groups arise naturally in deformation theory: $H_\pi^2(M)$ may be interpreted as the moduli space of formal infinitesimal deformations of π , while $H_\pi^3(M)$ may be interpreted as the space of obstructions of such deformations [Dufour and Zung, 2005].

Example 3.31. Using the fact that $H_{dR}^1(\mathbb{R}^{2n}) = 0$, we will show in Remark 3.38 that $H_\pi^1(\mathbb{R}^{2n}) = 0$ (where the Poisson structure is the canonical one). It implies that on \mathbb{R}^{2n} equipped with the Poisson bracket defined in Equation (3.2), every Poisson vector field is hamiltonian, i.e. descends from a smooth function. However, contrary to de Rham cohomology which is locally trivial on a smooth manifold, Poisson cohomology needs not be locally trivial on a Poisson manifold because the Poisson structure needs not be non-degenerate.

An alternative view on Poisson vector fields can be made through *Lie derivatives*. First, define the Lie derivative of a vector field Y along the vector field X by the Lie bracket:

$$\mathcal{L}_X(Y) = [X, Y] \tag{3.20}$$

Then, to be consistent with Schouten-Nijenhuis bracket, it implies that on smooth functions, Lie derivatives act as derivations:

$$\mathcal{L}_X(f) = [X, f]_{SN} = X(f) \tag{3.21}$$

More generally, on any multivector field P , the Lie derivative acts as:

$$\mathcal{L}_X(P) = [X, P]_{SN}$$

Then, one notices that the condition (3.5) of X being a Poisson vector field (with respect to the Poisson bivector π) is equivalent to the following equality:

$$\mathcal{L}_X(\pi) = 0$$

Exercise 3.32. Show that the condition that a bivector field B is a Poisson bivector is equivalent to the following identity:

$$\mathcal{L}_{B\sharp(df)}(B) = 0 \quad \text{for every smooth function } f$$

To conclude this subsection, we show that the Lie derivative can naturally act on differential forms. Using Equations (3.20) and (3.21), one can deduce that the Lie derivative \mathcal{L}_X of a vector field X acts on differential one-forms since it should satisfy a kind of ‘derivation property’:

$$X(\xi(Y)) = \mathcal{L}_X(\xi)(Y) + \xi(\mathcal{L}_X(Y)) \quad (3.22)$$

for every $\xi \in \Omega^1(M)$ and $Y \in \mathfrak{X}(M)$. Defining the *interior product on differential forms* as the linear operator defined, for every $x \in \mathfrak{X}(M)$, as:

$$\begin{aligned} \iota_X : \Omega^\bullet(M) &\longrightarrow \Omega^{\bullet-1}(M) \\ \eta &\longmapsto \iota_X \eta = \eta(X, \dots) \end{aligned}$$

we notice that Equation (3.22) is equivalent to writing:

$$\mathcal{L}_X(\xi)(Y) = X(\xi(Y)) - \xi([X, Y]) = Y(\xi(X)) + d\xi(X, Y) = (d\iota_X(\xi) + \iota_X d\xi)(Y)$$

This allows us to find the following characterization:

Definition 3.33. *The Lie derivative of a vector field $X \in \mathfrak{X}(M)$ is the unique derivation of both $\mathfrak{X}^\bullet(M)$ and $\Omega^\bullet(M)$, defined on any multivector field P and differential form η as:*

$$\begin{aligned} \mathcal{L}_X(P) &= [X, P]_{SN} \\ \mathcal{L}_X(\eta) &= d\iota_X(\eta) + \iota_X d\eta \end{aligned}$$

Then, we have a nice result involving Lie derivatives, which is often used in geometry:

Proposition 3.34. *For any two vector fields $X, Y \in \mathfrak{X}(M)$, one has:*

$$[\mathcal{L}_X, \mathcal{L}_Y] = \mathcal{L}_{[X, Y]}$$

Proof. This is a direct consequence of the properties of the Schouten-Nijenhuis bracket or the operators d and ι_X . \square

3.2 Properties of Poisson bivectors

Let M be a smooth manifold and let $B = \frac{1}{2}b^{ij}\partial_i \wedge \partial_j$ be a (non-necessarily Poisson) bivector field. The local components $(b^{ij})_{1 \leq i, j \leq n}$ are smooth functions $b^{ij} : x \longrightarrow b^{ij}(x)$ on U . For simplicity, the latter expression $b^{ij}(x)$ will be denoted b_x^{ij} . Such functions define a smooth function:

$$\begin{aligned} \bar{B} : U &\longrightarrow \mathfrak{gl}_n(\mathbb{R}) \\ x &\longmapsto \begin{pmatrix} b_x^{11} & b_x^{12} & \dots & b_x^{1n} \\ b_x^{21} & & & b_x^{2n} \\ \dots & & & \dots \\ b_x^{n1} & \dots & \dots & b_x^{nn} \end{pmatrix} \end{aligned}$$

where $n = \dim(M)$. Because $b^{ij} = -b^{ji}$, this function takes values in skew-symmetric $n \times n$ matrices. We say that the matrix is a *contravariant tensor* because indices appear as exponents

and any change of coordinates induce a transformation of the functions b^{ij} following that of the components of vector fields: $b^{kl} = \frac{\partial x^k}{\partial x^i} \frac{\partial x^l}{\partial x^j} b^{ij}$ (summation is implicit). The *rank* of B at a point x is the rank of the corresponding matrix $\overline{B}(x)$. The rank is obviously invariant under coordinate change. We say that a bivector field is *non-degenerate* when it has maximal rank n at every point of M . Since the rank of an anti-symmetric matrix is even, it means that such situation can occur only when M is even dimensional.

To make further sense of \overline{B} , let us introduce the notation for the natural pairing between differential two-forms and bivector fields on M . Indeed, the pairing $\langle \cdot, \cdot \rangle$ between covector and tangent vectors on M can be extended to decomposable differential 2-forms and bivector fields by the following identity:

$$\langle \xi \wedge \eta, X \wedge Y \rangle = 2(\xi(X)\eta(Y) - \xi(Y)\eta(X)) \quad (3.23)$$

The factor 2 comes from the fact that we have two wedges products on the left-hand side, compared to Equation (A.17) for example. These conventions imply that, for a bivector B , we have:

$$B(f, g) = \frac{1}{2} \langle df \wedge dg, B \rangle \quad (3.24)$$

for any two smooth functions $f, g \in C^\infty(M)$. Indeed, in local coordinates the bivector field B reads $B = \frac{1}{2} b^{ij} \partial_i \wedge \partial_j$. Then, applying Equation (3.23) to the right hand side of Equation (3.24) gives (summation on repeated indices is implicit):

$$\frac{1}{2} \langle df \wedge dg, B \rangle = \frac{1}{4} b^{ij} \langle df \wedge dg, \partial_i \wedge \partial_j \rangle = \frac{1}{2} b^{ij} (\partial_i(f)\partial_j(g) - \partial_j(f)\partial_i(g)) = \frac{1}{2} b^{ij} \partial_i \wedge \partial_j(f, g) = B(f, g)$$

Remark 3.35. Equation (3.24) straightforwardly applies to $B = \pi$ a Poisson bivector field, although in the litterature the left hand side is often written $\pi(df, dg)$ (not to be confused with $\langle df \wedge dg, B \rangle$ then).

Using the pairing between differential two-forms and bivector fields defined in Equation (3.23), the bivector field B induces a vector bundle morphism:

$$\begin{aligned} B^\# : T^*M &\longrightarrow TM \\ (x, \xi_x) &\longmapsto \frac{1}{2} \langle \xi_x \wedge d(-), B_x \rangle \end{aligned}$$

where B_x denotes the evaluation of the bivector field B at x . The term on the right-hand side should be read as follows:

$$\frac{1}{2} \langle \xi_x \wedge d(-), B_x \rangle : f \longmapsto \frac{1}{2} \langle \xi_x \wedge df|_x, B_x \rangle$$

One can check that it is indeed a derivation of smooth functions. More generally, evaluating any differential form η on $B^\#(\xi)$ corresponds to the following pairing:

$$\eta(B^\#(\xi)) = \frac{1}{2} \langle \xi \wedge \eta, B \rangle \quad (3.25)$$

The definition of $B^\#$ has been made so that, when evaluated on exact differential forms (every sufficiently local section of T^*M is exact), it is the unique vector bundle morphism satisfying:

$$B^\#(df) = B(f, -)$$

The right hand-side is a vector field on M (or at least an open set U), so the smooth map $B^\#$ indeed takes values in the tangent bundle and defines a vector bundle morphism. Since in

local coordinates, the right hand side of Equation (3.25) reads $\xi_{i,x} b_x^{ij} \partial_j(f)$ – where the $\xi_{i,x}$ are the components of the covector ξ_x in the basis dx^1, \dots, dx^n – the rank of the map $B^\#$ is the rank of the map $\bar{B} : M \rightarrow \mathfrak{gl}_n(\mathbb{R})$. The morphism $B^\#$ extends as a vector bundle morphism $\wedge^i T^*M \rightarrow \wedge^i TM$, for $1 \leq i \leq n$, compatible with the wedge product. It means that, setting the action of $B^\#$ on $\Omega^0(M) = \mathfrak{X}^0(M) = C^\infty(M)$ to be the identity map, $B^\#$ extends to a morphism of graded commutative algebras $B^\# : \Omega^\bullet(M) \rightarrow \mathfrak{X}^\bullet(M)$:

$$B^\#(\eta \wedge \mu) = B^\#(\eta) \wedge B^\#(\mu)$$

This perspective on bivector fields is quite useful regarding the relationship between de Rham cohomology and Poisson cohomology. Indeed, if $B = \pi$ is a Poisson bivector field, then the Hamiltonian vector field associated to the smooth function f is precisely $X_f = \pi^\#(df)$. More generally, the vector bundle morphism $\pi^\# : \wedge^i T^*M \rightarrow \wedge^i TM$ commutes with the respective differentials:

Proposition 3.36. *The graded commutative algebras morphism $\pi^\# : \Omega^\bullet(M) \rightarrow \mathfrak{X}^\bullet(M)$ is a chain map:*

$$\pi^\# \circ d_{\text{dR}} = d_\pi \circ \pi^\#$$

The proof is made by induction on the form degree, and can be found in Proposition 2.1.3 in [Dufour and Zung, 2005]. Then, closed (resp. exact) differential form are sent to closed (resp. exact) multivector fields. The chain map $\pi^\# : (\Omega^\bullet(M), d_{\text{dR}}) \rightarrow (\mathfrak{X}^\bullet(M), d_\pi)$ is an algebra homomorphism, and then induces a homomorphism between cohomology groups that we denote $\pi^\#$ as well:

Corollary 3.37. *For any Poisson manifold (M, π) , there is a natural homomorphism, called the Lichnerowicz homomorphism, between the de Rham cohomology and the Poisson cohomology:*

$$\pi^\# : H_{\text{dR}}^\bullet(M) \rightarrow H_\pi^\bullet(M)$$

Remark 3.38. Then, if π is a non-degenerate bivector field, the Lichnerowicz homomorphism is an isomorphism. This Corollary proves that $H_\pi^1(\mathbb{R}^{2n}) = 0$ so that every Poisson vector field on \mathbb{R}^{2n} (equipped with the standard Poisson bracket) is a Hamiltonian vector field, and that this Poisson structure is ‘rigid’ in the sense that $H_\pi^2(\mathbb{R}^{2n}) = 0$.

The importance of the vector bundle morphism $\pi^\# : T^*M \rightarrow TM$ is the following: for every $x \in M$ its image in $T_x M$ spans the directions taken by the hamiltonian vector fields at x . As it is, this might be useless, but actually it allows us to understand that hamiltonian vector fields do not necessarily span the entire tangent space, and thus that the transport along these vector fields are constraints in some directions. Hence, for a physical hamiltonian, it means that the Poisson structure on M *constraints* the set of reachable points in the phase space, given an initial point. In particular, if the Poisson bivector is degenerate at a point x , there is no bijection between T_x^*M and $T_x M$ and the integral curves of Hamiltonian vector fields passing through x will not be able to reach every point in the neighborhood of x . This can be explained by the fact that the vector bundle morphism $\pi^\#$ defines an integrable distribution, as the following proposition shows:

Proposition 3.39. *Let (M, π) be a Poisson manifold. Then T^*M is Lie algebroid – called the cotangent Lie algebroid, with anchor $\pi^\# : T^*M \rightarrow TM$ and with Lie bracket:*

$$\begin{aligned} [\cdot, \cdot]_{T^*M} : \Omega^1(M) \times \Omega^1(M) &\longrightarrow \Omega^1(M) \\ (\alpha, \beta) &\longmapsto [\alpha, \beta]_{T^*M} = \mathcal{L}_{\pi^\#(\alpha)}(\beta) - \mathcal{L}_{\pi^\#(\beta)}(\alpha) - \frac{1}{2}d(\langle \alpha \wedge \beta, \pi \rangle) \end{aligned}$$

Remark 3.40. Usually, the last term on the last hand side is often written as $d(\pi(\alpha, \beta))$, where the bivector $\pi \in \mathfrak{X}^2(M)$ is here seen as a bilinear form on $\Omega^1(M)$. We chose to use the pairing given by Equation (3.23) for it seems more transparent; see Remark 3.35 for a comparison between the two notations.

Proof. We already know that $\pi^\#$ is vector bundle morphism and the bracket $[\cdot, \cdot]_{T^*M}$ is obviously skew-symmetric. Then we only need to show that the bracket satisfies the Jacobi identity and that it is compatible with the anchor map in the sense that they satisfy the Leibniz rule. Since every differential one-form is locally exact, and that the bracket is defined only locally, we will evaluate both the Jacobi identity and the Leibniz rule on exact differential one-forms. Then we can observe the following fact:

Lemma 3.41. *On exact differential one forms, the bracket $[\cdot, \cdot]_{T^*M}$ satisfies the following identity:*

$$[df, dg]_{T^*M} = d\{f, g\}$$

for every $f, g \in \mathcal{C}^\infty(M)$.

Exercise 3.42. Prove this lemma by using Proposition (3.33), the definition of $\pi^\#$ and the properties of the Lie derivatives given in Proposition 3.33.

Let us now show that $[\cdot, \cdot]_{T^*M}$ satisfies the Jacobi identity on exact differential one-forms; let $f, g, h \in \mathcal{C}^\infty(M)$, then by Lemma 3.41 we obtain:

$$\begin{aligned} [df, [dg, dh]_{T^*M}]_{T^*M} &= [df, d\{g, h\}]_{T^*M} \\ &= d\{f, \{g, h\}\} \\ &= d(\{\{f, g\}, h\} + \{g, \{f, h\}\}) \\ &= [d\{f, g\}, dh]_{T^*M} + [dg, d\{f, h\}]_{T^*M} \\ &= [[df, dg]_{T^*M}, dh]_{T^*M} + [dg, [df, dh]_{T^*M}]_{T^*M} \end{aligned}$$

Notice that the Jacobi identity for $[\cdot, \cdot]_{T^*M}$ is a consequence of the Jacobi identity for the Poisson bracket. Now let us check the Leibniz rule (4.86):

$$\begin{aligned} [df, gdh]_{T^*M} &= \mathcal{L}_{\pi^\#(df)}(gdh) - \mathcal{L}_{\pi^\#(gdh)}(df) - \frac{1}{2}d(\langle df \wedge gdh, \pi \rangle) \\ &= \mathcal{L}_{X_f}(gdh) - \mathcal{L}_{gX_h}(df) - \frac{1}{2}d(g\langle df \wedge dh, \pi \rangle) \\ &= \mathcal{L}_{X_f}(g)dh + g\mathcal{L}_{X_f}(dh) - d(gX_h(f)) - d(g\{f, h\}) \\ &= X_f(g)dh + g\{f, h\} - d(g\{h, f\}) - d(g\{f, h\}) \\ &= \pi^\#(df)(g)dh + g[df, dh]_{T^*M} \end{aligned}$$

We used the definition of the Lie derivative as given by Definition 3.33, as well as the definition of Hamiltonian vector fields (see Equation (3.6)). \square

The fact that T^*M is a Lie algebroid over M implies that the image of the anchor map $\pi^\#$ is a (possibly singular) smooth involutive distribution on M . If the distribution is regular – i.e. has constant rank – Frobenius theorem 2.68 implies that it is integrable to a regular foliation. If the distribution is singular – i.e. if its rank is not constant – then, because it is finitely generated and involutive, it turns out that it also integrates, to what is called a *singular foliation*. The latter notion generalizes the notion of regular foliation in the following way:

Definition 3.43. A singular foliation is a partition $\bigsqcup_{\alpha} L_{\alpha}$ of M by disjoint connected weakly embedded submanifolds L_{α} called leaves, such that the induced distribution $D : x \mapsto T_x L_{\alpha(x)}$ is smooth. Here $\alpha(x)$ denotes the index α such that L_{α} is the unique leaf passing through x .




Figure 15: Two examples of partition of \mathbb{R}^2 : the first one consists of horizontal lines for $x \neq 0$ and the vertical line for $x = 0$, while the second one has points on the vertical axis. The figure on the left is *not* a singular foliation because any tangent vector to the submanifold at $x = 0$ does not satisfy the smoothness condition: any smooth extension of ∂_y in any neighborhood of the origin necessarily contains a vertical part, which is then not tangent to the horizontal leaves outside the vertical axis. On the contrary, the figure on the right is a singular foliation.

An alternative formulation, closer to that relying on foliated atlases for regular foliations, involves the notion of *distinguished atlas*. We say that M admits a distinguished atlas (with respect to a partition $\mathcal{L} = \bigsqcup_{\alpha} L_{\alpha}$ of M into immersed submanifolds, if for every $x \in M$ there exists a chart (U, φ) such that [Stefan, 1974]:

1. $\varphi(U)$ decomposes as a product of connected open sets $\varphi(U) = V \times W \subset \mathbb{R}^p \times \mathbb{R}^{n-p}$;
2. $\varphi(x) = (0, 0)$;
3. for any $L \in \mathcal{L}$, $\varphi(L \cap U) = V \times l_L$, with $l_L = \{y \in W \mid \varphi^{-1}(0, y) \in L\}$.

In particular, the last condition implies that the leaves intersecting U have higher than or equal dimension to that passing through x . It is equivalent to requiring that the map $x \mapsto \dim(L_x)$ (where L_x is the leaf passing through x), going from the topological space M – equipped with the distinguished atlas – to the integer, is continuous. Since the target space has the discrete topology, the map is lower semi-continuous, hence the result. Moreover, it also implies that the map is locally constant on a dense subset of M , which means that the singular leaves are quite rare actually.

Example 3.44. The distribution defined in Example 2.64, although integrable, is *not* a singular foliation because the leaf passing through the origin $(0, 0)$ has higher dimension than its neighbors. Thus it does not admit a distinguished atlas as item 3. fails to be satisfied.

Frobenius’ result about integrability can then be generalized to singular distributions thanks to Hermann’s theorem:

Theorem 3.45. Hermann Theorem. A locally finitely generated singular smooth distribution D on a smooth manifold is integrable (to a singular foliation) if and only if it is involutive.

Remark 3.46. The first assumption, that the distribution is locally finitely generated means the following: for every point $x \in M$, there exist an open neighborhood U of x and a finite number of smooth sections $X_1, \dots, X_m \in \Gamma(U, D)$ such that, for every open set V such that $\bar{V} \subset U$, the space of smooth sections $\Gamma(V, D)$ is generated as a $C^\infty(V)$ -module by the restrictions of X_1, \dots, X_m to V . The definition seems complicated but it is made so that the corresponding notion is local.

The idea behind integration of a singular smooth distribution is that the smooth manifold M is foliated by a set of weakly embedded submanifolds called *leaves* such that, given any point x , the tangent space to the leaf through x – denoted L_x – coincides with D_x :

$$T_x L_x = D_x$$

This identity being actually true for every point y of the leaf: $T_y L_x = D_y$. Since the rank of the distribution jumps, the dimension of the leaves will jump as well. A reservoir of examples of integrable distributions come from the following observation:

Proposition 3.47. *The (possibly singular) distribution generated by the anchor of a Lie algebroid is integrable.*

Proof. Let A be a Lie algebroid with anchor map ρ , and set $D_x = \text{Im}(\rho(A_x))$. This is a smooth distribution because each element X_x of D_x admits a preimage $a_x \in A_x$, and it is then sufficient to take any smooth section of A passing through a_x , and to project it to $\mathfrak{X}(M)$ via ρ . The image is a vector field X such that $X_y \in D_y$ for every y in some neighborhood of x . The distribution is locally finitely generated because A is a vector bundle of finite rank, so it admits local frames that induce local generators of $\Gamma(D)$. Finally, it is involutive because $\rho : \Gamma(A) \rightarrow \mathfrak{X}(M)$ is a homomorphism of Lie algebras. Then, by Hermann Theorem 3.45, the distribution D is integrable to a (possibly singular) foliation. \square

Since, for a Poisson manifold (M, π) , Proposition 3.39 implies that T^*M is a Lie algebroid, the vector bundle morphism $\pi^\#$ defines an integrable generalized distribution $D_\pi = \text{Im}(\pi^\#) \subset TM$. The (possibly singular) foliation integrating this distribution is called the *characteristic foliation*. Since the rank of the distribution D at the point x equates that of the image of $\pi^\#$ at x and is thus even, the leaves of the foliation induced by a Poisson bivector field will always be even dimensional. We will now explain that they are, in fact, *symplectic manifolds*:

Definition 3.48. *A symplectic manifold is a smooth, even dimensional manifold, equipped with a non-degenerate closed two-form ω .*

Remark 3.49. Here, non-degeneracy means that the canonical vector bundle morphism $\omega^\flat = \iota_X(\omega) : TM \rightarrow T^*M$ induced by ω by contraction with tangent vectors is an isomorphism of vector bundles.

Example 3.50. For every smooth manifold M , the cotangent bundle T^*M is naturally a symplectic manifold: let denote q^i the local coordinate functions on M and p_i the local coordinate functions on the fibers of T^*M , i.e. $p_i(dx^j) = \delta_i^j$. Then the differential 2-form $\omega \in \Omega^2(T^*M)$ defined as $\omega = \sum_{i=1}^n dp_i \wedge dq^i$ is a non-degenerate closed 2-form on T^*M . This result shows that the isomorphism $\omega^\flat : TM \rightarrow T^*M$ then associates the tangent vector $\frac{\partial}{\partial q^i}$ to dp_i . In Hamiltonian mechanics, the coordinate function p_i is the *conjugate momentum* associated to q^i . Hence, symplectic manifolds represent the canonical setup to do classical mechanics (when it is well-defined). When working in a physical context we may call M the *configuration space* and T^*M the *phase space*. In particular the phase space \mathbb{R}^{2n} presented in Example 3.3 actually corresponds to $T^*\mathbb{R}^n$.

Now let us draw the relationship between symplectic manifolds and Poisson manifolds. At this point, we need not assume that ω is a closed differential form, although we still assume that it is non-degenerate. Then, the vector bundle morphism $\omega^\flat : TM \longrightarrow T^*M$ can be inverted. Its inverse is thus a vector bundle morphism $B^\sharp : T^*M \longrightarrow TM$ satisfying:

$$\omega^\flat \circ B^\sharp(\alpha) = \alpha$$

for every differential one-form $\alpha \in \Omega^1(M)$. We denote X_f the vector field $B^\sharp(df)$ (we will soon see that it is not contradictory with the earlier notation). Then, by construction we have:

$$\omega(X_f, X_g) = \omega^\flat(X_f)(X_g) = \omega^\flat \circ B^\sharp(df)(X_g) = df(X_g) = df(B^\sharp(dg))$$

for any two smooth functions $f, g \in \mathcal{C}^\infty(M)$. Antisymmetry with respect to f and g is implicit everywhere. Then, by Equations (3.24) and (3.25), there exists a unique non-degenerate bivector field B (hence the notation) such that:

$$\omega(X_f, X_g) = -\frac{1}{2}\langle df \wedge dg, B \rangle = -B(f, g) \quad (3.26)$$

This bivector field is actually not any bivector field:

Proposition 3.51. *ω is a symplectic form if and only if B is a Poisson bivector, that is to say:*

$$d\omega = 0 \quad \Longleftrightarrow \quad [B, B]_{SN} = 0$$

Proof. Let $f, g, h \in \mathcal{C}^\infty(M)$, and we set $X_f = B^\sharp(df)$, $X_g = B^\sharp(dg)$ and $X_h = B^\sharp(dh)$. Then, by Equation (1.31) the de Rham derivative of ω satisfies:

$$d\omega(X_f, X_g, X_h) = X_f(\omega(X_g, X_h)) - \omega([X_f, X_g], X_h) + \circlearrowleft \quad (3.27)$$

where \circlearrowleft symbolizes circular permutation of the three functions. Using successively Equations (3.26), (3.25) and (3.24), one deduces that the first term on the right hand side of Equation (3.27) is $-B(f, B(g, h))$. On the other hand, the second term on the right-hand side of Equation (3.27) can be rewritten as:

$$\begin{aligned} -\omega([X_f, X_g], X_h) &= \omega(X_h, [X_f, X_g]) \\ &= \omega^\flat \circ B^\sharp(dh)([X_f, X_g]) \\ &= [B(f, -), B(g, -)](h) \\ &= B(f, B(g, h)) - B(g, B(f, h)) \end{aligned}$$

Thus, noticing that $-B(g, B(f, h)) = -B(B(h, f), g)$ and writing explicitly the circular permutation, Equation (3.27) can be rewritten:

$$d\omega(X_f, X_g, X_h) = -B(B(f, g), h) - B(B(g, h), f) - B(B(h, f), g)$$

On the right-hand side, one can recognize minus the Schouten-Nijenhuis bracket of B with itself, so that:

$$d\omega(X_f, X_g, X_h) = -[B, B]_{SN}(f, g, h)$$

This prove the claim. □

Remark 3.52. A Hamiltonian vector field on a symplectic manifold is a vector field X such that there exists a smooth function f such that:

$$\omega(X, \cdot) = df \quad (3.28)$$

Then we have that $X = B^\sharp(df)$ but since B is a Poisson bivector by Proposition 3.51, we deduce that $X = \{f, \cdot\} = X_f$, the usual Hamiltonian vector field associated to f . Hence, on a symplectic manifold, the definition of Hamiltonian vector fields using the symplectic form and that using the Poisson bracket or bivector are equivalent.

Thus a symplectic manifold is a Poisson manifold where the Poisson bivector is non-degenerate. Conversely, using Proposition (3.51), one can show the converse statement: any non-degenerate bivector field B on a smooth manifold M gives rise to a non-degenerate differential 2-form ω which is closed – i.e. symplectic – if and only if B is a Poisson bivector. We can summarize these results in the following general statement:

Proposition 3.53. *Let M be an even dimensional smooth manifold. Then there is a one-to-one correspondence between non-degenerate Poisson structures and symplectic structures on M .*

Let M be an even dimensional smooth manifold, equipped with a symplectic form ω , to which correspond a non-degenerate Poisson bivector π . Let us now determinate the relationship between $\omega = \frac{1}{2}\omega_{kl}dx^k \wedge dx^l$ and $\pi = \frac{1}{2}\pi^{ij}\frac{\partial}{\partial x^i} \wedge \frac{\partial}{\partial x^j}$ in local coordinates x^1, \dots, x^n . Evaluating ω on the hamiltonian vector fields $X_{x^i} = \pi^{ij}\frac{\partial}{\partial x^j}$, Equation (3.26) is equivalent to:

$$\omega_{kl}\pi^{ik}\pi^{jl} = -\pi^{ij} \quad (3.29)$$

where summation on repeated indices is implicit. We denote the coefficients of the inverse matrix of $\bar{\pi} = (\pi^{rs})_{rs}$ by π_{rs} , with indices at the bottom to allow contractions, so that $\pi^{rs}\pi_{st} = \delta_t^r$. Then, multiplying both sides of Equation (3.29) with π_{jm} and summing over j we obtain:

$$\pi^{ik}\omega_{km} = \delta_m^i \quad (3.30)$$

Thus, the components ω_{kl} turns out to precisely be π_{kl} , i.e. the coefficients of ω form a matrix that is the inverse matrix of $\bar{\pi}$. Thus, a symplectic form and its associated non-degenerate Poisson bivector somehow represent dual, equivalent pictures.

We have so far shown that when the characteristic foliation of a Poisson manifold consists of one leaf – i.e. when the Poisson bivector is non-degenerate – then the leaf is a symplectic manifold. We want to generalize this result to degenerate Poisson bivectors, by studying the local picture of symplectic manifolds. Recall that the standard Poisson bivector on \mathbb{R}^{2n} is of the form (3.17). It is a non-degenerate Poisson bivector, and the corresponding symplectic form is given in Example 3.50:

$$\pi = \sum_{i=1}^n \frac{\partial}{\partial q^i} \wedge \frac{\partial}{\partial p_i} \quad \iff \quad \omega = \sum_{i=1}^n dp_i \wedge dq^i$$

One can check that their respective components π^{ij} and ω_{kl} satisfy Equation (3.30). It turns out that this structure is quite central in symplectic geometry, because every symplectic manifold is locally symplectomorphic to \mathbb{R}^{2n} :

Theorem 3.54. Darboux theorem. *Let (M, ω) be a symplectic manifold and let $x \in M$. Then there exists local coordinates (q^i, p_i) centered at x , with respect to which the symplectic form ω is expressed as:*

$$\omega = \sum_{i=1}^n dp_i \wedge dq^i$$

In other words, Darboux theorem states that locally, every symplectic manifold locally looks the same. It implies that there are no local invariants in symplectic geometry, contrary to Riemannian geometry for example. The above result occurs when M is symplectic or, equivalently, when it is a non-degenerate Poisson manifold, so that the characteristic foliation consists of one, unique leaf: the total manifold. By Proposition 3.53 we can reformulate Darboux theorem in terms of non-degenerate Poisson structures:

Theorem 3.55. Darboux theorem (Poisson version). *Let (M, π) be a Poisson manifold manifold such that π is non-degenerate, and let $x \in M$. Then there exists local coordinates (q^i, p_i) centered at x , with respect to which the symplectic form ω is expressed as:*

$$\pi = \sum_{i=1}^n \frac{\partial}{\partial q^i} \wedge \frac{\partial}{\partial p_i}$$

Remark 3.56. This theorem sheds light on why the standard Poisson structure on \mathbb{R}^{2n} is ‘rigid’ in the sense that $H_\pi^2(\mathbb{R}^{2n}) = 0$ and more generally every non-degenerate Poisson bivectors (by Remark 3.38). This is because any small (formal) deformation of such Poisson bivector is still non-degenerate, so they locally still look like the standard structure on \mathbb{R}^{2n} . Thus, we cannot ‘deform’ them.

Now what happens when the Poisson bivector is degenerate, i.e. when its rank does not equate the dimension of the manifold at every point? In that case, the generalized distribution D_π associated to the Poisson bivector π is integrable and its leaves are even dimensional. The following important result generalizing Darboux theorem 3.54 to Poisson manifolds sheds light on what happens locally:

Theorem 3.57. Weinstein splitting theorem. *Let (M, π) be a Poisson manifold of dimension n and let $x \in M$ be an arbitrary point. Denote the rank of the Poisson bivector π at x by $2r$, and let $s = n - 2r$. Then, there exists local coordinates $q^1, \dots, q^r, p_1, \dots, p_r, z_1, \dots, z_s$ centered at x , such that the Poisson bivector reads:*

$$\pi = \sum_{i=1}^r \frac{\partial}{\partial q^i} \wedge \frac{\partial}{\partial p_i} + \sum_{1 \leq k, l \leq s} \phi_{kl} \frac{\partial}{\partial z_k} \wedge \frac{\partial}{\partial z_l} \quad (3.31)$$

where the functions $\phi_{kl} = -\phi_{lk}$ are smooth functions, which depend on $z = (z_1, \dots, z_s)$ only, and which vanish when $z = 0$.

Weinstein’s theorem is not a result about local coordinates, but a result about the possibility of choosing a special subset of local coordinates satisfying some nice property regarding the Poisson bivector. It is a result of foliation theory that leaves are weakly embedded submanifolds. Then by Proposition 2.55, it is always possible to choose, in a vicinity of the point x , coordinates adapted to the leaf L_x : the first $2r$ coordinates are local coordinates on L_x , while the last s coordinates represent transversal ones. In particular, the zero locus of the last s coordinates represent the leaf through x in that vicinity, see Figure 16. Weinstein’s theorem states that, additionally, a choice of such local coordinates can be made so that, in a vicinity of the point x , the rank of the Poisson bivector field has constant rank $2r$ on the leaf through x , this rank coinciding by definition with the dimension of L_x . This implies in turn that the restriction of the Poisson bivector to L_x is a non-degenerate Poisson bivector $\pi|_{L_x}$ and its form is the standard one, of Theorem 3.55. By Proposition 3.53, this makes L_x a symplectic manifold. This fact being true for every point x and thus every leaf of the characteristic foliation, we have finally obtained a full characterization of the latter:

Proposition 3.58. *The leaves of the characteristic foliation of a Poisson manifold are symplectic manifolds.*

Example 3.59. The linear Poisson structure defined on the dual of a Lie algebra \mathfrak{g} induces a foliation of \mathfrak{g}^* by symplectic leaves. These actually correspond to the coadjoint orbits of \mathfrak{g} on \mathfrak{g}^* . Polynomial functions on \mathfrak{g}^* (i.e. elements of the universal enveloping algebra of \mathfrak{g}) that are constant along these orbits are called Casimir operators. This convention explains why, in Poisson geometry, functions whose hamiltonian vector field is zero are called Casimirs.




Figure 16: Representation of the local coordinates centered at the point x appearing in Weinstein splitting theorem 3.57. The z -coordinates are transversal to the leaf L_x passing through x .

Example 3.60. As a particular case of the last example, the Poisson bivector field associated to the Poisson structure of Example 3.5 is the following:

$$\pi = z\partial_x \wedge \partial_y + x\partial_y \wedge \partial_z + y\partial_z \wedge \partial_x$$

where we transformed capital letters into small ones. The symplectic leaves are the concentric spheres or radius r (2-dimensional) and the origin (0-dimensional).

Example 3.61. The Poisson manifold (\mathbb{R}^3, π) defined in Example 3.26 induces a distribution D_π generated by the following three hamiltonian vector fields:

$$X_x = x\frac{\partial}{\partial z}, \quad X_y = y\frac{\partial}{\partial z} \quad \text{and} \quad X_z = -x\frac{\partial}{\partial x} - y\frac{\partial}{\partial y}$$

This distribution is integrable into a singular foliation: the singular leaves are points of coordinates $(0, 0, z)$ because the three vectors fields vanish, while the regular leaves are 2-dimensional vertical planes escaping radially from the vertical axis because then X_z is radial and either X_x , X_y or both are vertical (see Figure 17).

Let us work in the half-space with $x > 0$, and use polar coordinates $(x, y, z) \mapsto (r, \theta, z)$, where $r = \sqrt{x^2 + y^2} > 0$ and $\theta = \arctan(\frac{y}{x}) \in]-\frac{\pi}{2}, \frac{\pi}{2}[$. Then, the constant vectors $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ become respectively:

$$\frac{\partial}{\partial x} = \cos(\theta)\frac{\partial}{\partial r} - \frac{\sin(\theta)}{r}\frac{\partial}{\partial \theta} \quad \text{and} \quad \frac{\partial}{\partial y} = \sin(\theta)\frac{\partial}{\partial r} + \frac{\cos(\theta)}{r}\frac{\partial}{\partial \theta}$$

Then, since $x = r \cos(\theta)$ and $y = r \sin(\theta)$, the Poisson bivector of Example 3.26 becomes:

$$\pi = r\frac{\partial}{\partial r} \wedge \frac{\partial}{\partial z} \tag{3.32}$$

Now, if one performs the following new change of radial coordinate: $r \mapsto \rho = \ln(r)$, one then obtains $\frac{\partial}{\partial r} \mapsto \frac{1}{r}\frac{\partial}{\partial \rho}$, so that the Poisson bivector (3.32) becomes:

$$\pi = \frac{\partial}{\partial \rho} \wedge \frac{\partial}{\partial z} \tag{3.33}$$




Figure 17: The singular leaves of the characteristic foliation of the Poisson bivector $\pi = (x \partial_x + y \partial_y) \wedge \partial_z$ are points on the z -axis (in orange, 0-dimensional submanifolds), while the regular leaves are vertical, radial planes (in purple, 2-dimensional submanifolds).

We have then found the expression of the Poisson bivector π in a set of coordinates (ρ, θ, z) adapted to the situation, although they are *not* those of Weinstein splitting theorem for they are not centered at any point. Expression (3.32) is valid even for $x \leq 0$ (unless $x = y = 0$), because there is no dependence on θ .

Now, to make the connection explicit with Equation (3.31), let (x_0, y_0, z_0) be a point such that $x_0 > 0$, let $\rho_0 = \ln(\sqrt{x_0^2 + y_0^2})$ and $\theta_0 = \arctan(\frac{y_0}{x_0})$. Denoting $q = \rho - \rho_0$, $p = z - z_0$ and $\varphi = \theta - \theta_0$, we have a set of (local) coordinates centered at the point (x_0, y_0, z_0) , such that (q, p) span the leaf through (x_0, y_0, z_0) – a vertical radial plane – and φ encodes the transversal direction and vanishes on the leaf. Moreover, since we have $\frac{\partial}{\partial \rho} = \frac{\partial}{\partial q}$, $\frac{\partial}{\partial z} = \frac{\partial}{\partial p}$ and $\frac{\partial}{\partial \theta} = \frac{\partial}{\partial \varphi}$, one can write Equation (3.33) as:

$$\pi = \frac{\partial}{\partial q} \wedge \frac{\partial}{\partial p} + \sin(\varphi) \frac{\partial}{\partial \varphi} \wedge \frac{\partial}{\partial \varphi}$$

The last term automatically vanishes because of the antisymmetry of the wedge product. However, we have nonetheless provided a smooth function $\phi : \varphi \mapsto \sin(\varphi)$ which only depends on φ and vanishes on the level set $\varphi = \theta - \theta_0 = 0$, and managed to write the Poisson bivector as in Equation (3.31). Thus, the set of (local) coordinates (q, p, φ) are those whose existence is claimed by Weinstein splitting theorem. For other regular point with $x \leq 0$, one uses the same argument. For singular points (on the z -axis), the Poisson bivector is zero.

Notice that, since a symplectic manifold is always even dimensional, when the smooth manifold M is odd-dimensional, there will necessarily be zero-dimensional leaves (hence trivial symplectic manifolds). In the case where the rank of π is locally constant at x – i.e. on some open

neighborhood U of x – then 1. the foliation induced by the Poisson bivector on U is regular, and 2. there exist s Casimirs such that the symplectic leaves correspond to the level sets of the Casimirs (and then can be taken to be the coordinates z^k). That is why one often call the local coordinates $(q^1, \dots, q^r, p_1, \dots, p_r, z^1, \dots, z^s)$ *Casimir-Darboux coordinates*. Knowing that a singular foliation forms a partition of the ambient manifold M , a corollary of Proposition 3.58 is the following:

Corollary 3.62. *Every point of a Poisson manifold is contained in a unique symplectic leaf.*

We conclude this subsection by the following very beautiful and nice result: one can show that the Poisson bracket can be entirely reconstructed from the data of the symplectic forms on the leaves of the characteristic foliation. One defines a *smooth family of symplectic leaves* on a manifold M as the data of a singular foliation such that each leaf L is a symplectic manifold (L, ω_L) , and such that for every $f \in C^\infty(M)$, the family of tangent vectors $(X_{f,x})_{x \in M}$ defined at each point by $\omega_{L,x}(X_{f,x}, -) = df_x$ (where L is the leaf through x) is a smooth vector field on M . A Poisson manifold obviously induces a smooth family of symplectic leaves, and Vaisman has shown the converse statement [Vaisman, 1994]:

Theorem 3.63. *Let M be a smooth manifold equipped with a smooth family \mathcal{L} of symplectic leaves. Then there exists a unique Poisson structure on M such that the characteristic foliation coincide with the foliation \mathcal{L} (as well as the symplectic structures on the leaves).*

Proof. See Theorem 2.14 in [Vaisman, 1994]. One implication has been proven by the discussion surrounding Weinstein’s splitting theorem, the other implication relies on defining $\{f, g\} = X_f(g)$ (since the hamiltonian vector fields are smooth). \square

3.3 Submanifolds and reduction in Poisson geometry

The study of submanifolds in Poisson geometry is slightly more intricate than in differential geometry, because one needs to evaluate if the Poisson bracket originally defined on the ambient manifold M descends to the submanifold $S \subset M$. In this section, unless otherwise stated, the word ‘submanifold’ designates any kind of submanifolds: immersed, weakly embedded or embedded.

Definition 3.64. *A Poisson submanifold of a Poisson manifold M is a submanifold $S \xrightarrow{\iota} M$ admitting a Poisson structure such that the inclusion map ι is a Poisson map.*

The immersed or (weakly) embedded submanifold S can always be seen as the image of a injective immersion/weak embedding/smooth embedding $F : N \longrightarrow M$, such that $S = F(N)$. Then one may equivalently consider that the submanifold S is a Poisson submanifold if N is a Poisson manifold and F is a Poisson map. Denoting $\{ \cdot, \cdot \}$ (resp. $\{ \cdot, \cdot \}_N$) the Poisson bracket on M (resp. on N), this definition implies that the Poisson submanifold $S = F(N)$ is characterized by the fact that:

$$\{F^*(f), F^*(g)\}_N = F^* (\{f, g\}) \quad (3.34)$$

for every $f, g \in C^\infty(M)$. We shall see that in terms of bivector fields, Equation (3.34) can be restated as the fact that the Poisson bivector π defined on M is tangent to S at every point of S : $\pi_x \in \wedge^2 T_x S \subset \wedge^2 T_x M$ for every $x \in S$. There are additional equivalent characterizations of Poisson submanifolds, both geometric and algebraic:

Proposition 3.65. *Let (M, π) be a Poisson manifold and let S be a submanifold of M . The following are equivalent:*

1. S is a Poisson submanifold;
2. $\pi|_S$ takes values in $\wedge^2 TS$;
3. $\pi^\sharp(TS^\circ) = 0$;
4. $\pi^\sharp(T^*M|_S) \subset TS$;
5. all Hamiltonian vector fields are tangent to S .

Remark 3.66. The notation TS° stands for the *annihilator* of TS . It is the vector bundle over S consisting of all the covectors vanishing on TS . More precisely, for every $x \in S$, one sets:

$$T_x S^\circ = \{\xi_x \in T_x^* M \mid \xi_x(T_x S) = 0\}$$

If M is n -dimensional and if S is a r -dimensional submanifold, TS° is a rank $n - r$ vector bundle over S .

Proof. 1. \iff 2. Suppose that the submanifold S is obtained as the image of an injective immersion $F : N \hookrightarrow M$ (weak and smooth embeddings are injective immersions). Then we define $\wedge^2 TS$ as the pushforward of the vector bundle $\wedge^2 TN$ on M via $F_* \wedge F_*$, and we have $\wedge^2 TS \subset \wedge^2 TM|_S$.

First, assume that S is a Poisson submanifold of M , i.e. that N admits a Poisson structure $\{\cdot, \cdot\}_N$, and that F is a Poisson map. In full generality, Equation (3.34) can be rewritten in terms of Poisson bivectors as:

$$\langle F^*(df \wedge dg), \pi_N \rangle_N = F^*(\langle df \wedge dg, \pi \rangle_M) \quad (3.35)$$

where π (resp. π_N) is the Poisson bivector corresponding to $\{\cdot, \cdot\}$ (resp. $\{\cdot, \cdot\}_N$). On the left hand-side the pairing is taken with respect to TN and T^*N , while on the right-hand side it is taken with respect to TM and T^*M . Equation (3.35) is to be understood as an equality on N or, equivalently, on $S = F(N)$. Restricting Equation (3.35) to S has the following two consequences: one can rewrite the left-hand side as $\langle df \wedge dg, F_* \wedge F_*(\pi_N) \rangle_M$, while dropping F^* on the right-hand side:

$$\langle df \wedge dg, F_* \wedge F_*(\pi_N) \rangle_M \Big|_S = \langle df \wedge dg, \pi \rangle_M \Big|_S \quad (3.36)$$

where both sides here have to be understood as the restriction to S of the underlying smooth functions. Since the functions f and g are arbitrary, one obtains that, on S , $\pi|_S = F_* \wedge F_*(\pi_N)$, which proves item 2 since by definition $\wedge^2 TS = F_* \wedge F_*(\wedge^2 TN)$.

Conversely, still assuming that $S = F(N)$ is a submanifold of M , then item 2. implies that there exists a bivector field π_N on N such that $\pi|_S = F_* \wedge F_*(\pi_N)$. Since F is an injective immersion, it is unique. Moreover, for every open set $U \subset M$ the bivector field π_N satisfies Condition (3.16) on $F^*(C^\infty(U))$. Let us show that this implies that Condition (3.16) is satisfied in the neighborhood of every point of N . Let $x \in S$ and let x^1, \dots, x^n be local coordinates adapted to the submanifold S in a neighborhood of x , in the sense of Proposition 2.57. That is to say, there exists a connected coordinate chart $V \subset N$ centered at $y = F^{-1}(x)$ in N and a coordinate chart (U, φ) centered at x such that:

$$\varphi(U \cap F(V)) = \varphi(U) \cap \{\mathbb{R}^k \times 0\}$$

In other words, if the dimension of N is k , we can assume that the first k coordinates x^1, \dots, x^k of the chart φ are those parametrizing both V and $U \cap F(V) \subset S$, so that the function F

becomes $(x^1, \dots, x^k) \mapsto (x^1, \dots, x^k, 0, \dots, 0)$. Then any function $g \in \mathcal{C}^\infty(V)$ can be written as the pull-back of a smooth function on $\mathcal{C}^\infty(U)$: let $\mu : (x^1, \dots, x^n) \mapsto (x^1, \dots, x^k)$ be the projection along the last $n - k$ coordinates, and let $f = g \circ \mu$. Then $g = f \circ F = F^*(f)$. Since $\mathcal{C}^\infty(V) = F^*(\mathcal{C}^\infty(U))$, then π_N satisfies Condition (3.16) on V . This result being true in the neighborhood of each point of N , we deduce that π_N is a Poisson bivector. Then, since Equation (3.36) holds for arbitrary functions f and g , implying in turn that Equation (3.35) holds, the map $F : N \hookrightarrow M$ is a Poisson map, turning S into a Poisson submanifold of M .

2. \iff 5. Again suppose that S is obtained (at least) as an injective immersion. Let $x \in S$ and let x^1, \dots, x^n be local coordinates adapted to the submanifold S in a neighborhood of x , as in the proof of the last item. In particular, letting V be a sufficiently small neighborhood of $y = F^{-1}(x)$ as in Proposition 2.57, the first k coordinates x^1, \dots, x^k parametrize $V \simeq F(V)$, while the last $n - k$ coordinates are transverse to $F(V)$. Then the Poisson bivector π can be decomposed as:

$$\pi = \frac{1}{2} \sum_{i,j=1}^k \pi^{ij} \frac{\partial}{\partial x^i} \wedge \frac{\partial}{\partial x^j} + \sum_{i=1}^k \sum_{j=k+1}^n \pi^{ij} \frac{\partial}{\partial x^i} \wedge \frac{\partial}{\partial x^j} + \frac{1}{2} \sum_{i,j=k+1}^n \pi^{ij} \frac{\partial}{\partial x^i} \wedge \frac{\partial}{\partial x^j} \quad (3.37)$$

For any smooth function $f \in \mathcal{C}^\infty(M)$, the associated hamiltonian vector field is then:

$$X_f = \sum_{i,j=1}^k \pi^{ij} \frac{\partial f}{\partial x^i} \frac{\partial}{\partial x^j} + \sum_{i=1}^k \sum_{j=k+1}^n \pi^{ij} \left(\frac{\partial f}{\partial x^i} \frac{\partial}{\partial x^j} - \frac{\partial f}{\partial x^j} \frac{\partial}{\partial x^i} \right) + \sum_{i,j=k+1}^n \pi^{ij} \frac{\partial f}{\partial x^i} \frac{\partial}{\partial x^j} \quad (3.38)$$

Only the first term on the right-hand side of Equation (3.37) is a section of $\wedge^2 TS$. Thus, if item 2. holds, then the second and third sums are zero on S , proving, using Equation (3.38), that every Hamiltonian vector field is necessarily tangent to S . Conversely, if item 5. holds, the last term of Equation (3.38) necessarily vanishes on S as it is not tangent to S , while in the parenthesis there is a term tangent to S and the other is not. Applying successively Equation (3.38) to the coordinate functions x^1, \dots, x^k give:

$$X_{x^l} = \sum_{j=1}^k \pi^{lj} \frac{\partial}{\partial x^j} + \sum_{j=k+1}^n \pi^{lj} \frac{\partial}{\partial x^j}$$

Since these hamiltonian vector fields have to be tangent to S for every $1 \leq l \leq k$, we deduce that on S , we have $\pi^{lj} = 0$ for very $1 \leq l \leq k$ and $k + 1 \leq j \leq n$. This proves that the Poisson bivector field π reduces to the first term of Equation (3.37) on the submanifold S , i.e. item 2.

3. \iff 4. Let $\xi \in \Omega^1(M)$ such that $\xi_x \in T_x S^\circ$ for every $x \in S$ and let η be a differential 1-form on M . Then by Equation (3.25) we have, for every $x \in S$:

$$\xi_x(\pi_x^\#(\eta_x)) = \frac{1}{2} \langle \eta_x \wedge \xi_x, \pi_x \rangle = -\eta_x(\pi_x^\#(\xi_x))$$

This identity being true for every differential one-form η on M and every ξ taking values in TS° on S , this implies that the right-hand side equals 0 if and only if item 3. holds and the left-hand side equals zero if and only if item 4. holds. Then item 3. is equivalent to item 4.

4. \iff 5. The direct implication is straightforward, while for the reverse implication, assume that every Hamiltonian vector field is tangent to S . Since through every point of the cotangent bundle of M passes an exact form, then for every point $x \in S$ and $\xi_x \in T_x^* M$, there exists a smooth function f defined on M such that $df_x = \xi_x$. Then $\pi_x^\#(\xi_x) = \pi_x^\#(df_x) = X_{f,x}$, which shows that $\pi^\#(\xi_x)$ actually takes values in $T_x S$. Since this is true for every point of the cotangent bundle over the submanifold S , one deduces that item 4. holds. \square

Example 3.67. Obvious examples of Poisson submanifolds are the symplectic leaves of the characteristic foliation induced by π^\sharp . They have the property that their Poisson bracket is non-degenerate. More generally, Poisson submanifolds are unions of (open subsets of) symplectic leaves (see Proposition 3.99).

Example 3.68. As seen in Example 3.61, using polar coordinates allow to write the Poisson bivector $\pi = (x\partial_x + y\partial_y) \wedge \partial_z$ on \mathbb{R}^3 as $\pi = r\partial_r \wedge \partial_z$. One then straightforwardly sees that the 2-dimensional symplectic leaves (the vertical radial planes) are Poisson submanifolds of (\mathbb{R}^3, π) .

Exercise 3.69. This exercise is a continuation of Example 3.60, where the symplectic leaves are the concentric spheres in \mathbb{R}^3 and aims at showing that they are indeed Poisson submanifolds. The hemi-sphere of radius $r > 0$ located in the positive y half-space admits adapted spherical coordinates (r, θ, φ) on it, where $r > 0$ is the distance from the origin, $\theta \in]0, \pi[$ is the angle between the positive z axis and the vector and $\varphi \in]0, \pi[$ is the angle between the x axis and the projection of the vector on the Oxy plane. Show that in these spherical coordinates the Poisson bivector of Example 3.60 reads:

$$\pi = \frac{1}{r\sin(\theta)} \partial_\theta \wedge \partial_\varphi$$

and deduce from it that the hemi-sphere of radius $r > 0$ equipped with this Poisson bivector is a Poisson submanifold of (\mathbb{R}^3, π) (we know that it should be, as it is (a submanifold of) the level set of the Casimir element $C = x^2 + y^2 + z^2 - r$ of π).

As seen earlier, one can always characterize geometric objects by algebraic ones and vice-versa. This is the goal of the following proposition:

Proposition 3.70. *Let S be a Poisson submanifold of the Poisson manifold M . Then, the multiplicative ideal:*

$$\mathcal{I}_S = \left\{ f \in \mathcal{C}^\infty(M) \text{ such that } f|_S \equiv 0 \right\}$$

is a Lie ideal of the Lie algebra $(\mathcal{C}^\infty(M), \{ \cdot, \cdot \})$.

Proof. Since every hamiltonian vector field is tangent to S on S , then for any smooth functions $f \in \mathcal{C}^\infty(M)$ and $g \in \mathcal{I}_S$, one has by definition of TS $X_f(g) = 0$ on S , which can be equivalently be written as $\{f, g\}(x) = 0$ for every $x \in S$, that is to say: $\{f, g\} \in \mathcal{I}_S$. This proves that \mathcal{I}_S is a Lie algebra ideal with respect to the Poisson bracket. \square

The condition stated in Proposition 3.70 is not sufficient to characterize Poisson manifolds, unless they are embedded. Indeed, for immersed or weakly embedded submanifolds, the tangent space at a point does not necessarily coincide with the set of tangent vectors on M at that point that vanish on I_S (see counter-Example 2.60):

$$T_x S \subset \left\{ X_x \in T_x M \mid X_x(f) = 0 \text{ whenever } f \in \mathcal{I}_S \right\} \quad (3.39)$$

The fact that \mathcal{I}_S is a Poisson ideal in $\mathcal{C}^\infty(M)$ means that for every $f \in \mathcal{I}_S$ and $g \in \mathcal{C}^\infty(M)$, the smooth function $X_g(f) = \{g, f\}$ vanishes on S , which implies that Hamiltonian vector fields belong to the set on the right hand-side of Equation (3.39). When S is an embedded submanifold, we can conclude that these hamiltonian vector fields are tangent to S , and hence that S is a Poisson submanifold.

Example 3.71. If C is a Casimir function on a Poisson manifold M , then the levels sets of regular values of C are closed embedded submanifolds of M (see Theorem 2.45). Let S be such the level set of such a regular value $\lambda \in \mathbb{R}$, then it is a closed embedded submanifold of M . The ideal of functions vanishing on S is then spanned by the function $x \mapsto C(x) - \lambda$ (see Theorem

1.1 in [Henneaux and Teitelboim, 1992] or pages 95-96 of [Sudarshan and Mukunda, 1974]), and we write $\mathcal{I}_S = \langle C - \lambda \rangle$. This forms a Lie ideal of $(\mathcal{C}^\infty(M), \{.,.\})$ (independently of λ being a regular value or not), as the following argument shows: any smooth function f satisfies $\{f, C - \lambda\} = \{f, C\} = -X_C(f)$, which vanishes on M by definition of C being a Casimir function, hence in particular it vanishes on S , so $\{f, C - \lambda\} \in \mathcal{I}_S$. This also shows that any hamiltonian vector field X_f vanishes on \mathcal{I}_S and, since S is an embedded submanifold because λ is a regular value, inclusion (3.39) becomes an equality. These facts imply that every hamiltonian vector fields are tangent to S , proving that it is a Poisson submanifold of M by item 5. of Proposition 3.65. For example the level sets of the Casimir element of Exercice 3.11 correspond to concentric spheres in \mathbb{R}^3 , and coincide with the symplectic leaves of $\mathfrak{so}_3(\mathbb{R})$, i.e. its coadjoint orbits.

Poisson submanifolds are actually very rare. As in symplectic geometry, there are weaker notions of submanifolds in Poisson geometry, that possess specific features leading to important applications in mathematical physics: *Poisson-Dirac* submanifolds and *coisotropic* submanifolds. The Poisson bracket of the ambient manifold descends on the former via the so-called *Poisson-Dirac reduction*, while one has to further take a quotient of the latter to define a Poisson bracket: this is the topic of *coisotropic reduction*.

The notion of Poisson-Dirac submanifold relies on the following notion: let S be a – immersed or (weakly) embedded – submanifold, $f \in \mathcal{C}^\infty(S)$ and $x \in S$, then a *local extension of f at x* is the data (V, U, \tilde{f}) of an open neighborhood V of x in S (which then can be embedded into M via Proposition 2.57), an open neighborhood $U \subset M$ of x in M such that $V \subset S \cap U$, and a smooth function $\tilde{f} \in \mathcal{C}^\infty(U)$, such that \tilde{f} and f coincide on V : $\tilde{f}|_V = f|_V$.

Lemma 3.72. *Let S be a – immersed or (weakly) embedded – submanifold of M :*

1. *every smooth function on S can be locally extended;*
2. *if S is closed and embedded, then every smooth functions on S admit global extensions.*

Proof. The proof of the second statement can be found in Lemma 2.27 in [Lee, 2003] and Proposition 1.36 in [Warner, 1983], and heavily rely on the closedness of the submanifold. This statement can be alternatively be described as the following isomorphism: $\mathcal{C}^\infty(S) \simeq \mathcal{C}^\infty(M) / \mathcal{I}_S$. Let us now prove the first statement: let $x \in S$ and let V be an open neighborhood of x ; Proposition 2.57 tells us that V forms a slice of U , i.e. a closed embedded submanifold of U . In that case, applying the already proven second statement, one can extend $f \in \mathcal{C}^\infty(V)$ to a smooth function \tilde{f} on U . \square

Next, we say that a local extension (V, U, \tilde{f}) of a function f at x is *horizontal* if the Hamiltonian vector field $X_{\tilde{f}}$ is tangent to S , i.e. if $X_{\tilde{f},y} \in T_y S$ for every $y \in V$. Although every function on S admits local extensions, it may not be true that it admits horizontal local extensions. Poisson-Dirac submanifolds are precisely those submanifolds in Poisson geometry which have such a property:

Definition 3.73. *A Poisson-Dirac submanifold of a Poisson manifold M is a submanifold $S \xhookrightarrow{\iota} M$ which is such that for every point $x \in S$, every smooth function f on S admits an horizontal local extension (V, U, \tilde{f}) at x .*

Remark 3.74. The definition comes from subsection 5.3.2 of [Laurent-Gengoux et al., 2013]. It implies in particular that the pull-back to S of the so-called *Dirac structure* on M corresponding to the Poisson bivector π is a Dirac structure on S . See Section 6 of these [lectures notes](#).

Obviously, every Poisson submanifold is a Poisson-Dirac submanifold since *every* hamiltonian vector field is tangent to a Poisson submanifold. Definition 3.73 however shows that this condition has been profoundly weakened for Poisson-Dirac submanifolds. The main interest of the latter – and the definition has been explicitly chosen to this purpose – is that the Poisson bracket on M descends to S in a unique way to turn S into a Poisson manifold in its own way:

Proposition 3.75. Poisson-Dirac reduction. *Let S be a Poisson-Dirac submanifold of the Poisson manifold $(M, \{.,.\})$. Then there exists a unique Poisson bracket $\{.,.\}_S$ on S such that for every $x \in S$ and every two smooth functions $f, g \in C^\infty(S)$, one has:*

$$\{f, g\}_S = \{\tilde{f}, \tilde{g}\}|_V \quad (3.40)$$

for any horizontal local extensions (V, U, \tilde{f}) and (V, U, \tilde{g}) of f and g at x .

Proof. The proof can be found in Proposition 5.24 of [Laurent-Gengoux et al., 2013]. □

Example 3.76. This example is taken from [Fernandes, 2005]: let M be a Poisson manifold and let G be a Lie group properly acting on M via Poisson diffeomorphisms. Then the fixed points set M^G is a Poisson-Dirac submanifold of M .

A Poisson-Dirac submanifold possesses at most one Poisson structure satisfying Equation (3.40), and this Poisson structure is completely determined by the Poisson structure of M . Notice that the Poisson bracket on S is defined from picking up two local extensions whose hamiltonian vector field is tangent to S . It does not work with every local extension, although any other choice of extensions (such that their hamiltonian vector field is tangent to S) gives the same result in Equation (3.40). The fact that not every extension would satisfy Equation (3.40) can be explained from the following observation: contrary to Poisson submanifolds, where the Poisson bivector, restricted to S , takes values in $\Lambda^2 TS$ (see item 2. of Proposition 3.65), on a Poisson-Dirac submanifold S the Poisson matrix $(\pi^{ij})_{ij}$, with respect to a choice of coordinates adapted to S (see e.g. Proposition 2.57 or Proposition 1.35 in [Warner, 1983]) can be decomposed into blocks and take the form:

$$(\pi^{ij})_{ij} = \begin{pmatrix} A & B \\ -B^t & D \end{pmatrix} \quad (3.41)$$

Then one can show that on S , the anti-diagonal components B and $-B^t$ identically vanish so that the Poisson bivector reduces to two independent terms: $\pi|_S = \pi_1 + \pi_2$, where π_1 corresponds to the A component in the matrix and takes values in $\Lambda^2 TS$, while π_2 corresponds to the D component. Thus, the Poisson bracket on S by Proposition 3.75 corresponds to π_1 , although on S the Poisson bivector $\pi|_S$ contains another component π_2 , which only vanishes when evaluated on local extensions whose hamiltonian vector fields are tangent to S . A nice presentation of this issue (in a slightly less general case, however allowing to split TM into a direct sum) is made in the discussion surrounding Lemma 2.15 in these [lecture notes](#). Although Poisson-Dirac submanifolds are very useful for the possibility that Poisson-Dirac reduction offer, Definition 3.73 is a bit obscure so that it is not very clear what does it look like in geometric terms. This is the role of the next proposition:

Proposition 3.77. *Let S be a submanifold of a Poisson manifold M . Then, the following are equivalent:*

1. S is a Poisson-Dirac submanifold;
2. for every $\alpha \in \Omega^1(S)$ there exists open sets $V \subset S$ and $U \subset M$ such that $V \subset S \cap U$, and a differential one-form $\tilde{\alpha} \in \Omega^1(U)$ such that $\alpha|_V = \iota^*(\tilde{\alpha}|_V)$ and $\pi^\#(\tilde{\alpha})$ is tangent to S ;

3. for each $x \in S$, there exist local coordinates on M centered at x such that, if the matrix of π with respect to these coordinates can be written in a block form, then there exists a neighborhood V of x in S such that the matrix is diagonal by block on V ;
4. $TS \cap \pi^\#(TS^\circ) = 0$ and the bivector field π_S induced from π on S via Equation (3.40) is smooth.

Proof. For item 2. see Lemma 3.29 of these [lectures notes](#), for item 3. see Proposition 5.25 in [[Laurent-Gengoux et al., 2013](#)], while for item 4. see subsection 9.2 in [[Crainic and Fernandes, 2004](#)]. □

This proposition is similar to Proposition 3.65, when S is a Poisson-Dirac submanifold. Since such a submanifold is precisely defined from the behavior of Hamiltonian vector fields of local extensions, the counterpart of item 5. of Proposition 3.65 is item 1. of Proposition 3.77. Item 3. of Proposition 3.65 corresponds to item 2. of Proposition 3.77, via a slight reformulation because not every Poisson-Dirac submanifold S admits a normal bundle N such that $T_x M = T_x S \oplus N_x$ for every $x \in S$ and $\pi|_S$ takes values in $\Lambda^2 TS \oplus \Lambda^2 N$. The fact that the Poisson bivector, restricted to S , does not coincide with the component of $\pi|_S$ in $\Lambda^2 TS$, implies in particular that the inclusion $\iota : S \hookrightarrow M$ is certainly not a Poisson map (except of course if S is a Poisson submanifold). Item 3. is useful to further segregate different kinds of submanifolds within the family of Poisson-Dirac submanifolds. Poisson submanifolds form one extremity of this family, for which $\pi^\#(TS^\circ) = 0$. The other extremity is represented by the following condition:

Definition 3.78. Let S be a Poisson-Dirac submanifold of a Poisson manifold M . We say that S is a cosymplectic submanifold (or a Poisson transversal⁸) if $TM|_S = TS \oplus \pi^\#(TS^\circ)$

Example 3.79. In Example 3.61, any one-dimensional submanifold of $\mathbb{R}^3 - \{z - \text{axis}\}$ that is transversal to the symplectic leaves is a cosymplectic submanifold of (M, π) , on which the Poisson structure is zero (see [[Crainic et al., 2021](#)]).

Exercise 3.80. This exercise is taken from [[Crainic et al., 2021](#)]. Let (M, π) be a Poisson manifold and let S be a cosymplectic submanifold. Let π_S be the induced Poisson structure obtained by Poisson-Dirac reduction (see Proposition 3.75). Show that if C is a Casimir element of π , then $C|_S$ is a Casimir element of π_S .

For Poisson submanifolds, $\pi^\#(TS^\circ)$ has rank zero, while for cosymplectic manifolds, it is of maximal rank $n - \dim(S)$. Subsection 2.2 of these [lecture notes](#) give plenty of informations on cosymplectic manifolds. Their main use is that their class contain what are called second class constraint surfaces in Hamiltonian mechanics. To any physical system in hamiltonian mechanics corresponds a configuration manifold Q (with local coordinates the generalized coordinates) and a canonically associated phase space T^*Q (with fiber coordinates the conjugate momenta). The phase space is a symplectic manifold, characterized by the canonical symplectic form $\omega = \sum_i dp_i \wedge dq^i$, dual to a non-degenerate Poisson bivector $\pi = \sum_i \frac{\partial}{\partial q^i} \wedge \frac{\partial}{\partial p_i}$. A state of the physical system corresponds to a point in the phase space. The equations of motions then govern the evolution of the state of the system and, accordingly, the trajectory of the point to which it is associated.

Sometimes, it may happen that the states that the physical system can occupy are *constrained* (by some physical constraint, such as e.g. the length of the thread of the pendulum). A *constraint* is thus a smooth function on T^*M such that physical states are points of its zero level set. A physical system may admit several (possibly functionally dependent) constraints ϕ_1, \dots, ϕ_m , so

⁸For a quick overview of the various notions of transversals in Poisson geometry, see the beginning of Section 6 in [[Bursztyn et al., 2019](#)].

that the physical state is contained to the *constraint surface* $\Sigma = \Phi^{-1}(0)$, where $\Phi : T^*Q \longrightarrow \mathbb{R}^r$ is uniquely defined as $\Phi(x) = (\phi_1(x), \dots, \phi_m(x))$. There are two main kinds of constraints: *first-class constraints* and *second-class constraints*. First class constraints are those constraints whose Poisson bracket with any function vanishing on Σ is zero; we denote them $\varphi_1, \dots, \varphi_s$. Second-class constraints are those which are not first-class, and are often denoted χ_1, \dots, χ_r (so that $r + s = m$). In particular it means that for any second class constraint χ_k , there exist at least another second class constraint χ_l such that $\{\chi_k, \chi_l\} \neq 0$ on Σ . We define the zero level set of the second-class constraints Σ_0 – it obviously includes Σ . Then Dirac has shown that at least in the neighborhood of the zero level set of the second-class constraints Σ_0 , one can define a Poisson bracket on T^*Q (or at least on some tubular neighborhood of Σ_0), because the matrix of functions $C = (\{\chi_k, \chi_l\})_{k,l}$ is invertible:

$$\{f, g\}_{Dirac} = \{f, g\} - \{f, \chi_k\}(C^{-1})^{kl}\{\chi_l, g\} \quad (3.42)$$

This bracket, called the *Dirac bracket*, is such that the second-class constraint become Casimirs of this new bracket and that Σ_0 is a symplectic leaf – hence a Poisson submanifold – of $(T^*Q, \{\cdot, \cdot\}_{Dirac})$.

Exercise 3.81. Show that any second class constraint χ_l is a Casimir element of the Dirac bracket.

Let us now explain in geometric terms what is happening. Let M be a symplectic manifold, whose corresponding non-degenerate Poisson bracket is denoted $\{\cdot, \cdot\}$. Let $\Phi : M \longrightarrow \mathbb{R}^r$ (where $r \leq \dim(M)$) be a smooth map and assume that 0 is a *regular value* of Φ . It means that $\Phi_{*,x} : T_x M \longrightarrow T_{\Phi(x)}\mathbb{R}^r$ is surjective for every $x \in \Phi^{-1}(0)$. Since Φ is a smooth map then the map $x \mapsto \text{rk}(\Phi_{*,x})$ is lower semi-continuous, so it means that there exists an open neighborhood U of the origin of \mathbb{R}^r such that Φ_* is surjective on the open set $\Phi^{-1}(U)$. By the regular level set Theorem 2.45, the level sets of every points of U are closed embedded submanifolds of M , which form a regular foliation of $\Phi^{-1}(U)$. We denote $\Sigma_0 = \Phi^{-1}(0)$ the level set of 0. Decomposing the map Φ on the basis of \mathbb{R}^r : $\Phi(x) = (\chi_1(x), \chi_2(x), \dots, \chi_r(x))$ gives r smooth functions $\chi_i \in C^\infty(M)$, $1 \leq i \leq r$, called *constraints*. They are said *regular* when the smooth map Φ has constant rank so that the level sets of Φ on U are closed embedded submanifolds (Theorem 2.45), and *irreducible* when they are functionally independent, i.e. if there are smooth functions f^i such that $\sum_i f^i \chi_i = 0$ then all the f^i are necessarily of the form $f^i = \sum_j \sigma^{ij} \chi_j$ with $\sigma^{ij} = -\sigma^{ji}$ (see Definition 4.42). The conjunction of the irreducibility and the regularity conditions can be coined as a unique condition: $d\chi_1 \wedge \dots \wedge d\chi_r \in \Gamma(\wedge^r T^*M)$ is nowhere vanishing on $\Phi^{-1}(U)$.

Let C be the anti-symmetric matrix of functions whose i, j -th component is:

$$C_{ij} = \{\chi_i, \chi_j\}$$

We further assume that:

$$\det(C) \neq 0 \text{ on } \Sigma_0 \quad (3.43)$$

This condition on the smooth functions $(\chi_i)_i$ characterizes *second class constraints*, and Σ_0 is called the second class constraints surface. Since condition (3.43) is an open condition, there exists a tubular neighborhood $V \subset \Phi^{-1}(U)$ of Σ_0 (because Σ_0 is embedded, see Theorem 10.19 in [Lee, 2003]) such that $\det(C) \neq 0$ on the whole of V . Let us show that this condition is central in the properties of Σ_0 :

Proposition 3.82. *The second class constraint surface Σ_0 is a cosymplectic submanifold of $(M, \{\cdot, \cdot\})$. In particular it is not a Poisson submanifold.*

Proof. Just for the sake of the exercise, let us first show that Σ_0 is *not* a Poisson submanifold. Since it is embedded in M , by the discussion below Proposition 3.70, the condition for Σ_0

to be a Poisson submanifold is that the multiplicative ideal \mathcal{I}_{Σ_0} of smooth functions on M vanishing on Σ_0 is a Lie subalgebra of $(\mathcal{C}^\infty(M), \{.,.\})$. Since Σ_0 is an embedded submanifold, the multiplicative ideal \mathcal{I}_{Σ_0} is generated – as a sub-algebra of $\mathcal{C}^\infty(M)$ – by the second class constraints χ_1, \dots, χ_r (see Theorem 1.1 in [Henneaux and Teitelboim, 1992] or pages 95-96 of [Sudarshan and Mukunda, 1974]). Then, if it ever occurred that \mathcal{I}_{Σ_0} was a Lie ideal of $\mathcal{C}^\infty(M)$, then it would mean that $\{\mathcal{I}_{\Sigma_0}, \mathcal{I}_{\Sigma_0}\} \subset \mathcal{I}_{\Sigma_0}$. In other words, every Poisson bracket $\{\chi_i, \chi_j\}$ would vanish on Σ_0 . But this would contradict condition (3.43). Hence Σ_0 is not a Poisson submanifold.

Now let us proof that on the contrary, Σ_0 is quite far from being a Poisson submanifold, as it is a cosymplectic submanifold. Let us first proof that it is Poisson-Dirac. Since Σ_0 is a closed embedded submanifold, by Lemma 3.72, we know that every smooth function $f \in \mathcal{C}(\Sigma_0)$ admits a global extension, i.e. a smooth function $F \in \mathcal{C}^\infty(M)$ on M such that $F|_S = f$. Actually, we have the following isomorphism: $\mathcal{C}^\infty(\Sigma_0) \simeq \mathcal{C}^\infty(M) / \mathcal{I}_{\Sigma_0}$. So, any other choice of function $F + \sum_i \lambda^i \chi_i$ (where the λ^i 's are smooth functions on M) is another global extension for f . Let us find such an extension which is horizontal, i.e. whose hamiltonian vector field is tangent to Σ_0 . Let us set $\theta_i = \{F, \chi_i\}$ and let us set $\tilde{f} = F - \theta_k C^{kl} \chi_l$ (summation implied), where for simplicity the C^{kl} denote the coefficients of the inverse matrix C^{-1} . Then one has:

$$X_{\tilde{f}}(\chi_i) = \{F - \theta_k C^{kl} \chi_l, \chi_i\} = \underbrace{\{F, \chi_i\}}_{=\theta_i} - \underbrace{(\{\theta_k, \chi_i\} C^{kl} + \theta_k \{C^{kl}, \chi_i\}) \chi_l}_{\text{vanishes on } \Sigma_0 \text{ because of } \chi_l} - \underbrace{\theta_k C^{kl} C^{li}}_{=\theta_k \delta_i^k = \theta_i}$$

which then vanishes on Σ_0 . Since the multiplicative ideal \mathcal{I}_{Σ_0} is generated by the constraints χ_i , it means that $X_{\tilde{f}}(\mathcal{I}_{\Sigma_0}) \subset \mathcal{I}_{\Sigma_0}$, i.e. all functions in $X_{\tilde{f}}(\mathcal{I}_{\Sigma_0})$ vanish on Σ_0 . Then, since Σ_0 is embedded, we have the equality (see Lemma 2.58):

$$T_x \Sigma_0 = \left\{ X_x \in T_x M \mid X_x(g) = 0 \text{ whenever } g \in \mathcal{I}_{\Sigma_0} \right\} \quad (3.44)$$

Since $X_{\tilde{f}}$ belongs to the right-hand side, it means that it is tangent to Σ_0 . Hence, the smooth function $\tilde{f} = F - \theta_k C^{kl} \chi_l$ is a horizontal (global) extension of f . This proves that Σ_0 is a Poisson Dirac submanifold of M .

Now, as the constraints χ_i generate \mathcal{I}_{Σ_0} , and that equality (3.44) holds, we deduce that the differential one-forms $d\chi_i$ form a frame of $T\Sigma_0^\circ$, for they are independent and $d\chi_i(X) = X(\chi_i) = 0$ if and only if the vector field X takes values in $T\Sigma_0$. Then, since the Poisson bivector field π is non-degenerate, it sends the rank r subbundle $T\Sigma_0^\circ$ to a rank r subbundle of $T\Sigma_0$. Since the rank of the vector bundle $T\Sigma_0$ is $n - r$ and that a Poisson-Dirac submanifold satisfies $T\Sigma_0 \cap \pi^\#(T\Sigma_0) = 0$, we conclude that $TM|_{\Sigma_0} = T\Sigma_0 \oplus \pi^\#(T\Sigma_0)$. In other words, Σ_0 is a cosymplectic submanifold of M . \square

Since Σ_0 is a Poisson-Dirac submanifold of $(M, \{.,.\})$, we denote by $\{.,.\}_{\Sigma_0}$ the Poisson bracket inherited by Σ_0 via Poisson-Dirac reduction. More generally using the same arguments as in the proof of Proposition 3.82, one can show that every level sets of the smooth map Φ are cosymplectic submanifolds of $(M, \{.,.\})$ (at least on $\Phi^{-1}(U)$), so they all inherit the a Poisson structure from that on M via Poisson-Dirac reduction. Now, the Dirac bracket defined in Equation (3.42) is another choice of Poisson structure⁹ on M (or at least on some tubular neighborhood V of Σ_0), relative to which the second class constraints χ_i are Casimirs. Then, by Example 3.71, we deduce that the level sets of Φ are the symplectic leaves of $(M, \{.,.\}_{Dirac})$ (but *not* of $(M, \{.,.\})$). Then, the second-class constraint surface Σ_0 is a cosymplectic submanifold of

⁹The meaning of the Dirac bracket in the context of Poisson and symplectic geometry was first explained by ?niatycki [Śniatycki, 1974].

$(M, \{.,.\})$ but is a Poisson submanifold of $(M, \{.,.\}_{Dirac})$ (or at least the tubular neighborhood V). What is even more interesting is the following result:

Proposition 3.83. *For simplicity assume that $\{.,.\}_{Dirac}$ is defined on the whole of M . Then the Poisson structure on Σ_0 making it a Poisson submanifold of $(M, \{.,.\}_{Dirac})$ is precisely the Poisson bracket $\{.,.\}_{\Sigma_0}$ inherited from $\{.,.\}$ via Poisson-Dirac reduction.*

Proof. We need to show that for any two smooth functions $f, g \in C^\infty(\Sigma_0)$, one has on Σ_0 :

$$\{\iota^*(f), \iota^*(g)\}_{\Sigma_0} = \{f, g\}_{Dirac}|_{\Sigma_0} \quad (3.45)$$

Since the second class constraints χ_i are Casimirs elements of the Dirac bracket, we have, for every smooth functions $f, g \in C^\infty(M)$:

$$\{f, g\}_{Dirac}|_{\Sigma_0} = \{f - \lambda^i \chi_i, g - \mu^j \chi_j\}_{Dirac}|_{\Sigma_0} \quad (3.46)$$

for any family of functions λ^i, μ^j (notice that the equality only holds on Σ_0). In particular, one can make special choices of λ^i and μ^j as in the proof of Proposition 3.82 so that the hamiltonian vector fields of $f - \lambda^i \chi_i$ and $g - \mu^j \chi_j$ are tangent to Σ_0 . Then, the very definition of the Dirac bracket implies that we have the following equality:

$$\{f - \lambda^i \chi_i, g - \mu^j \chi_j\}_{Dirac}|_{\Sigma_0} = \{f - \lambda^i \chi_i, g - \mu^j \chi_j\}|_{\Sigma_0} \quad (3.47)$$

Again, the identity holds only on Σ_0 . Now, the fact that f and $f - \lambda^i \chi_i$ coincide on Σ_0 can be written as $\iota^*(f) = \iota^*(f - \lambda^i \chi_i) \in C^\infty(\Sigma_0)$, where $\iota : \Sigma_0 \mapsto M$ is the inclusion map. Moreover, since we have that $C^\infty(\Sigma_0) \simeq C^\infty(M) / \mathcal{I}_{\Sigma_0}$, the smooth function f is a global extension of $\iota^*(f)$, and $f - \lambda^i \chi_i$ is an horizontal one. Then, together with Equations (3.40), (3.46) and (3.47), we deduce that:

$$\begin{aligned} \{\iota^*(f), \iota^*(g)\}_{\Sigma_0} &= \{\iota^*(f - \lambda^i \chi_i), \iota^*(g - \mu^j \chi_j)\}_{\Sigma_0} \\ &= \{f - \lambda^i \chi_i, g - \mu^j \chi_j\}|_{\Sigma_0} \\ &= \{f - \lambda^i \chi_i, g - \mu^j \chi_j\}_{Dirac}|_{\Sigma_0} \\ &= \{f, g\}_{Dirac}|_{\Sigma_0} \end{aligned}$$

which is Equation (3.45), as desired. \square

Another way of making sense of Proposition 3.83 is the following: for any two smooth functions $f, g \in C^\infty(\Sigma_0)$, one has:

$$\{f, g\}_{\Sigma_0} = \{\tilde{f}, \tilde{g}\}|_{\Sigma_0} = \{\hat{f}, \hat{g}\}_{Dirac}|_{\Sigma_0} \quad (3.48)$$

where on the one hand, $\tilde{f}, \tilde{g} \in C^\infty(M)$ are any horizontal local extensions of f, g (they are required in Poisson-Dirac reduction), and on the other hand $\hat{f}, \hat{g} \in C^\infty(M)$ are any local extensions of f, g (since in that case $\iota^*(\hat{f}) = f$ and $\iota^*(\hat{g}) = g$, making Equation (3.45) valid). This is the formalized content of Theorem 2.5 in [Henneaux and Teitelboim, 1992]. Let us give a final, alternative formulation: since the second class constraint surface Σ_0 is a cosymplectic submanifold of $(M, \{.,.\})$, the tangent bundle restricted to Σ_0 is a direct sum of the two subbundles $T\Sigma_0$ and $\pi^\#(T\Sigma_0^\circ)$, so that one can see the term $\{-, \chi_k\}(C^{-1})^{kl}\{\chi_l, -\}$ in the formula (3.42) defining the Dirac bracket as a bivector field taking values in $\pi^\#(T\Sigma_0^\circ)$, which precisely compensates the

block D in formula (3.41). Then, the matricial representation of the Dirac bracket with respect to adapted local coordinates around a point $x \in \Sigma_0$ becomes

$$\begin{pmatrix} A & B \\ -B^t & 0 \end{pmatrix} \quad (3.49)$$

Thus, to create the Dirac bracket one has removed the lower right component of the original Poisson bracket represented matricially in formula (3.41). Moreover we can see from the above matrix (Equation (3.49)) that on Σ_0 , the bivector associated to $\{.,.\}_{Dirac}$ reduces to A , which takes values in $\wedge^2 T\Sigma_0$, as is characteristic for a Poisson submanifold. See subsection 5.1 of [Calvo et al., 2010] for more details on this background story.

Example 3.84. Let $M = T^*\mathbb{R}^2$ be the cotangent bundle of \mathbb{R}^2 , and let denote the coordinate functions (x, y, p_x, p_y) . The canonical (non-degenerate) Poisson bracket on M is then $\pi = \frac{\partial}{\partial x} \wedge \frac{\partial}{\partial p_x} + \frac{\partial}{\partial y} \wedge \frac{\partial}{\partial p_y}$. Let us set $\chi_1 = p_y$ and $\chi_2 = p_x + x - 2y$; these two smooth functions on M make $\Phi = (\chi_1, \chi_2) : M \longrightarrow \mathbb{R}^2$ a submersion. Then, the level set of Φ at 0 is a 2-dimensional plane in M that we denote Σ_0 . The Poisson bracket of χ_1 and χ_2 is:

$$\{\chi_1, \chi_2\} = 2$$

So in particular, denoting $C_{ij} = \{\chi_i, \chi_j\}$, one obtains:

$$C = \begin{pmatrix} 0 & 2 \\ -2 & 0 \end{pmatrix}$$

which is constant on the whole of M . Then $\det(C) \neq 0$, and the inverse matrix is:

$$C^{-1} = \begin{pmatrix} 0 & -\frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix}$$

Being defined over the whole of M , the Dirac bracket will be defined on the whole of M :

$$\begin{aligned} \{f, g\}_{Dirac} &= \{f, g\} - \{f, \chi_1\} \times \left(-\frac{1}{2}\right) \times \{\chi_2, g\} - \{f, \chi_2\} \times \frac{1}{2} \times \{\chi_1, g\} \\ &= \{f, g\} + \frac{1}{2} \frac{\partial f}{\partial y} \left(-\frac{\partial g}{\partial x} + \frac{\partial g}{\partial p_x} - 2\frac{\partial g}{\partial p_y}\right) - \frac{1}{2} \left(-\frac{\partial f}{\partial x} + \frac{\partial f}{\partial p_x} - 2\frac{\partial f}{\partial p_y}\right) \frac{\partial g}{\partial y} \end{aligned}$$

One can check that χ_1 and χ_2 are Casimir elements of the Dirac bracket.

Example 3.85. Another example of a situation where the matrix is invertible on the whole of M is the following: let $M = T^*\mathbb{R}^3$ and let denote the coordinate functions (x, y, z, p_x, p_y, p_z) . Let us define the following four linear functions:

$$\chi_1 = x + y, \quad \chi_2 = p_x, \quad \chi_3 = p_y + p_z \quad \text{and} \quad \chi_4 = z - x$$

The level set of $\Phi = (\chi_1, \chi_2, \chi_3, \chi_4) : M \longrightarrow \mathbb{R}^4$ at 0 is a 2-dimensional plane, that we denote Σ_0 . This plane is not a Poisson submanifold of M (with respect to its canonical non-degenerate Poisson structure) because the Poisson brackets of the constraints χ_i do not all vanish on Σ_0 . Indeed, the matrix C whose i, j -th component is $\{\chi_i, \chi_j\}$ is:

$$C = \begin{pmatrix} 0 & 1 & 1 & 0 \\ -1 & 0 & 0 & 1 \\ -1 & 0 & 0 & -1 \\ 0 & -1 & 1 & 0 \end{pmatrix}$$

It has determinant 4 and is invertible on the whole of M , with inverse matrix:

$$C^{-1} = \frac{1}{2} \begin{pmatrix} 0 & -1 & -1 & 0 \\ 1 & 0 & 0 & -1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & -1 & 0 \end{pmatrix}$$

The corresponding Dirac bracket is so that the constraints χ_i are Casimirs elements, and the plane Σ_0 is a Poisson submanifold of $(M, \{.,.\}_{Dirac})$.

Example 3.86. Let $M = T\mathbb{R}^2$ and let $\chi_1 = xy - 1$ while $\chi_2 = p_x$. The smooth map $\Phi = (\chi_1, \chi_2) : M \rightarrow \mathbb{R}^2$ is so that Φ_* is surjective on $\Phi^{-1}(0)$. Then, the preimage $\Sigma_0 = \Phi^{-1}(0)$ is a closed embedded submanifold of M . It has two connected components because the preimage of the first constraint χ_1 has such. The Poisson bracket of the two constraints is:

$$\{\chi_1, \chi_2\} = y$$

Interestingly, this Poisson bracket vanishes on the hyperplane of equation $y = 0$, but this hyperplane does not intersect Σ_0 so the Poisson bracket never vanishes on Σ_0 . Hence, the matrix C is:

$$C = \begin{pmatrix} 0 & y \\ -y & 0 \end{pmatrix}$$

Then $\det(C) \neq 0$ on Σ_0 (not on the whole of M) because $y \neq 0$ on the surface, and the inverse matrix is (only defined in a neighborhood of Σ_0):

$$C^{-1} = \begin{pmatrix} 0 & -\frac{1}{y} \\ \frac{1}{y} & 0 \end{pmatrix}$$

The Dirac bracket will then be defined only in a tubular neighborhood of Σ_0 :

$$\begin{aligned} \{f, g\}_{Dirac} &= \{f, g\} - \{f, \chi_1\} \times \left(-\frac{1}{y}\right) \times \{\chi_2, g\} - \{f, \chi_2\} \times \frac{1}{y} \times \{\chi_1, g\} \\ &= \{f, g\} + \frac{1}{y} \left(\frac{\partial f}{\partial p_x} y + \frac{\partial f}{\partial p_y} x \right) \frac{\partial g}{\partial x} - \frac{1}{y} \frac{\partial f}{\partial x} \left(\frac{\partial g}{\partial p_x} y + \frac{\partial g}{\partial p_y} x \right) \end{aligned}$$

The constraints χ_1 and χ_2 are Casimirs of this bracket.

Coming back to our problem in Hamiltonian mechanics, first-class constraints define a submanifold Σ in the embedded cosymplectic submanifold Σ_0 . This submanifold is however not a Poisson-Dirac submanifold because first-class constraints satisfy a nullity condition on Σ : $\{\varphi_p, \varphi_q\} = 0$ (so Σ does not satisfy item 4. of 3.77). Rather, the submanifold Σ is what is called a *coisotropic submanifold*. Poisson brackets cannot descend to them, but under some circumstances, to a quotient of them, through a procedure called *Poisson reduction*, also called *coisotropic reduction*. The notion of coisotropy is well-known in symplectic geometry, and is attached to submanifolds S whose symplectic orthogonal $TS^{\perp\omega}$ is a sub-bundle of TS . Since, for a non-degenerate Poisson structure, one has $\pi^\#(TS^\circ) = TS^{\perp\omega}$, the condition that a submanifold is coisotropic is straightforwardly transported to the realm of Poisson geometry:

Definition 3.87. A coisotropic submanifold of a Poisson manifold (M, π) is a submanifold $S \xrightarrow{\iota} M$ such that $\pi^\#(TS^\circ) \subset TS$.

Example 3.88. Any codimension 1 submanifold S of a Poisson manifold is coisotropic because TS° is 1-dimensional, implying that the right-hand side of Equation (3.25) is zero, implying in turn that $\pi^\#(TS^\circ) \subset TS$.

Example 3.89. An interesting example of a coisotropic submanifold is provided by a theorem of A. Weinstein [Weinstein, 1988]: A smooth map $\varphi : (M_1, \pi_1) \longrightarrow (M_2, \pi_2)$ is a Poisson map if and only if its graph $\text{Gr}(\varphi) \subset M_2 \times M_1^-$ is a coisotropic submanifold (where M_1^- is the smooth manifold M_1 equipped with the opposite Poisson structure $-\pi_1$). This statement is the Poisson equivalent of the well-known result in symplectic geometry stating that if M_1, M_2 are symplectic manifolds, then a diffeomorphism $\varphi : (M_1, \omega_1) \longrightarrow (M_2, \omega_2)$ is a symplectomorphism if and only if its graph $\text{Gr}(\varphi) \subset M_2 \times M_1^-$ is a Lagrangian submanifold.

There are two distinguished sub-families of coisotropic submanifolds: those for which $\pi^\#(TS^\circ) = 0$, i.e. Poisson submanifolds, and on the other extreme those for which $\pi^\#(TS^\circ) = TS$; they are called *Lagrangian submanifolds* as they correspond to their counterparts in symplectic geometry. Obviously, given the condition established in Definition 3.87 and item 4. of Proposition 3.77, the intersection of the set of coisotropic submanifolds and the set of Poisson-Dirac submanifolds is precisely the set of Poisson submanifolds. As for the other kinds of submanifolds, coisotropic submanifolds have equivalent alternative definitions:

Proposition 3.90. *Let S be a submanifold of a Poisson manifold M . Then, the following are equivalent:*

1. S is a coisotropic submanifold;
2. for every smooth function $f \in \mathcal{C}^\infty(M)$ vanishing on some open set $V \subset S$, the Hamiltonian vector field X_f is tangent to S at every point of V ;
3. $\langle \wedge^2 TS^\circ, \pi \rangle = 0$, where $\langle \cdot, \cdot \rangle$ is the pairing between T^*M and TM .

Proof. The direction 1. \implies 2. is straightforward because $f|_V = 0$ means that $df \in TS^\circ|_V$, so let us turn to the direction 2. \implies 1. Let f be such a function vanishing on V and suppose $X_{f,x} \in T_x S = T_x V$ for every point $x \in V$. Let $\xi \in \Gamma(TS^\circ)$ then, one has on V :

$$0 = \xi(X_f) = -df(\pi^\#(\xi)) \quad (3.50)$$

We know for sure that $df \in \Gamma(TS^\circ|_V)$ but the fact that, upon shrinking it, V is an embedded submanifold of M (see Proposition 2.57), implies that $TV^\circ = (TS^\circ)|_V$ is spanned by the point-wise evaluation of exact differential one-forms df for those functions $f \in \mathcal{C}^\infty(M)$ vanishing on V . Since Equation (3.50) holds for every such function, and every $\xi \in \Gamma(TS^\circ)$, one deduces that $\pi^\#(\xi)$ is necessarily a tangent vector to S at every point of V . The proof of the equivalence 1. \iff 3. is straightforward, using Equation (3.25). \square

Remark 3.91. Notice that in the second item of Proposition 3.90, we did not ask f to vanish on S but on an open set of V precisely because we needed to characterize $TS^\circ|_V = TV^\circ$ as spanned by the exact differential one-forms df . And to do that we needed at least an embedded submanifold, which is true for V but not necessarily for S if immersed.

We have another characterization of coisotropic submanifolds, mimicking Proposition 3.70 for Poisson submanifolds. Indeed, it admits the following counterpart for coisotropic submanifolds:

Proposition 3.92. *Let S be a coisotropic submanifold of the Poisson manifold M . Then, the multiplicative ideal:*

$$\mathcal{I}_S = \left\{ f \in \mathcal{C}^\infty(M) \text{ such that } f|_S \equiv 0 \right\}$$

is a Poisson subalgebra of the Poisson algebra $(\mathcal{C}^\infty(M), \cdot, \{ \cdot, \cdot \})$.

The proof of Proposition 3.92 is a straightforward application of Definition 3.87. Notice however that, as for Poisson submanifolds, the converse implication – that the ideal \mathcal{I}_S being a Lie subalgebra of $\mathcal{C}^\infty(M)$ implies that S is a coisotropic submanifold of M – is true only when S is embedded in M .

Example 3.93. Taken from [Crainic et al., 2021]: let \mathfrak{g} be a finite dimensional real Lie algebra and \mathfrak{g}^* be the associated linear Poisson manifold, described in Example 3.4. Let $\xi \in \mathfrak{g}^*$, then the definition of the linear Poisson structure on \mathfrak{g}^* implies that, for any two elements $x, y \in \mathfrak{g}$:

$$\{\bar{x}, \bar{y}\}(\xi) = \overline{[x, y]}(\xi) = \xi([x, y]) \quad (3.51)$$

where \bar{x} is the notation used in Example 3.4 to symbolize the linear form on \mathfrak{g}^* defined as $\bar{x}(\xi) = \xi(x)$. Using Equations (3.24) and (3.25), the left-hand side of Equation (3.51) can be written as:

$$\{\bar{x}, \bar{y}\} = d\bar{y}(\pi^\#(d\bar{x})) \quad (3.52)$$

Let V be a subspace of the Lie algebra \mathfrak{g} , and let $V^\circ \subset \mathfrak{g}^*$ be its annihilator, that we will denote S in the following. Then, the annihilator of TS is spanned by the elements $d\bar{x}$ for every $x \in V$. This implies that S is a coisotropic submanifold of \mathfrak{g}^* if and only if, for every $x, y \in V$, the right-hand side of Equation (3.52) – evaluated at a point $\xi \in S = V^\circ$ – vanishes, i.e. if and only if the right-hand side of Equation (3.51) vanishes for every $\xi \in V^\circ$. This implies in turn that V° is a coisotropic submanifold of \mathfrak{g}^* if and only if V is a Lie subalgebra of \mathfrak{g} . Since \mathcal{I}_{V° is generated by $V^{\circ\circ} = V$ (because \mathfrak{g} is finite dimensional), we deduce that V° is a coisotropic submanifold of \mathfrak{g}^* if and only if \mathcal{I}_{V° is a Lie subalgebra of $\mathcal{C}^\infty(\mathfrak{g}^*)$. As a final remark, notice that V° is a Poisson submanifold if and only if V is a Lie ideal, if and only if \mathcal{I}_{V° is a Lie ideal of $\mathcal{C}^\infty(\mathfrak{g}^*)$.

Contrary to what happens for Poisson-Dirac submanifolds, coisotropic submanifolds are rarely equipped with an induced Poisson bracket. Rather, one may only have a *Poisson reduction* on a quotient of coisotropic submanifolds. Let us first define what is meant by this concept (we use the terminology of subsection 5.2.2 in [Laurent-Gengoux et al., 2013]):

Definition 3.94. *Let $(M, \{.,.\})$ be a Poisson manifold and S be a submanifold, N a smooth manifold and $p : S \rightarrow N$ a surjective submersion:*

$$\begin{array}{c} S \subset M \\ \downarrow p \\ N \end{array}$$

We say that the triple (M, S, N) is Poisson reducible if there exists a Poisson structure $\{.,.\}_N$ on N such that, for all open subsets $V \subset S$ and $U \subset M$ such that $V \subset U \cap S$, and for all functions $f, g \in \mathcal{C}^\infty(p(V))$, one has:

$$\{f, g\}_N(p(x)) = \{\tilde{f}, \tilde{g}\}(x) \quad (3.53)$$

for every $x \in V$, and arbitrary local extensions $\tilde{f}, \tilde{g} \in \mathcal{C}^\infty(U)$ of functions $f \circ p|_V$ and $g \circ p|_V$.

Remark 3.95. Poisson reduction is a particular case of what is called Marsden-Ratiu reduction on Poisson manifold, which also generalize Mayer-Marsden-Weinstein reduction on Hamiltonian G-spaces. See these [lectures notes](#), as well as this paper [Falceto and Zambon, 2008].

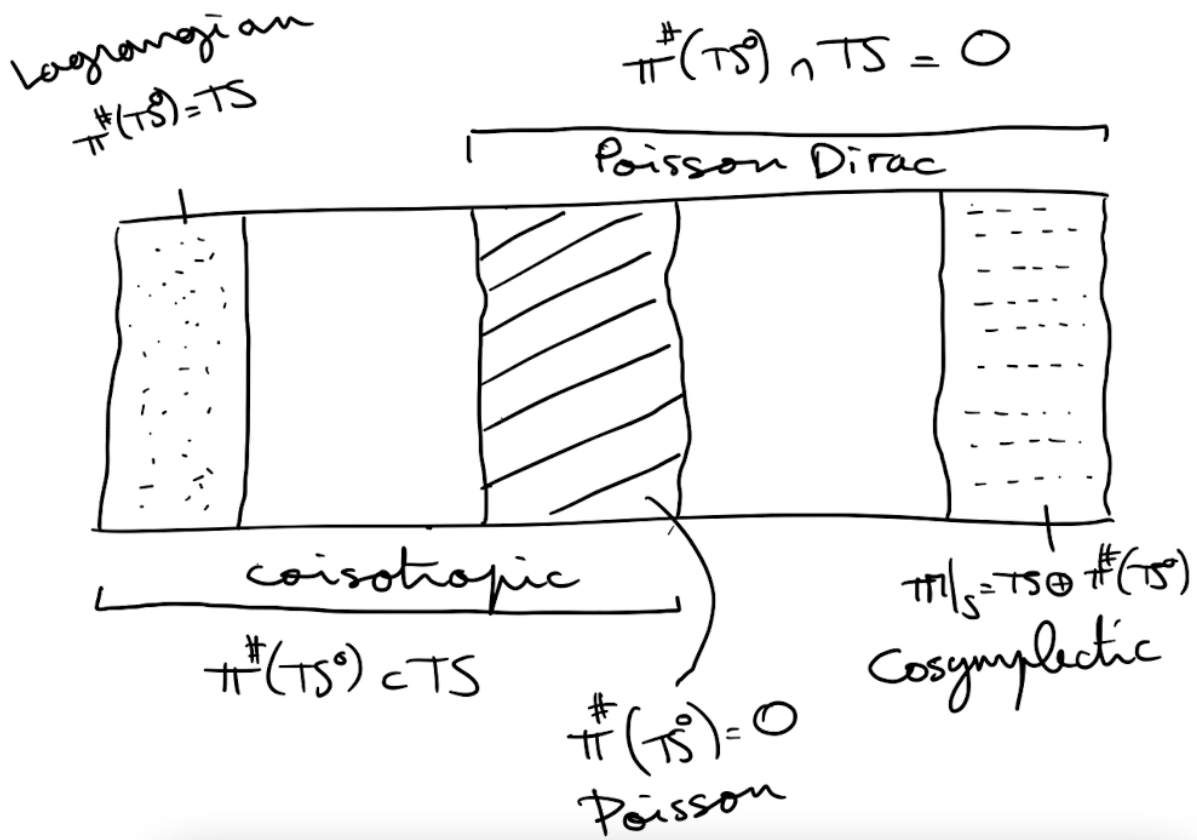


Figure 18: Schematic map of the various families of submanifolds in Poisson geometry. Poisson submanifolds are both coisotropic and Poisson-Dirac submanifolds. Cosymplectic and Lagrangian submanifolds are opposite to Poisson submanifolds in their respective families. See [Zambon, 2011] for additional informations about relationships between various kinds of submanifolds in Poisson geometry.

Example 3.96. If S is a submanifold of M , and f, g are two smooth functions on S , admitting local extensions \tilde{f} and \tilde{g} , then $\iota^* \tilde{f} = f$ and $\iota^* \tilde{g} = g$. Assuming that S is a Poisson submanifold, we set $N = S$, $p = \text{id}_S$, so that Equation (3.34) becomes (3.53). This property being true for every smooth functions f, g , the triple (M, S, S) is Poisson reducible.

It turns out – by Proposition 5.11 of [Laurent-Gengoux et al., 2013] – that a triple (M, S, N) satisfying the conditions of Definition 3.94 is Poisson reducible if and only if:

1. for every function $\tilde{f} \in \mathcal{C}^\infty(U)$ whose restriction to V is constant of the fibers of p , the hamiltonian vector field $X_{\tilde{f}}$ is tangent to S at every point of V ;
2. for every pair of functions $\tilde{f}, \tilde{g} \in \mathcal{C}^\infty(U)$ whose restriction to V is constant of the fibers of p , the restriction of their Poisson bracket to V is constant of the fibers of p .

Then, one can show that in such a case S is a coisotropic submanifold of M . It is thus legitimate to ask under which circumstances a coisotropic S allows a Poisson reduction to some quotient of itself. The following argument provides such an example of assumptions fitting the situation met often in constrained Hamiltonian systems.

An alternative approach Poisson reducibility involves the use of Poisson algebras and coisotropic ideals, and is such that geometric Poisson reduction is a particular case when the Poisson algebra is $\mathcal{C}^\infty(M)$ and the corresponding coisotropic ideal is the ideal \mathcal{I}_S of functions vanishing on a coisotropic submanifold S . By Proposition 3.92, it is a Lie subalgebra of $\mathcal{C}^\infty(M)$. I Poisson reduction of a Poisson structure from M to a coisotropic submanifold S then consists in first defining the space of function:

$$\left(\mathcal{C}^\infty(M)/\mathcal{I}_S\right)^{\mathcal{I}_S}$$

If S is embedded in M , then the above algebra of functions is isomorphic to the following:

$$\left(\mathcal{C}^\infty(S)\right)^{\mathcal{I}_S}$$

We recognize here the algebra of smooth functions on S which are invariant under the flow of Hamiltonian vector fields of functions of \mathcal{I}_S . In physical terms, they are the gauge invariant functions on the constraint surface, that is to say, the physical observables. Now the question is to determine if this algebra of functions is an algebra of function of a Poisson manifold, i.e. if there exists a Poisson manifold N such that its algebra of smooth functions $\mathcal{C}^\infty(N)$ is isomorphic to $\simeq (\mathcal{C}^\infty(S))^{\mathcal{I}_S}$. Such a Poisson manifold N may not exist, in particular if the leaf space of the foliation defined by the Hamiltonian vector fields of \mathcal{I}_S is not a smooth manifold. This is precisely the assumption appearing in the following statement:

Proposition 3.97. *Let S be a coisotropic submanifold of M , and assume that $\pi^\#(TS^\circ)$ has constant rank over S (i.e. defines a regular smooth distribution). Then it is integrable in the sense of Frobenius and if the space of leaves N of the corresponding regular foliation is a smooth manifold, the triple (M, S, N) is Poisson reducible.*

Proof. We will show that items 1. and 2. above are satisfied (see also Remark 5.15 in [Laurent-Gengoux et al., 2013]). In the present context, p is the quotient map sending S to the leaf space N , so the fibers of p are the leaves.

First, let $\tilde{f} \in \mathcal{C}^\infty(U)$ such that it is constant on the leaves, and let ξ a differential one-form taking values in TS° on S . Since S is a coisotropic submanifold, $\pi^\#(\xi)$ is a vector field taking values in the regular integrable distribution, that is to say it is tangent to the leaves. Since \tilde{f} is constant along the leaves, $d\tilde{f}(\pi^\#(\xi)) = \pi^\#(\xi)(\tilde{f}) = 0$ on S . By Equation (3.25), the left-hand side of the former equation is equal to $-\xi(\pi^\#(d\tilde{f}))$. Then it vanishes on S and since ξ takes values in TS° , and that the vanishing of $-\xi(\pi^\#(d\tilde{f}))$ is valid for any such ξ , we deduce that $X_{\tilde{f}} = \pi^\#(d\tilde{f})$ is tangent to S .

Secondly, let $\tilde{f}, \tilde{g} \in \mathcal{C}^\infty(U)$ be two smooth functions which are constant along the leaves and let ξ be a differential one form taking values in TS° . Then, by the first point just proven, $X_{\tilde{f}}$ and $X_{\tilde{g}}$ are tangent to S , so is their Lie bracket, and we have:

$$0 = \xi([X_{\tilde{f}}, X_{\tilde{g}}]) = \xi(X_{\{\tilde{f}, \tilde{g}\}}) = \xi(\pi^\#(d\{\tilde{f}, \tilde{g}\})) = -d\{\tilde{f}, \tilde{g}\}(\pi^\#(\xi)) = -\pi^\#(\xi)(\{\tilde{f}, \tilde{g}\})$$

Since by definition, $\pi^\#(\xi)$ is a vector field tangent to the leaves, and that $\pi^\#(TS^\circ)$ generate all such tangent vector fields, we deduce that the Poisson bracket $\{\tilde{f}, \tilde{g}\}$ is constant along the leaves, as required. \square

Remark 3.98. Proposition 3.97 gives a geometric grounding to the algebraic approach to Poisson reduction, also called *Sniatycki-Weinstein reduction*. There is an alternative algebraic version of Poisson reduction, which instead of first reducing the algebra of functions from $\mathcal{C}^\infty(M)$ to $\mathcal{C}^\infty(S)$, and then to take the invariant functions, goes the other way around: first defining the

normalizer of \mathcal{I}_S in $\mathcal{C}^\infty(M)$ (i.e. the functions invariant under the flow of the Hamiltonian vector fields of \mathcal{I}_S), and then only take the quotient by \mathcal{I}_S . This procedure is called *Dirac reduction* and is described in the discussion preceding Definition 5.37. See Section 1 of [Blacker et al., 2022] for an overview of the various approaches in Poisson reduction, and how they coincide when the constraint surface is smoothly embedded.

This proposition is quite useful to study Hamiltonian under constraints. We have seen earlier that second-class constraints define an embedded cosymplectic submanifold of a Poisson manifold M . On the other hand, first class constraints define a coisotropic submanifold of M (here we assume M to be a symplectic manifold). A quick way to see this is by using the converse of Proposition 3.70, which holds for embedded submanifolds, which is an assumption largely met in most cases. Thus, assume that we have s constraints $\varphi_1, \dots, \varphi_s$ which are *irreducible* – i.e. for any smooth functions f^i such that $\sum_i f^i \varphi_i = 0$ then all the f^i are necessarily of the form $f^i = \sum_j \sigma^{ij} \varphi_j$ with $\sigma^{ij} = -\sigma^{ji}$ – and *regular* – i.e. their zero level set $\Sigma = \bigcap_{i=1}^s \varphi_i^{-1}(0)$ defines an embedded submanifold Σ of M . We can reformulate both properties into one condition, using the Regular Level Set Theorem 2.45: the smooth map $\Phi = (\varphi_1, \dots, \varphi_s) : M \rightarrow \mathbb{R}^s$ is a submersion on its zero level set Σ , i.e. the differential s -form $d\varphi_1 \wedge \dots \wedge d\varphi_s$ is nowhere vanishing on it. This proves that Φ_* is surjective on this level set (actually on a tubular neighborhood), proving in turn that Σ is an embedded submanifold of M .

Being first-class means that $\{\varphi_i, \varphi_j\} = 0$ on Σ for every $1 \leq i, j \leq s$, which is actually equivalent to saying that $\{\varphi_i, f\} = 0$ for every $f \in \mathcal{I}_S$, because every such function is functionally dependent on the constraints since Σ is an embedded submanifold (see Theorem 1.1 in [Henneaux and Teitelboim, 1992] or pages 95-96 of [Sudarshan and Mukunda, 1974]). But this is just the condition that \mathcal{I}_S is a Lie subalgebra of $\mathcal{C}^\infty(M)$. Being embedded, this implies that Σ is a coisotropic submanifold of M . Since the differential one forms $d\varphi_i$ span $T\Sigma^\circ$, the hamiltonian vector fields $X_{\varphi_i} = \pi^\#(d\varphi_i)$ span $\pi^\#(T\Sigma^\circ)$ and define a regular distribution on Σ (the rank of π is constant over Σ). By Frobenius theorem this distribution is integrable and the leaf space Σ_{ph} is called the *reduced phase space* because its points are the physical states of the system: on the one hand they all satisfy the constraints, and on the other hand we have got rid of the gauge freedom (symbolized by the leaves of the foliation). By Proposition 3.97, if the reduced phase space is a smooth manifold, the Poisson bracket of M descends to Σ_{ph} .

However, in most situation, we have a mixed set of constraints, i.e. some of them are first-class and some of them are second-class. Then, the strategy to obtain the physical phase space is first, to perform a Poisson-Dirac reduction on the second-class constraint surface Σ_0 , which is a cosymplectic embedded submanifold of $(T^*Q, \{\dots\})$ (and a symplectic leaf of $(T^*Q, \{\dots\}_{Dirac})$), and second, to consider the first-class constraint surface Σ as a coisotropic submanifold of $(\Sigma_0, \{\dots\}_{\Sigma_0})$ (or equivalently of $(T^*Q, \{\dots\}_{Dirac})$ because the former second-class constraint become first class with respect to the Dirac bracket). It turns out that Σ is a presymplectic submanifold of $(T^*Q, \{\dots\}_{Dirac})$ because the restriction of the symplectic form is degenerate there. See this chapter as well as [Gotay et al., 1978] for a clear presentation of this latter approach. Then, the leaf space of the regular foliation generated by the vector fields $X_{\varphi_i} = \{\varphi_i, \cdot\}_{Dirac}$ on Σ is the physical phase space Σ_{ph} of the theory. One eventually obtains a symplectic structure on this reduced phase space by proceeding to a Poisson reduction from Σ to Σ_{ph} . Notice that physicists have found a way of performing Poisson reduction while also circumventing the complicated quotient procedure: this is called the *BRST-BFV formalism* or *homological Poisson reduction* and it is quite useful, as it uses simple cohomological techniques instead of using quotients of the coisotropic submanifold. See Section 5.2 for more details.

To conclude this section, let us discuss a bit more the relationship between the symplectic leaves of a Poisson manifold and its submanifolds. We know from Theorem 3.63 that there is a one-to-one correspondence between Poisson structures on M and smooth families of symplectic

leaves on M . Then a way of defining a Poisson structure on a given submanifold S of M would be to study the properties of the intersection of S with the symplectic leaves of M :

Proposition 3.99. *Let M be a Poisson manifold and let $S \subset M$ be a submanifold. Then we have the following statements:*

1. *S is a Poisson submanifold if and only if for each symplectic leaf L , the intersection $S \cap L$ is an open set of L ;*
2. *S is a Poisson-Dirac submanifold if and only if for each symplectic leaf L , the intersection $S \cap L$ is clean¹⁰ and a symplectic submanifold of L , such that these symplectic structure turn the connected components of the intersections $S \cap L$ into a smooth family of symplectic leaves on S , when L ranges over the symplectic leaves of M .*

In both cases, the symplectic leaves induced by the Poisson bivector π_S on S are the connected components of the intersections $S \cap L$, where L ranges over all symplectic leaves of M . Finally, for coisotropic submanifolds, one has the following statement:

3. *assuming that S has clean intersection with all the symplectic leaves of M , S is a coisotropic submanifold if and only if for each symplectic leaf L , the intersection $S \cap L$ is a coisotropic submanifold of L .*

Proof. For Poisson submanifolds, the proof can be found in Proposition 2.12 in [Laurent-Gengoux et al., 2013] or in Proposition 3.26 of these lectures notes. For Poisson-Dirac submanifolds, the proof can be found in Proposition 6 of [Crainic and Fernandes, 2004], and for coisotropic submanifolds it can be found in Proposition 3.29 of the same lectures notes. \square

Remark 3.100. The latter statement is more stringent because in general coisotropic submanifolds are far from havin clean intersections with symplectic leaves. See Remark 1. of [Zambon, 2011].

Remark 3.101. There exist plenty of other kinds of submanifolds in Poisson geometry, e.g. *isotropic submanifolds* are those submanifolds S such that $TS \subset \pi^\#(TS^\circ)$, *pre-Poisson submanifolds* are those submanifolds S such that the vector bundle $TS + \pi^\#(TS^\circ)$ has constant rank, etc.

¹⁰A *clean intersection* of two submanifolds S and L means that $S \cap L$ is a submanifold satisfying the following condition: $T(S \cap L) = TS \cap TL$. It implies, by the implicit function theorem, that for every $x \in S \cap L$, there exists open neighborhoods $U \subset S$ and $V \subset L$ such that $U \cap V$ is an open neighborhood of x in $S \cap L$. See e.g. this page or the proof of Proposition 5.26 in [Laurent-Gengoux et al., 2013].

4 A geometric perspective on the canonical Hamiltonian formalism

At the beginning of the XXth century, physicists realized that the equations of motion of quantum mechanics resemble the Hamiltonian formulation of classical mechanics. This similarity has led physicists to find a way of ‘quantizing’ existing classical physical theories in their Hamiltonian form in order to find out what would a quantum field theory look like. The program for guessing the quantum description of systems from a classical Hamiltonian formulations is called *canonical quantization* because it relies the “canonical” (i.e. Hamiltonian) form of classical mechanics. Under this perspective, classical mechanics appeared as a limit of non-relativistic quantum mechanics, formulated in terms of a Hamiltonian and of position and momenta operators. Physicists then were hoping to develop the canonical formalism associated with classical Hamiltonian mechanics to relativistic field theories (see e.g. such a justification in 1932 [Rosenfeld, 1932]).

In particular, a possible goal was to obtain quantum electrodynamics by quantizing Maxwell electromagnetism, and some quantum theory of gravity by quantizing general relativity. Unfortunately, both of those theories possess inner symmetries (gauge symmetries and coordinate invariance) which prevent to straightforwardly obtain the Hamiltonian from the Lagrangian, as is usually possible in classical mechanics. Indeed, it has been shown that if a Lagrangian is covariant under a set of symmetries – i.e. if its expression stays invariant – then the Legendre transform from the Lagrangian to the Hamiltonian is not invertible. On the contrary, one has to add several *constraints* in the hamiltonian picture to account for the non-invertibility of the Legendre transform. Existence of constraints characterize physical theories with internal symmetries such as gauge symmetries.

Several competing approaches were developed in the 1940s to tackle the problem of quantizing field theories. The *path integral* formulation of Feynman – already based on an idea by Dirac from the early 1930s – has been praised because manifest Lorentz covariance is easier to achieve than in the operator formalism of canonical quantization. Another advantage of the path integral formulation is that it is in practice easier to guess the correct form of the Lagrangian of a theory, which naturally enters the path integrals, than the Hamiltonian, that is usually derived from the Lagrangian. Peter Bergmann and Paul Dirac proposed in the late 40s-early 50s an alternative approach to quantization that sticks to the traditional, historical quantization methods that were initially developed at the beginning of the 20th century [Bergmann and Brunings, 1949, Dirac, 1950].

This canonical quantization procedure relied on obtaining first the Hamiltonian corresponding to the given Lagrangian characterizing the action principle, and then quantize the Hamiltonian as well as the various position and momenta operators, together with the several *constraints* emerging from the procedure. More generally, any smooth function f of the canonical coordinates should be sent to an operator via a linear quantization map $\mathcal{Q} : f \mapsto \mathcal{Q}(f)$ having natural properties. In this latter step, Dirac requires that the Lie bracket of operators and the Poisson – or Dirac – bracket of observables (smooth functions on the phase space) obeys the following compatibility condition:

$$\mathcal{Q}_{\{f,g\}} = \frac{1}{i\hbar}[\mathcal{Q}_f, \mathcal{Q}_g] \quad (4.1)$$

Notice however that, although classical physics seems to be a limit of quantum physics (e.g. when $\hbar \rightarrow 0$), canonical quantization is only approximate in the sense that one cannot fully deduce a quantum theory from a classical one. In particular, Dirac emphasized that the quantum theory obtained through canonical quantization should be taken as a mere possibility (among others) and the classical theory is used to develop the intuition about this quantum theory.

Although the procedure seems perfectly viable on the paper, and that the first part of the procedure is well-known, there are several drawbacks. First, and this is not restricted to canonical quantization, there is actually no unique way of quantizing a classical theory. Indeed, one usually promotes the position q^k and conjugate momenta p_k coordinates to operators Q^k, P_k on a Hilbert space, and require that their Lie Bracket is proportional to $i\hbar$ as in Equation (4.1), but there may exist alternative choice of coordinates that would thus give other quantized operator (see the introduction of [Bergmann and Goldberg, 1955]). Moreover, when one has a product of conjugates coordinates – such as $qp = pq$, say – there is no standard way of assigning an operator because the operators associated to p_k and q^k do not commute. There exists a convention specified by Weyl, which comes close to achieve this task, but a no-go theorem by Groenewold proves that there is no quantization scheme such that Equation (4.1) is satisfied at any polynomial order. This is why the canonical quantization proposed by Dirac is then usually used as an heuristics or performed only for unambiguous classical theories for which the Hamiltonian has nice properties and for which physical intuition works well to fully determine the quantum theory.

The main concern of Dirac’s quantization procedure is the treatment of *constraints*. These constraints are smooth functions which emerge as a consequence of the fact that the Lagrangian of the system is *singular* (i.e. its hessian is singular). In that case physical solutions of the equations of motions live on a submanifold of the phase space T^*Q called the *constraint surface*, which is the zero level set of these smooth functions called constraints. These are obtained when passing from the Lagrangian formalism to the Hamiltonian formalism via the Legendre transform. Bergmann and Dirac Thus, one needs to keep track of the constraints when going through quantization. As we will see, the quantization scheme \mathcal{Q} obviously sends every constraint ϕ_a to an operator, but it does not say what convention one should impose on the action of $\Phi_a = \mathcal{Q}(\phi_a)$ on the vectors of the Hilbert space. Moreover the type of the constraint – first-class or second-class – often implies different outcomes in the quantization which are difficult to handle in a practical way. Additional procedures have been developed to handle this problem which arise as soon as one wants to quantize a gauge theory: the *BV formalism* on the one hand (in the Lagrangian picture) and the *BFV/BRST formalism* (in the Hamiltonian picture)¹¹. The present chapter is devoted to study the canonical Hamiltonian formalism following the steps of Bergmann and Dirac, in order to prepare the Hamiltonian theory to be quantized. We will mostly rely on the following texts: [Sudarshan and Mukunda, 1974], [Sundermeyer, 1982], [Henneaux and Teitelboim, 1992], together with the incredibly pedagogical [Matschull, 1996] and [Rothe and Rothe, 2010]. Other useful main resources [Bergmann and Brunings, 1949, Dirac, 1950, Bergmann and Goldberg, 1955, Dirac, 1964, Śniatycki, 1974, Gitman and Tyutin, 1990] are of historical interest¹² and may also be pedagogical on particular aspects of the topic. See also [Salisbury, 2012] for a historical focus on Bergmann’s work.

4.1 Lagrangian and Hamiltonian formalism from a geometric point of view

We begin the review of Dirac’s canonical formalism with a non-relativistic physical model. Passing from classical mechanics to field theories can be done by considering that the discrete index labelling the coordinates of classical mechanics becomes continuous: fields are labelled by the space-time point at which they are evaluated. Then, let us start with a given *configuration space* represented by a n -dimensional oriented smooth manifold Q (possibly with boundary). In

¹¹BV stands for Batalin-Vilkovisky and BFV stands for Batalin-Fradkin-Vilkovisky. BRST stands for Becchi, Rouet, Stora and Tyutin.

¹²We add further historical – and possibly much less known – references here: [Bergmann, 1949, Anderson and Bergmann, 1951, Dirac, 1951, Dirac, 1958, Shanmugadhasan, 1963, Kundt, 1966, Künzle, 1969, Shanmugadhasan, 1973, Gotay et al., 1978], and [Earman, 2003] for a more philosophically leaning description.

this section the points of Q are denoted q – instead of x . The local coordinate functions on Q are denoted by q^i – instead of x^i – and can express the position of several particles, the length of a spring, the charge of a capacitor etc. That is why they are called *generalized coordinates*. Let us now fix a trivializing chart U of both TQ and T^*Q , admitting local coordinates q^i on the base U . The tangent bundle TQ over U admits fiberwise coordinate functions $v^i : TQ \rightarrow \mathbb{R}$ that are a mere rewriting of the constant covector fields on Q denoted dq^i . In particular for every tangent vector $X \in T_qQ$, $v^i(X) = v^i(X^j \frac{\partial}{\partial q^j}) = X^i$. That is why we will often denote tangent vectors at q as $v \in T_qQ$, so that by abuse of notation, we would identify the components of v in the basis $\frac{\partial}{\partial q^i}$ with v^i .

The cotangent bundle T^*Q over U also admits fiberwise coordinate functions denoted p_i and defined as expected: $p_i(\xi) = p_i(\xi_j dq^j) = \xi_i$ for any covector field ξ . For this reason, the coordinates p_i can be identified to the locally defined constant vector fields $\frac{\partial}{\partial q^i}$. In particular we set $p_i(v^j) = \delta_i^j$ so that the p_i are the dual coordinates to the v^i , explaining why the former are called *conjugate momenta*. This also justifies that we call T^*Q the *phase space* – sometimes denoted P – since it contains the configurations as well as the momenta of the configuration space Q . We will often denote covector fields as the letter p , so that a point in the cotangent bundle T^*Q would be denoted (q, p) . By abuse of notation, we identify the components of p (resp. v) in the basis dq^i (resp. $\frac{\partial}{\partial q^i}$) with p_i (resp. v^i). Since the tangent and cotangent bundles need not be trivial vector bundles, both v^i and p_i are only defined locally on Q . More precisely, the coordinates q^i are local coordinates on the trivializing neighborhood U of q , which in turn implies that the coordinate v^i and p^i are fiberwise linear coordinates globally defined on the fiber.

Example 4.1. The cotangent bundle represent the natural setup to do Hamiltonian mechanics. Let us illustrate this property by analyzing the pendulum (of mass m and length L) from a Poisson/symplectic geometry perspective. The physical system is parametrized by the angle θ so that we set the space of all possible angles – i.e. the *configuration space* – to be the circle S^1 . The conjugate momentum to the generalized coordinate $q = \theta$ is denoted p so that it is interpreted as the fiberwise linear coordinate on the phase space T^*S^1 . The symplectic structure on this cotangent bundle is the standard one, i.e. $\omega = dp \wedge dq$, where $q = \theta$. The corresponding non-degenerate Poisson structure on T^*S^1 is thus given by:

$$\{f, h\} = \frac{\partial f}{\partial q} \frac{\partial h}{\partial p} - \frac{\partial f}{\partial p} \frac{\partial h}{\partial q}$$

for every two smooth functions $f, h \in \mathcal{C}^\infty(T^*S^1)$.

Let us define the following smooth function on T^*S^1 :

$$H = \frac{p^2}{2mL^2} + mgL(1 - \cos(\theta)) \quad (4.2)$$

where g is a constant positive parameter that may be fixed at 9,8 if one wants to reproduce the gravitational force equivalent. We call this function (4.2) the "Hamiltonian of the system" and compute its hamiltonian vector field $X_H \in \mathfrak{X}(T^*S^1)$:

$$X_H = \{H, -\} = -\frac{p}{mL^2} \frac{\partial}{\partial \theta} + mgL \sin(\theta) \frac{\partial}{\partial p} \quad (4.3)$$

An *integral curve* of the vector field $-X_H$ is a smooth path $\gamma : \mathbb{R} \rightarrow T^*S^1, t \mapsto (\theta(t), p(t))$ which is such that the tangent vector $\dot{\gamma}(t) = \dot{\theta}(t) \frac{\partial}{\partial \theta} + \dot{p}(t) \frac{\partial}{\partial p}$ at the point $\gamma(t) = (\theta(t), p(t))$ is equal to $-X_H|_{(\theta(t), p(t))}$. Alternatively, it corresponds to the level set of the smooth function H .

By isolating the two components $\dot{\theta}(t)$ and $\dot{p}(t)$ forming $\dot{\gamma}(t)$ at the point $\gamma(t) = (\theta(t), p(t))$ and equating them to that of $-X_H$ at the same point, one has, for every t :

$$\begin{aligned}\dot{\theta}(t) &= -X_H(\theta) = \{\theta, H\} = \frac{\partial H}{\partial p} \\ \dot{p}(t) &= -X_H(p) = \{p, H\} = -\frac{\partial H}{\partial \theta}\end{aligned}$$

Thus, the integral curves of the vector field $-X_H$ are precisely those path $\gamma : \mathbb{R} \rightarrow T^*S^1$ whose components $\theta(t)$ and $p(t)$ satisfy the Hamilton equations of motion. This implies that such integral curves are the physical solutions of the Hamilton equations which means that, starting from a point (q_0, p_0) on the phase space T^*S^1 , the physical motion of the pendulum obliges to follow the integral curve of the vector field $-X_H$ passing through (q_0, p_0) . More abstractly, we say that $-X_H$ points towards the flow of physical time¹³. Drawing such integral curves using the expression (4.3) gives the well-known phase portrait, Fig. 20.

We have thus seen in Example 4.1 that the mathematics developed in Poisson geometry is well-adapted to describe physical systems in the Hamiltonian formalism. However, this was only possible because every point of the phase space could be used as an initial condition. Sometimes in physics, it may happen that not every point of the phase space can be chosen to be a set of initial conditions. In that case one cannot straightforwardly apply Hamiltonian formalism to the model, and a more refined formalism is required: *constrained Hamiltonian formalism*. We will spend the rest of this section on this topic.

Definition 4.2. A Lagrangian is a fiberwise convex smooth function $L \in C^\infty(TQ)$ on the tangent bundle of Q . By fiberwise convex, we mean that, for every $q \in Q$, the function $L(q, -) : T_qQ \rightarrow \mathbb{R}$ is a smooth convex function, i.e. it is such that its Hessian symmetric matrix (written in local coordinates):

$$\mathcal{H}_{ij}(q, v) = \frac{\partial^2 L}{\partial v^i \partial v^j}(q, v) \quad (4.4)$$

has non negative determinant for every $v \in T_qQ$.

Remark 4.3. In a different coordinate chart V , the change of basis $v^i \mapsto v'^j$ is given by a section $g : U \cap V \rightarrow GL_n(\mathbb{R}^n)$ taking values in the $n \times n$ invertible real matrices. In particular, such a section has either positive determinant or negative determinant but cannot pass through 0. The determinant of the Hessian then stays of the same sign because whatever the sign of $\det(g)$, it comes to the square in the expression of the Hessian when we change the basis.

Recall that, here, we consider that $v \in T_qQ$ and we identify the coordinate functions $v^i : T_qQ \rightarrow \mathbb{R}$ with the components of v in the basis $\frac{\partial}{\partial q^i}$. To any smooth path $\gamma : \mathbb{R} \rightarrow M$, one can associate a tangent vector at the point $\gamma(t)$, which we denote $\dot{\gamma}(t) \in T_{\gamma(t)}M$ (see Section 1.1). One can then evaluate the Lagrangian function along this path: $t \mapsto L(\gamma(t), \dot{\gamma}(t))$. A priori, one can always pick up any kind of path on Q , but physicists have a recipe to determine which kind of path would correspond to the time evolution of the physical system whose state is encoded by the generalized coordinates q . Indeed, such a path γ should satisfy some differential equations called the *Euler-Lagrange equations*, under appropriate boundary conditions. Any other choice of path would be considered as non-physical. They proceed as follows: the (non-relativistic) physical model is characterized by a so-called *action*, which depends exclusively on the choice of path γ :

$$S(\gamma) = \int_{\mathbb{R}} L(\gamma(t), \dot{\gamma}(t)) dt$$

¹³The minus sign comes down to the choice of defining the hamiltonian vector field of a smooth function f as $\{f, -\}$ and not as $\{-, f\}$, although the latter convention exists.




Figure 19: Phase portrait of the pendulum build from a purely symplectic/Poisson geometry perspective. The horizontal axis represents the angular coordinate $q = \theta$ between $-\pi$ and π , while the vertical axis represents the conjugate momenta p . The arrow heads represent the direction of $-X_H$ (flow of physical time) and the lines its corresponding integral curves. The *separatrix* is actually made of four submanifolds: 2 points (singular leaves) at $\theta = \pm\pi$ and $p = 0$, while the upper (resp. lower) red line is directed toward the right (resp. the left) but never reaches π (resp. $-\pi$). There is an additional singular leaf at $(0, 0)$. Hence this phase portrait is indeed a singular foliation, integrating the distribution generated by $-X_H$. Picture taken from Wolfram Alpha.

Often the path admits well-defined boundary conditions so the integral converges. To stick with physicists' notation, we will now write the time dependency of the Lagrangian with respect to the chosen path as $L(q, \dot{q})$ instead of $L(\gamma(t), \dot{\gamma}(t))$, where \dot{q} denotes the time derivative of the generalized coordinate $q = \gamma(t)$ at time t , which geometrically corresponds to the vector $\dot{\gamma}(t)$

tangent to the curve γ at time t .

Assuming that smooth path γ corresponding to physical evolution are extrema of the action – i.e. stationary points, one requires that an infinitesimal variation of the action with respect to an infinitesimal change of path would vanish if the original path is a physical path. More precisely, assume that γ_0 is a smooth path in Q corresponding to a physical evolution of the system, then $S(\gamma_0)$ should be an extremum of the function S , and thus the infinitesimal variations of S around γ_0 should be zero:

$$0 = \delta S = \int_{\mathbb{R}} \delta L dt$$

where the variation should be understood to be taken at γ_0 (stationary point of the action). Computing the variation of L with respect to infinitesimal change of path – i.e. with respect to coordinates q and v – and with respect to the *fixed* boundary conditions¹⁴, gives the following identity :

$$\delta S = - \sum_{i=1}^n \int_{\mathbb{R}} E_i(q, \dot{q}, \ddot{q}) \delta q^i dt$$

where the $E_i(q, \dot{q}, \ddot{q})$ are defined as:

$$E_i(q, \dot{q}, \ddot{q}) = \frac{d}{dt} \frac{\partial L(q, \dot{q})}{\partial v^i} - \frac{\partial L(q, \dot{q})}{\partial q^i} \quad (4.5)$$

for every $1 \leq i \leq n$. Hence, a smooth path γ_0 corresponding to a physical evolution of the system (given appropriate initial state and boundary conditions), being a stationary point of the action, should make Equation (4.5) vanish when $(q, \dot{q}) = (\gamma_0(t), \dot{\gamma}_0(t))$.

In other words, a path corresponding to a physical evolution of the system should necessarily satisfy the infamous *Euler-Lagrange equations*:

$$\frac{d}{dt} \frac{\partial L}{\partial v^i} - \frac{\partial L}{\partial q^i} = 0 \quad (4.6)$$

for every $1 \leq i \leq n$. Conversely, we will consider that solutions of these equations – i.e. smooth paths $\gamma : \mathbb{R} \rightarrow M$ such that $(\gamma(t), \dot{\gamma}(t))$ are solutions of the Euler-Lagrange equations – are precisely the paths characterizing physical evolution of the system. Now, since we assume that the Lagrangian does not have explicit time dependency, expanding the time derivative in the Euler-Lagrange equations (4.6) gives the following equivalent set of $2n$ first-order differential equations

$$v^i = \dot{q}^i \quad \text{and} \quad \mathcal{H}_{ij}(q, v) v^j = \frac{\partial L}{\partial q^i} - \frac{\partial^2 L}{\partial v^i \partial q^j} v^j \quad (4.7)$$

where we have assumed that some path $\gamma : \mathbb{R} \rightarrow TQ$ defines a solution, so that $(q(t), \dot{q}(t)) = (\gamma(t), \dot{\gamma}(t))$. Replacing the velocities v^i by their expression \dot{q}^i on the physical path, we obtain a set of n second-order differential equations:

$$\mathcal{H}_{ij}(q, \dot{q}) \ddot{q}^j = \frac{\partial L}{\partial q^i} - \frac{\partial^2 L}{\partial v^i \partial q^j} \dot{q}^j \quad (4.8)$$

that form necessary conditions for a path $t \mapsto \gamma(t) = q(t)$ to satisfy in order to be a physical path.

One then sees that the accelerations in Equation (4.8) are uniquely solvable in terms of the positions and the velocities if and only if the Hessian matrix $\mathcal{H}_{ij}(q, \dot{q})$ is invertible, i.e. if $\det(\mathcal{H}_{ij}(q, \dot{q})) \neq 0$. When the Hessian is invertible, one can apply its inverse to both sides of Equation (4.8) and obtain a set of n second-order differential equations of the form:

¹⁴Another variational principle – the Weiss action principle – allows the boundaries to vary and has its own merits, see Chapter 3 of [Sudarshan and Mukunda, 1974].

$\ddot{q}^i = \text{something depending only on } q \text{ and } \dot{q}$

Then the theory of ordinary differential equations says that, given a set of initial conditions, there is a unique solution of this Cauchy problem (at least) in a small neighborhood of these initial conditions. In other words, it means that the time evolution of the physical system – the evolution of the couple (q, \dot{q}) – is guaranteed to depend *only* on the initial conditions. However, when the Hessian matrix has vanishing determinant, the left hand side of Equation (4.8) vanishes so that the Cauchy problem does not admit a unique solution. The accelerations cannot be solved with respect to the velocities and some stay underdetermined. It appears that in this Lagrangian picture we are losing the well-known determinism of classical physics (which can be recovered in the Hamiltonian picture). Such a situation where the Hessian is not invertible happens when the Lagrangian admits local symmetries which involve *arbitrary functions of time*:

Definition 4.4. *A symmetry of a physical system is a diffeomorphism of the configuration space Q preserving the form of the equations of motion; they are said to be covariant under this symmetry. A gauge transformation of a physical system is a family $(\varphi_t)_t$ of local symmetries φ_t that can be prescribed independently (but smoothly) at each time t . Accordingly, a gauge theory is a physical theory in which the general solution to the equations of motion contains arbitrary functions of time.*

Remark 4.5. The second statement of Definition 4.4 can be explained from the following observation: gauge transformations are parametrized by arbitrary smooth functions of time, as opposed to rigid symmetry transformations. This has the following consequence: an initial state gives rise to several possible arbitrary different time evolutions, hence the second statement. See Chapter 3 of [Henneaux and Teitelboim, 1992] for a thorough treatment of gauge transformations.

Example 4.6. This example is taken from section 1.2 of [Gitman and Tyutin, 1990]. Let the configuration space be $Q = \mathbb{R}^2$ and let the Lagrangian be $L = \frac{1}{2}(v_x - y)^2$. Then the Euler-Lagrange equations amounts to only one equation:

$$\dot{x} = y$$

There is not enough equations to guarantee the unicity of the solutions, given a set

$$(x(0), v_x(0), y(0), v_y(0)) = (\alpha, \beta, \beta, \gamma)$$

of initial condition, for the general solution of the equation of motion is:

$$x(t) = \alpha + \beta t + \frac{1}{2}\gamma t^2 + \int_0^t \varphi(\tau) d\tau \quad \text{and} \quad y(t) = \beta + \gamma t + \varphi(t)$$

where φ is a smooth function of time such that $\varphi(0) = \varphi'(0) = 0$, and is absolutely arbitrary in other aspects. This function then encodes a gauge symmetry.

In general a symmetry would not leave the functional form of the Lagrangian invariant (see Chapter 4 of [Sudarshan and Mukunda, 1974]). However, a sufficient condition for a diffeomorphism to be a symmetry of the system is to leave the action functional invariant up to a total derivative. This can be seen from the fact that the action, being a smooth function integrating the Lagrangian over (space)-time, does not ‘see’ any total derivative. Thus the equations of motion stay unchanged. Concerning gauge transformations, the following result has been shown in the late 1940s:

Proposition 4.7. *A Lagrangian L admitting gauge transformations has a singular Hessian \mathcal{H} .*

Example 4.8. Using the Lagrangian of Example 4.6, one has $\mathcal{H}(q, v) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$.

For a discussion about Proposition 4.7, see Appendix A of [Rothe and Rothe, 2010] which is a modern reformulation of [Bergmann and Brunings, 1949], or page 87-88 of [Sudarshan and Mukunda, 1974]. In that case, as is explained in Chapter 2 of [Rothe and Rothe, 2010], one has to dig into the constraints that the Lagrangian imposes on the system by carefully studying the null eigenvectors of the Hessian matrix. This opens the treatment of the quantization of gauge theories via the Batalin-Vilkovisky formalism. Notice however that in Dirac's canonical quantization procedure, one quantize the theory from the Hamiltonian perspective because in quantum mechanics the Hamiltonian has a central role. Let us give a bit more details on how hamiltonian mechanics enter the picture.

Definition 4.9. Let $L : TQ \longrightarrow \mathbb{R}$ be a Lagrangian (assumed to be a fiberwise convex function) and define the canonical hamiltonian to be the following function on the generalized tangent bundle¹⁵ $\mathbb{T}Q = TQ \oplus T^*Q$:

$$H_c(q, v, p) = \langle p, v \rangle_q - L(q, v) \quad (4.9)$$

where $\langle p, v \rangle_q$ denotes the pairing between T_q^*M and T_qM .

This function is called the *canonical hamiltonian* because it corresponds to the usual definition of the hamiltonian for unconstrained systems. Let U be a trivializing chart of both TQ and T^*Q and let q^i, v^i and p_i the corresponding local coordinates on the base, and on the fibers of $TQ|_U$ and $T^*Q|_U$, respectively. Since $\langle p, v \rangle_q = \sum_{i=1}^n p_i v^i$, by differentiating the canonical hamiltonian with respect to p_i one obtains:

$$v_i = \frac{\partial H_c}{\partial p^i} \quad (4.10)$$

Let us compute the derivative of H_c with respect to v^i :

$$\frac{\partial H_c}{\partial v^i} = p_i - \frac{\partial L}{\partial v^i} \quad (4.11)$$

Then, the points of $\mathbb{T}_q Q$ for which $\frac{\partial H_c}{\partial v^i} = 0$ are those such that:

$$p_i = \frac{\partial L}{\partial v^i} \quad (4.12)$$

Notice that the set of $2n$ first-order differential equations (4.7) are equivalent to the following set of equations, called the *implicit Euler-Lagrange equations*:

$$v^i = \dot{q}^i, \quad p_i = \frac{\partial L}{\partial v^i} \quad \text{and} \quad \dot{p}_i = \frac{\partial L}{\partial q^i} \quad (4.13)$$

The second one is *not* a differential equation but an algebraic one. The equivalence can be straightforwardly calculated, and the Equations (4.13) can be obtained as the variation of the following action, where the p_i have the role of Lagrange multipliers in what is called the *Hamilton-Pontryagin action* [Yoshimura and Marsden, 2006a]:

$$S = \int L(\gamma(t), v) + \langle p, \dot{\gamma}(t) - v \rangle_{\gamma(t)} dt \quad (4.14)$$

The set of Equations (4.13) can then be recasted using the canonical Hamiltonian and Equations (4.10) and (4.12):

$$\dot{q}^i = \frac{\partial H_c}{\partial p_i}, \quad \frac{\partial H_c}{\partial v^i} = 0 \quad \text{and} \quad \dot{p}_i = -\frac{\partial H_c}{\partial q^i} \quad (4.15)$$

¹⁵The denomination *generalized tangent bundle* comes from the work of Hitchin and Gualtieri on generalized geometry [Gualtieri, 2003], while it is called *Pontryagin bundle* in the work of Yoshimura and Marsden [Yoshimura and Marsden, 2006b, Yoshimura and Marsden, 2006a]

These equations descend from the variation of the following action:

$$S = \int \langle p, \dot{\gamma}(t) \rangle_{\gamma(t)} - H_c(\gamma(t), v, p) dt \quad (4.16)$$

which is actually a rewriting of Equation (4.14). Then, we see how Hamiltonian can be a very efficient way of recasting Euler-Lagrange equations (4.6) into first-order differential equations.

Under the integral sign on the right-hand side of Equation (4.16), we recognize the well-known relationship between Hamiltonian and Lagrangian. Indeed, in classical mechanics, the Hamiltonian is the *Legendre transform* of the Lagrangian. Usually the Legendre transform of a convex function $v \mapsto f(v)$ – with domain of definition I , that in the following we will take to be \mathbb{R} – is a smooth function $p \mapsto f^*(p)$ defined via evaluating the supremum of the concave function $v \mapsto pv - f(v)$ over $I = \mathbb{R}$, for each p such that this supremum is finite. Denoting I^* the subset of \mathbb{R} whose elements $p \in I^*$ are such that $\sup_{v \in \mathbb{R}} (pv - f(v)) < +\infty$, one sets:

$$f^*(p) = \sup_{v \in \mathbb{R}} (pv - f(v)) \quad (4.17)$$

Under the assumption that the derivative of f is invertible there is an explicit formula for f^* :

$$f^*(p) = pv - f(v) \Big|_{v=(f')^{-1}(p)} \quad (4.18)$$

where here one really should understand v and p as real numbers so it makes sense to have $(f')^{-1}(p)$. Equation (4.18) is the kind of formula one usually uses in thermodynamics [Zia et al., 2009], where Helmholtz free energy A and Gibbs free energy G are obtained by performing Legendre transforms (up to a sign) of the internal energy U and enthalpy H , respectively. There, we usually do not explicitly check that the derivative of U and H with respect to the entropy is invertible although it is implicitly used when we do the Legendre transform using Formula (4.18) instead of Formula (4.17).

In our context, we precisely chose the Lagrangian to be convex so that we can take its Legendre transform. We will slightly extend the meaning of the latter by considering that it is a map from the tangent bundle to the cotangent bundle, thus providing an explanation for the formula $p_i = \frac{\partial L}{\partial v^i}$. The Lagrangian is supposed to be a convex function, i.e. its Hessian \mathcal{H} has non-negative determinant. The Legendre transform is then performed with respect to the coordinates v^i . In geometric terms, the Legendre transform between the Lagrangian and the Hamiltonian corresponds to performing a Legendre transform of the function $L(q, -) \in C^\infty(T_q Q)$ for every q :

Definition 4.10. *The Hamiltonian H_0 is the function on (a subset of) T^*Q defined as:*

$$H_0(q, p) = \sup_{v \in T_q M} (H_c(q, v, p)) \quad (4.19)$$

whenever such supremum exists.

We assume that the supremum varies smoothly when the base point q varies, hence H_0 is a smooth function that functionally depends only on the generalized coordinates q^i and on the conjugate momenta p_i . As for the rest of the section, we will use the Legendre transform from a more geometrical point of view. We will adopt an ‘in-between’ perspective where we mostly work in local coordinates over a trivializing chart $U \subset M$ to treat hamiltonian constraints (as physicists do), and at the same time we will adopt from time to time a global coordinate-free perspective to address issues that will inevitably arise along the way (as mathematicians do). We will mostly rely on the following resources: on the geometrical side, the Legendre transform

had been known since at least the seminal book of Abraham and Marsden [Abraham and Marsden, 1987] and had been investigated by Tulczyjew [Tulczyjew, 1977], as well as Marsden and Yoshimura [Yoshimura and Marsden, 2006b, Yoshimura and Marsden, 2006a] (see also most of references therein), while on the physical side there exist well established sources on constrained hamiltonian systems [Gitman and Tyutin, 1990], [Henneaux and Teitelboim, 1992], [Rothe and Rothe, 2010], see also these notes.

To provide a geometric flavour to this discussion, let us then generalize the Legendre transform to the tangent and cotangent bundles:

Definition 4.11. *The Legendre transform or fiber-derivative is a base point preserving smooth map from TQ to T^*Q (but not necessarily a vector bundle morphism) given by:*

$$\begin{aligned} \mathcal{L} : TQ &\longrightarrow T^*Q \\ (q, v) &\longmapsto \left(q, p : w \mapsto \left. \frac{d}{ds} \right|_{s=0} L(q, v + sw) \right) \end{aligned}$$

On the right-hand side, the element w is a tangent vector at q . Thus, the element p – image of v via \mathcal{L} – is a linear form on T_qQ , sending w to $\left. \frac{d}{ds} \right|_{s=0} L(q, v + sw)$. This definition does not depend on the local coordinates, but the function \mathcal{L} can be decomposed on the local frame dq^i as:

$$\mathcal{L}(q, v) = \sum_{i=1}^n \mathcal{L}_i(q, v) dq^i = \sum_{i=1}^n \frac{\partial L}{\partial v^i} dq^i$$

so that $\mathcal{L}_i = \frac{\partial L}{\partial v^i} \in \mathcal{C}^\infty(TQ)$ symbolize the components of the function and we indeed obtain again that $p_i(\mathcal{L}(q, v)) = \frac{\partial L}{\partial v^i}(q, v)$. As a base point preserving smooth map from TQ to T^*Q , the Legendre transform gives rise to a submanifold $N_{\mathcal{L}}$ of the total generalized tangent bundle $\mathbb{T}Q = TQ \oplus T^*Q$, defined as:

$$N_{\mathcal{L}} = \{(q, v, p) \mid (q, p) = \mathcal{L}(q, v)\} \subset \mathbb{T}Q$$

This submanifold is the disjoint union over the points $q \in Q$ of the graphs of the smooth maps $\mathcal{L}(q, -) : TQ \longrightarrow T^*Q$, i.e. $N_{\mathcal{L}} \cap \mathbb{T}_qQ = \text{Gr}(\mathcal{L}(q, -))$.

Seeing the Legendre transform from this geometrical viewpoint allows to retrieve the usual definition:

Lemma 4.12. *The submanifold $N_{\mathcal{L}}$ is the set of points $(q, v, p) \in \mathbb{T}Q$ such that v is a critical point of the smooth function $x \mapsto \langle p, x \rangle_q - L(q, x)$.*

Proof. Let $(q, v, p) \in \mathbb{T}Q = TQ \oplus T^*Q$, then $p_i(\mathcal{L}(q, v)) = \frac{\partial L}{\partial v^i}(q, v)$ if and only if v satisfies $\frac{\partial(\langle p, v \rangle_q - L(q, v))}{\partial v^i} = 0$, i.e. if and only if v is a critical point of $x \mapsto \langle p, x \rangle_q - L(q, x)$. \square

By Lemma 4.12, the restriction of the function H_c to $N_{\mathcal{L}}$ does not depend on v because for any given choice of pair (q, p) , any critical point v of $x \mapsto \langle p, x \rangle_q - L(q, x)$ gives the same critical value (because the supremum is unique). So, in particular:

$$\left. \frac{\partial H_c}{\partial v^i} \right|_{N_{\mathcal{L}}} = 0$$

This equation implies that the canonical hamiltonian induces a smooth function defined on the image of the Legendre transform $\text{Im}(\mathcal{L}) \subset T^*Q$:

$$H_0(q, p) = H_c(q, v, p) \quad \text{for any triple } (q, v, p) \in N_{\mathcal{L}} \tag{4.20}$$




Figure 20: The generalized tangent bundle $\mathbb{T}Q$ can be symbolically represented on a 2-dimensional sheet of paper with two axis, one for TQ and one for T^*Q . Then to any base-point preserving smooth function from the former to the latter corresponds a submanifold of $\mathbb{T}Q$. In the present case, for every $q \in Q$ the Legendre transform $\mathcal{L}(q, -)$ defines a graph – i.e. a submanifold – in the fiber T_qQ . The submanifold $\text{Gr}(\mathcal{L}(q, -))$ varies smoothly from fiber to fiber so that their union form a submanifold of the vector bundle $\mathbb{T}Q$, that is to say: $N_{\mathcal{L}} = \sqcup_{q \in Q} \text{Gr}(\mathcal{L}(q, -))$.

This latter equation can be summarized as:

$$H_0 = H_c|_{N_{\mathcal{L}}} \quad (4.21)$$

The notation is not innocent since Lemma 4.12 tells us that that H_0 is precisely the smooth function H_0 defined in Equation (4.19).

The Hamiltonian H_0 is not defined on the entirety of the cotangent bundle, except if the function \mathcal{L} is invertible – this condition is usually called the *Legendre condition* [Blohmann, 2023]. In that case:

$$H_0(q, p) = \langle p, \mathcal{L}^{-1}(q, p) \rangle_q - L(q, \mathcal{L}^{-1}(q, p)) \quad (4.22)$$

When it is not invertible, it is still possible to have an explicit expression for H_0 in terms of q and p but this requires to introduce local sections of the Legendre transform, see Equation (4.26). The condition for \mathcal{L} to be invertible goes down to the non-vanishing of the determinant of its Jacobian matrix $\mathcal{J}(q, v) = \left(\frac{\partial \mathcal{L}_i}{\partial v^j}\right)_{i,j}$. But this amounts to the non-vanishing of the determinant of the Hessian of the Lagrangian, for:

$$\frac{\partial \mathcal{L}_i}{\partial v^j} = \frac{\partial^2 L}{\partial v^i \partial v^j}$$

In other words, $\mathcal{J} = \mathcal{H}$ and, given the role of the Hessian matrix in Euler-Lagrange equations, we conclude that being able to solve for the accelerations in the Euler-Lagrange equations (4.8)

(the Hessian being invertible) and being able to solve for the velocities \dot{q}^i in terms of the positions q^j and the momenta p_j (the Jacobian matrix being invertible) are precisely equivalent. Then, in light of the discussion following Equation (4.8), one concludes that:

Proposition 4.13. *When the Lagrangian admits gauge transformations, the Legendre transform is not invertible.*

Example 4.14. Let us use Example 1 of [Rothe and Rothe, 2010], p. 8. The configuration manifold is $Q = \mathbb{R}^2$, and the Lagrangian is:

$$L(q, v) = \frac{1}{2}v_x^2 + v_x y + \frac{1}{2}(x - y)^2$$

where x, y are the standard coordinates on \mathbb{R}^2 and v_x, v_y are those on the tangent space. Fix $q = (x, y) \in Q$, then the Hessian of L is computed using Equation (4.4):

$$\mathcal{H}(q, v) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

This is obviously a singular matrix, which means that L is a singular Lagrangian, i.e. it admits a gauge transformation, given by the following transformations:

$$\delta x = \epsilon_x(t) \quad \text{and} \quad \delta y = \epsilon_y(t) \quad \text{such that} \quad \epsilon_y = \epsilon_x - \dot{\epsilon}_x$$

Since the Hessian is singular, we expect by the above discussion that the Legendre transform is not bijective. Indeed, applying the definition of the Legendre transform, one has:

$$\mathcal{L}_x(q, v) = \mathcal{L}(q, v)(\partial_x) = v_x + y \quad \text{and} \quad \mathcal{L}_y(q, v) = \mathcal{L}(q, v)(\partial_y) = 0$$

Then, we obtain that:

$$\text{Im}(\mathcal{L}) = \{(q, p) \text{ such that there exists } v \in T_q Q \text{ satisfying } p = (v_x + y)dx\} \subset T^*Q$$

One can straightforwardly check that the Jacobian of the Legendre transform coincides with the Hessian of the Lagrangian.

Example 4.15. Let us use Example 2 of [Rothe and Rothe, 2010], p. 8, first studied in [Henneaux and Teitelboim, 1992]. The configuration space is $Q = \mathbb{R}^3$, and the Lagrangian is:

$$L(q, v) = \frac{1}{2}(v_y - e^x)^2 + \frac{1}{2}(v_z - y)^2$$

At a given point $q = (x, y, z)$, the Hessian matrix is given by:

$$\mathcal{H}(q, v) = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

The matrix is singular, so is the Lagrangian, which means that it admits a gauge transformation. Indeed, it is given by:

$$\delta x = e^{-x} \frac{d^2}{dt^2} \alpha(t), \quad \delta y = \frac{d}{dt} \alpha(t) \quad \text{and} \quad \delta z = \alpha(t)$$

for any smooth function of the time $\alpha(t)$. The Legendre transform should be singular as well. It is given by:

$$\mathcal{L}_x(q, v) = 0, \quad \mathcal{L}_y(q, v) = v_y - e^x \quad \text{and} \quad \mathcal{L}_z(q, v) = v_z - y$$

Then, we obtain that:

$$\text{Im}(\mathcal{L}) = \{(q, p) \text{ such that there exists } v \in T_q Q \text{ satisfying } p = (v_y - e^x)dy + (v_z - y)dz\} \subset T^*Q$$

One can straightforwardly check that the Jacobian of the Legendre transform coincides with the Hessian of the Lagrangian.

4.2 Hamiltonian under constraints

Let us now dwell into the case where \mathcal{L} is possibly not invertible, by assuming however that the rank of the smooth function $\mathcal{L}(q, -) : T_q Q \rightarrow T_q^* Q$ is constant over Q , and we denote this rank $R_{\mathcal{L}}$, for some $1 \leq R_{\mathcal{L}} \leq n$. Let fix $q \in Q$ and let U be a trivializing chart of TQ (and hence of T^*Q as well). Let $v \in T_q Q$, then there exists a reindexing of the coordinates q^i (and thus of the coordinates v^i and p_i) such that:

1. the first $R_{\mathcal{L}}$ coordinates are labelled with a latin index from the beginning of the alphabet $1 \leq a \leq R_{\mathcal{L}}$, while the last $n - R_{\mathcal{L}}$ coordinates are labelled with a greek index from the beginning of the alphabet $R_{\mathcal{L}} + 1 \leq \alpha \leq n$, and
2. the minor $(\frac{\partial \mathcal{L}_a}{\partial v^b})_{1 \leq a, b \leq R_{\mathcal{L}}}$ of the Jacobian matrix \mathcal{J} is non-singular at (q, v) ¹⁶.

In other words, the $R_{\mathcal{L}}$ functions $\mathcal{L}_a \in \mathcal{C}^\infty(TQ|_U)$ are functionally independent in some open neighborhood $V \subset TQ$ of the point (q, v) ¹⁷. Then, the remaining $n - R_{\mathcal{L}}$ functions \mathcal{L}_α are functionally dependent on the former: for each $R_{\mathcal{L}} + 1 \leq \alpha \leq n$ and each base point q , there exists a functional relationship which smoothly depend on q :

$$\mathcal{L}_\alpha = \psi_\alpha(q, \mathcal{L}_a) \quad (4.23)$$

where we understand that each functional ψ_α depends on potentially all the \mathcal{L}_a . This argument is an adaptation of the proof of the Rank theorem in [Lee, 2003], see in particular Equation (7.9). Notice that the relabelling of coordinates utterly depends on the chosen tangent vector $(q, v) \in T_q Q$ for the functions \mathcal{L}_i may vary a lot over $TQ|_U$. Hence, the functional dependency (4.23) is in theory only defined locally, in the neighborhood of a given tangent vector, while at another point, we may have another reindexing and correspondingly another dependency. Moreover, the choice of a different minor in $\mathcal{J} = \mathcal{H}$ – equivalently, a different ordering at the same point (q, v) – gives a different set of independent functions \mathcal{L}_a , and thus different functions ψ_α . However, the number of independent functions would always stay equal to $R_{\mathcal{L}}$.

For every $q \in Q$ let us set $\Gamma_q = \text{Im}(\mathcal{L}(q, -))$ and $\Gamma = \text{Im}(\mathcal{L}) = \bigcup_{q \in Q} \Gamma_q$; it is a subset of T^*Q and we will now study its property. Any covector (q, p) lying in the subspace Γ_q satisfies:

$$p_i = \mathcal{L}_i(q, v) = \frac{\partial L}{\partial v^i}(q, v) \quad (4.24)$$

for some $(q, v) \in T_q Q$ and a local choice of coordinates q^i, v^i, p_i . Fix a covector $(\tilde{q}, \tilde{p}) \in \Gamma$ and a preimage (\tilde{q}, \tilde{v}) through the Legendre transform. Then from the discussion leading to Equation (4.23), there exists an open neighborhood $V \subset TQ$ of (\tilde{q}, \tilde{v}) and a reindexing of the coordinates q^i (and thus of the coordinates p_i as well) in two sets such that the coordinates of any covector $(q, p) \in \mathcal{L}(V)$ satisfy:

$$p_a = \mathcal{L}_a \quad \text{and} \quad p_\alpha = \psi_\alpha(q, p_a)$$

This is a mere rewriting of Equation (4.24), where we have replaced the terms \mathcal{L}_i by p_i since they coincide on Γ . Moreover, we have used Equation (4.23) to p_α in terms of the p_a and wrote $\psi_\alpha(q, p_a)$ instead of $\psi_\alpha(q^i, p_a)$ for simplicity.

This latter set of equations is a priori only valid on $\mathcal{L}(V)$. However, since on the open set V the functions \mathcal{L}_a are independent and coincide with the p_a on $\mathcal{L}(V)$, one can see the

¹⁶Since the matrix $\mathcal{J} = \mathcal{H}$ is symmetric, it is always possible to isolate such a minor.

¹⁷Here, by ‘functionally independent’ we mean that the only dependence is a minimal, or ‘trivial’ one, i.e. of the form $f^a \mathcal{L}_a = 0 \implies f^a = \sigma^{ab} \mathcal{L}_b$ with $\sigma^{ab} = -\sigma^{ba}$.

ψ_α as functions of p_α and locally extend them outside $\mathcal{L}(V)$ by replacing \mathcal{L}_α by p_α in their argument. See Equation (7.9) in the proof of the Rank theorem in [Lee, 2003] to understand the dependency of ψ_α in terms of independent functions. Let W be such a small neighborhood of (\tilde{q}, \tilde{p}) on which we formally extend these functions $\psi_\alpha \in \mathcal{C}^\infty(W)$ (it needs not contain the whole of $\mathcal{L}(V)$). Then one can define the following set of smooth functions on W :

$$\phi_\alpha(q, p) := p_\alpha - \psi_\alpha(q, p_\alpha) \quad \text{for every } R_{\mathcal{L}} + 1 \leq \alpha \leq n \quad (4.25)$$

called *primary constraints*. In particular these functions only depend on the generalized coordinates and on (part of) the conjugate momenta. The adjective *primary* denotes a further distinction between additional constraints that we will discuss next. The functions ϕ_α actually emerge naturally in the proof of the Rank theorem in [Lee, 2003]. Notice that the choice of a different minor in $\mathcal{J} = \mathcal{H}$ gives different independent coordinates and thus different primary constraints.

Remark 4.16. For reasons that will soon become clear, the triple $(W, p_\alpha, \phi_\alpha)$ is called a *constrained chart adapted to (\tilde{q}, \tilde{p})* (often we will omit to mention the dependency of these data on the original choice of point (\tilde{q}, \tilde{p})). Since the definition of such charts depend on the choice of preimage of (\tilde{q}, \tilde{p}) , every point of Γ might admit as many adapted constrained charts as it possesses preimages.

Example 4.17. Let draw on Example 4.6 to explain what a primary constraint looks like in that case. The Lagrangian being $L = \frac{1}{2}(v_x - y)^2$, one obtains:

$$\mathcal{L}_x = \frac{\partial L}{\partial v_x} = v_x - y \quad \text{and} \quad \mathcal{L}_y = \frac{\partial L}{\partial v_y} = 0$$

Then, we do not need to reindex the coordinates here, as we see which component \mathcal{L}_i is not independent. In particular, we have $\psi = 0$ and thus only one primary constraint, which is $\phi = p_y$. A similar argument holds for the Lagrangian of Example 4.14.

The choice of coordinates on Q has been made precisely so that the functions \mathcal{L}_α form a set of independent functions on V and that they parametrize the same subset of W as the first p_α coordinates (see the rank theorem [Lee, 2003]). Moreover, since each primary constraint ϕ_α involves linearly a different p_α , they form another independent set of functions, and since they altogether form an independent set of functions on W , it turns the constrained chart $(W, p_\alpha, \phi_\alpha)$ into a coordinate chart of T^*Q . Then, since the vanishing of the primary constraints is equivalent to the set of equations (4.25), we conclude that the primary constraints ϕ_α characterize the set $W \cap \mathcal{L}(V)$, in the sense that the smooth map $\Phi = (\phi_{R_{\mathcal{L}}+1}, \dots, \phi_n) : W \longrightarrow \mathbb{R}^{n-R_{\mathcal{L}}}$ has constant rank, because each constraint ϕ_α possesses a different, independent local coordinate p_α . Then, since Φ is surjective (one is free to chose any value for the p_α , whatever value for p_α has been chosen), it implies that it is a submersion (Theorem 7.14 in [Lee, 2003]).

Being the zero level set of a submersion, the set $W \cap \mathcal{L}(V)$ is a *closed embedded submanifold* of $W \subset T^*Q$ (Corollary 8.9 in [Lee, 2003] or Theorem 2.45). However, it does not imply that $W \cap \Gamma$ is an embedded submanifold of W , for Γ might be an immersed submanifold of T^*Q and have self intersections corresponding to the images through \mathcal{L} of open subsets of TQ located far from V . More precisely the primary constraints depend primarily on the choice of preimage of (\tilde{q}, \tilde{p}) . Although the matrix $\mathcal{J}(\tilde{q}, v)$ has rank $R_{\mathcal{L}}$ for every $v \in T_{\tilde{q}}M$, another choice of preimage (\tilde{q}, \tilde{v}) and of open set $V' \subset TQ|_U$ may imply another form of dependency from the components \mathcal{L}_i . That is to say: another reindexing of the coordinates q^i , as well as another dependency between the corresponding \mathcal{L}_α , leading to a redefinition of the ψ_α and hence of the primary constraints defined on another neighborhood W' of (\tilde{q}, \tilde{p}) . The vanishing of these new constraints would this turn make the set $W \cap W' \cap \mathcal{L}(V')$ – not necessarily coinciding with

$W \cap W' \cap \mathcal{L}(V)$ – a closed embedded submanifold. That observation would certainly not be sufficient to prevent Γ to be an immersed submanifold, with possible intersections. To avoid such annoying cases, physicists usually assume that the functions ϕ_α satisfy a so-called *regularity condition*:

Scholie 4.18. Regularity condition for constrained charts. *For every covector $(\tilde{q}, \tilde{p}) \in \Gamma$, and any constrained chart $(W, p_\alpha, \phi_\alpha)$ adapted to (\tilde{q}, \tilde{p}) , the subset $W \cap \Gamma$ is assumed to coincide with the zero level set of the primary constraints ϕ_α , i.e. $W \cap \Gamma = \bigcap_\alpha \phi_\alpha^{-1}(0)$.*




Figure 21: This is a situation we do not want: that different choices of preimages of (\tilde{q}, \tilde{p}) have neighborhoods V, V' whose image through the Legendre transform \mathcal{L} do not coincide in the vicinity of (\tilde{q}, \tilde{p}) . That is why we ask for the regularity condition, so that the primary constraint surface is an embedded submanifold.

This formulation is a slightly more mathematical version of that of [Batalin and Vilkovisky, 1984, Batalin and Vilkovisky, 1985], which gives physical justification for the need of the regularity condition. A consequence of the Regularity condition 4.18, for every point $(\tilde{q}, \tilde{p}) \in \Gamma$, and any constrained chart $(W, p_\alpha, \phi_\alpha)$, the coordinates (p_α, ϕ_α) form a set of local coordinates on W adapted to Γ . More precisely, the coordinates (q^i, p_α) form a local coordinate chart for Γ (because every point on $W \cap \Gamma$ can be retrieved from these data in a unique and smooth way using the smooth functions ψ_α), while the constraints ϕ_α are coordinate transverse to Γ . By Lemma 2.40, the above regularity condition on constrained charts then implies the following, more explicit version, that is the main assumption put forward by physicists (see alternative formulations on p. 7 of [Henneaux and Teitelboim, 1992]):

Scholie 4.19. Regularity condition on primary constraints. *The set $\Gamma = \text{Im}(\mathcal{L})$ is an embedded submanifold of T^*Q .*

Definition 4.20. *We call the embedded submanifold Γ (also denoted $\Gamma^{(1)}$) the primary constraint surface.*

Example 4.21. Using the Lagrangian of Example 4.14, one observes that the Legendre transform has rank 1 so the dimension of $\Gamma = \text{Im}(\mathcal{L})$ is 3 (because Q has dimension 2, to which we add 1 for the rank of \mathcal{L}). The only dependent function \mathcal{L}_α is \mathcal{L}_y and it vanishes. Thus, the only primary constraint is $\phi = p_y$, and the primary constraint surface is characterized by the vanishing of this constraint, i.e. $\Gamma = p_y^{-1}(0)$.

Example 4.22. Using the Lagrangian of Example 4.15, one observes that the Legendre transform has rank 2, so that the constraint surface Γ is a 5-dimensional submanifold of $T^*\mathbb{R}^3$. The only primary constraint is $p_x = 0$ so that $\Gamma = p_x^{-1}(0)$. On this submanifold, we have moreover $p_y = \mathcal{L}_y = v_y - e^x$ and $p_z = \mathcal{L}_z = v_z - y$. These identities – together with $p_x = 0$ – are also valid on the submanifold $N_{\mathcal{L}} \subset TQ$.

Remark 4.23. The regularity condition is widely met in physical systems, as Henneaux noticed at the very end of Section 4.4 in [Henneaux, 1990] that the regularity condition is usually fulfilled by all models of physical interest.

The regularity condition – together with the fact that the constraints ϕ_α are independent because each of them contains one and only one p_α – implies that they can be taken as local transverse coordinates to Γ . This has some quite important consequences, one of which is the fact that the ideal of functions vanishing on Γ is generated by the constraints:

Proposition 4.24. *Let $f \in C^\infty(T^*Q)$ be a smooth function that vanishes on Γ : $f|_\Gamma = 0$. Then, for every point $(\tilde{q}, \tilde{p}) \in \Gamma$ and any choice of constrained chart $(W, p_\alpha, \phi_\alpha)$ adapted to (\tilde{q}, \tilde{p}) , there exist functions $f_\alpha \in C^\infty(W)$ such that $f = \sum_\alpha f_\alpha \phi_\alpha$ on W .*

Proof. The proof is given in Theorem 1 and in the Appendix of Chapter 1 of [Henneaux and Teitelboim, 1992], or alternatively in section 3.3 of [Rothe and Rothe, 2010], where it was adopted from Chapter 8 of [Sudarshan and Mukunda, 1974], where it is shown that $f_\alpha = \frac{\partial f}{\partial p_\alpha}$. \square

Corollary 4.25. *If f and g coincide on Γ then there exist smooth functions h_α on W such that the following identity holds on W : $f = g + \sum_\alpha h_\alpha \phi_\alpha$.*


Notice that the statement of Proposition 4.24 is a local one, while in the references cited for the proof, the statement is a global one. The discrepancy comes from the fact that in physics textbooks, the constraints are defined globally over the phase space T^*Q . This is an extra assumption that is *not* a consequence of the Legendre transform. On the contrary, we have shown that using the Rank theorem, only local statement can be made on the form of the constraints. Then, while physicists usually think of constraints as a finite set of globally defined constraints, mathematicians should definitely think of them as a finitely generated subsheaf of the sheaf of smooth functions $C^\infty(M)$, which is so far locally free (but we'll see later that this condition might be broken when we define secondary constraints).

However, one can make "global" any locally defined constraint by assuming that it is extended outside W as the zero map, which can be done through the use of a bump function whose support is on W . Given that we have plenty of constrained charts, we end up with plenty of such extended constraints, defined over the whole phase space. Usually physicists assume that they are a finite number but a priori nothing can guarantee us this fact. Then, the regularity condition on the primary constraint surface reduces to the following statement found page 7 of [Henneaux and

[Teitelboim, 1992]: for every point $(\tilde{q}, \tilde{p}) \in \Gamma$ there exists a neighborhood W and a subset of all these extended constraints which can serve as transverse coordinates to Γ – and as a consequence they generate all the others constraints, as well as the ideal of functions vanishing on $W \cap \Gamma$.

The fact that Γ is an embedded submanifold of T^*Q implies in particular that the Legendre transform \mathcal{L} is a submersion. It then admits local sections¹⁸: for any point $(\tilde{q}, \tilde{p}) \in \Gamma$ and adapted constrained chart (W, p_a, ϕ_α) , there exists a smooth injective map $\nu : W \cap \Gamma \rightarrow TQ$ such that $\mathcal{L}(q, \nu(q, p)) = (q, p)$ for any $(q, p) \in W \cap \Gamma$. This map does actually depend only on $R_{\mathcal{L}}$ momenta, that we can chose to be the p_a , i.e. $\nu(q, p) = \nu(q^i, p_a)$. The image of ν are interpreted as the velocities that can be *solved* with respect to he momenta. There is thus a set of velocities that cannot be solved with respect to the momenta, and those physically correspond to the accelerations that cannot be solved with respect to the dynamical variables in Equation (4.8) (see pp. 93-94 in [Sudarshan and Mukunda, 1974] for example). This is a direct reformulation of the equivalence between the fact that the Hessian matrix of the Lagrangian is invertible if and only if the Legendre transform is invertible.

For simplicity assume (as physicists often do) that the primary constraints are globally defined, so that the section ν is globally defined. A choice of a section $\nu : \Gamma \rightarrow TM$ of the Legendre transform allows to find an explicit expression of H_0 in terms of q and p only, while we only had Equation (4.20) until now, which is defined using $N_{\mathcal{L}}$. Since H_c does not depends on ν over $N_{\mathcal{L}}$ (and hence, of the section ν), we deduce that the ‘restriction’ of the canonical hamiltonian H_c to $\nu(\Gamma) \oplus \Gamma$ gives an explicit formulation of the smooth function H_0 , as defined sloppily in (4.21). Indeed, the submanifold $\nu(\Gamma) \oplus \Gamma \subset TQ$ is by construction a submanifold of $N_{\mathcal{L}}$, so the following diagram explains how H_0 is defined as a function on $\Gamma = \text{Im}(\mathcal{L})$ only.



Let us find out a possible explicit expression of H_0 in local coordinates in a neighborhood of Γ . In a selected constrained chart (W, p_a, ϕ_α) , we have, for every $(q, p) \in \Gamma$:

$$H_0(q, p) = H_c(q, \nu(q, p), p) = p_a \nu^a(q, p) + \psi_\alpha \nu^\alpha(q, p) - L(q, \nu(q, p)) \quad (4.26)$$

Notice that we replaced the $n - R_{\mathcal{L}}$ conjugate momenta p_α by ψ_α because they are thus defined on the constraint surface Γ . Hence the hamiltonian H_0 does not depend on the $n - R_{\mathcal{L}}$ conjugate momenta p_α . By recalling that both the functions ψ_α and the section ν depends on the first $R_{\mathcal{L}}$ momenta only, we deduce that the hamiltonian H_0 does only depend on the first $R_{\mathcal{L}}$ coordinates. Then, the right-hand side of Equation (4.26), until now only valid on Γ , can be extended off to the whole of W , by adding and substracting $p_\alpha \nu^\alpha$:

$$H_0(q, p_a) = p_a \nu^a(q, p_a) + p_\alpha \nu^\alpha(q, p_a) - L(q, \nu(q, p_a)) - \phi_\alpha \nu^\alpha(q, p_a)$$

without any restriction on the p_a 's. But then, we can find an explicit expression of H_0 (we emphasized the explicit dependence in the first $R_{\mathcal{L}}$ coordinates):

$$H_0(q, p_a) = H_c(q, \nu(q, p_a), p) - \phi_\alpha \nu^\alpha(q, p_a) \quad \text{for any } (q, p) \in W \quad (4.27)$$

¹⁸To define a section of \mathcal{L} one needs only Γ to be a weakly embedded submanifold of T^*Q , because in that case one can show using Definition 2.53 that the Legendre transform \mathcal{L} defines a smooth map onto Γ , which is a necessary condition for \mathcal{L} to be a submersion. Being an immersed submanifold would certainly not be sufficient.

Notice that the primary constraints appear in the last term only because the Legendre transform is not invertible: compare for example with Equation (4.22) when it is invertible.

Because of Equation (4.11), H_0 does not depend on the choice of section ν . A careful discussion about this independence can be found in Proposition 1 (section 3.3) of [Rothe and Rothe, 2010]. Moreover, although the local expression of H_0 depends on the original choice of splitting between independent momenta p_a and dependent momenta p_α on Γ (and then ultimately on the choice of invertible minor of the matrix $\mathcal{J} = \mathcal{H}$), any other choice would give a function H'_0 that would coincide with H_0 on Γ . Some physicists emphasize that this discussion is purely local (see e.g. page 24 in [Gitman and Tyutin, 1990]) while other assume that local coordinates are actually global coordinates (i.e. they work on a vector space), so that Equation (4.27) is valid globally (see e.g. page 10 in [Henneaux and Teitelboim, 1992]). Under this assumption, the hamiltonian H_0 is a smooth function on the primary constraint surface, i.e. $H_0 \in \mathcal{C}^\infty(\Gamma)$.

Extending H_0 out of the constraint surface is actually necessary to proceed to Hamiltonian treatment of constrained systems. Indeed, there is no Poisson bracket on Γ so one cannot formally write Hamilton's equations with H_0 in their classical form. Equation (4.27) is a possible extension of H_0 off Γ in a small chart, but not the most general one because the left-hand side – hence the right-hand side as well – does not involve the last $n - R_{\mathcal{L}}$ coordinates. Replacing H_0 by a smooth function on the whole phase space would moreover solve the practical issue raised by the fact that in theory Equation (4.27) is only defined locally since sections are only local. A smooth function $H \in \mathcal{C}^\infty(T^*Q)$ which coincides with H_0 on Γ , i.e. such that $H|_\Gamma = H_0$, would be such a global smooth extension of H_0 to the whole of phase space, and would be a potential candidate to perform Hamiltonian analysis on T^*Q . On Equation (4.27) we see that H_0 can be written as the sum of a Hamiltonian and a linear combination of the primary constraints. The most general formula extending H_0 off Γ would have a similar structure:

Definition 4.26. *Assume that there is a finite number of globally defined constraints ϕ_α defining a closed embedded submanifold $\Gamma \subset T^*Q$, and let u^α be yet unspecified smooth functions on T^*Q (that physicists sometimes identify with velocities). Let H be a smooth function which coincides with H_0 on the primary constraint surface Γ . Then we define the total Hamiltonian to be the smooth function:*

$$H_T = H + u^\alpha \phi_\alpha \tag{4.28}$$

Remark 4.27. Since the parameters u^α are not (yet) fixed, and that the physical equations of motions do not depend on the (yet unfixed) parameters (see the discussion leading to Equation (4.39)), the total Hamiltonian can be considered as representing the equivalence class of all the Hamiltonian extending H outside of Γ (see section 15.2 in [Duncan and Duncan, 2012]). It might then be seen as a cocycle in a particular cohomology, yet to be found.

If Γ is not closed, the function H may not exist (see Lemma 3.72) but as mathematical physicists, we will assume such global extension always exists (this is the case in particular if we assume that the primary constraints are finite and defined globally). The choice of the map H is physically not relevant because physics only occurs on the constraint surface. We will soon see however that the primary constraints should be explicitly taken into account since their presence is necessary for the consistency of the dynamics. As implied by Proposition 4.24, another choice of hamiltonian $H \mapsto H'$ is equivalent to modifying the parameters u^α , as $H' = H + v^\alpha \phi_\alpha$. The total Hamiltonian defined in Equation (4.28) is then the most general form of Equation (4.27) valid outside Γ and by definition, $H_T|_\Gamma = H_0$. For more informations on the total Hamiltonian see e.g. the Corollary on page 31 of [Rothe and Rothe, 2010], or a similar but less general discussion on page 16 of [Dirac, 1964], or a more obscure but quite interesting approach in section 2.1 of [Gitman and Tyutin, 1990].

Remark 4.28. Obviously, if the Legendre transform is invertible, then H_0 is expressed as is Equation (4.22), and there are no primary constraints. Moreover, it is defined over the whole phase space so $H_T = H_0$, which is what is expected from a regular Hamiltonian system.

Example 4.29. We work on the primary constraint surface of Example 4.21. The canonical Hamiltonian H is given by:

$$H_c(q, v, p) = p_x v_x + p_y v_y - \left(\frac{1}{2} v_x^2 + v_x y + \frac{1}{2} (x - y)^2 \right)$$

On the constraint surface $\Gamma = \text{Im}(\mathcal{L})$, we know that $p_y = \mathcal{L}_y = 0$ and $p_x = \mathcal{L}_x = v_x + y$. By construction, these relations are also valid on $N_{\mathcal{L}}$. Thus, evaluating H_c on the latter gives:

$$H_0(q, p) = p_x(p_x - y) - \left(\frac{1}{2} (p_x - y)^2 + (p_x - y)y + \frac{1}{2} (x - y)^2 \right) = \frac{1}{2} p_x^2 - \frac{1}{2} x^2 + xy - y p_x \quad (4.29)$$

Alternatively, one would obtain this expression plugging in Equation (4.26) the following section of Γ to TQ is $\nu_x = p_x - y$ and $\nu_y = 0$. While H_0 is supposedly defined only over Γ , one can straightforwardly extend it to the whole phase space $T^*\mathbb{R}^2$ as a function H , and then define the total hamiltonian as:

$$H_T = H + u p_y \quad (4.30)$$

where $u \in \mathcal{C}^\infty(T^*\mathbb{R}^2)$ is still an unfixed smooth function acting as a parameter.

Example 4.30. We work on the primary constraint surface of Example 4.22. Evaluating the canonical Hamiltonian H_c on $N_{\mathcal{L}}$ gives:

$$H_0 = p_y(p_y + e^x) + p_z(p_z + y) - \left(\frac{1}{2} p_y^2 + \frac{1}{2} p_z^2 \right) = \frac{1}{2} p_y^2 + \frac{1}{2} p_z^2 + p_y e^x + y p_z \quad (4.31)$$

While this function is supposedly defined only over Γ , one can straightforwardly extend it to the whole phase space $T^*\mathbb{R}^2$ as a function H , so that the total Hamiltonian is:

$$H_T = H + u p_x$$

where $u \in \mathcal{C}^\infty(T^*\mathbb{R}^2)$ is still an unfixed smooth function acting as a parameter.

To justify the use of H_T , let us differentiate H_0 with respect to the canonical variables q_i and p^i . A detailed discussion about this can be found in Proposition 2 (section 3.3) of [Rothe and Rothe, 2010]. First, deriving Equation (4.26) with respect to p_a and noticing that $p_i = \frac{\partial L}{\partial v^i}$ on Γ , one obtains that the terms $p_a \frac{\partial \nu^a}{\partial p_i} + \psi_\alpha \frac{\partial \nu^\alpha}{\partial p_i}$ cancels out with $\frac{\partial L}{\partial v^j} \frac{\partial \nu^j}{\partial p_i}$ so that we obtain:

$$\frac{\partial H_0}{\partial p_a} = \nu^a + \frac{\partial \psi_\alpha}{\partial p_a} \nu^\alpha \quad (4.32)$$

We see that there is no contribution of the derivatives of ν with respect to p_a . Notice however that this observation is valid only on the primary constraint surface Γ , and thus so is Equation (4.32). By definition of ϕ_α , Equation (4.32) can be straightforwardly rewritten:

$$\frac{\partial H_0}{\partial p_a} = \nu^a - \frac{\partial \phi_\alpha}{\partial p_a} \nu^\alpha \quad (4.33)$$

Unfortunately the set of Equations (4.33) does not include the derivative with respect to the p_α since H_0 does not depend on them. However, relying on this fact and that $\frac{\partial \phi_\beta}{\partial p_\alpha} = \delta_\beta^\alpha$ on Γ , one may add a set of additional tautological equations:

$$\frac{\partial H_0}{\partial p_\alpha} = \nu^\alpha - \frac{\partial \phi_\beta}{\partial p_\alpha} \nu^\beta \quad (4.34)$$

Hence we notice that a priori Equations (4.33) and (4.34) do not involve time whatsoever.

Next, differentiating Equation (4.26) with respect to q^i and noticing that $p_i = \frac{\partial L}{\partial v^i}$ on Γ , one obtains that the terms $p_\alpha \frac{\partial \nu^\alpha}{\partial q^i} + \psi_\alpha \frac{\partial \nu^\alpha}{\partial q^i}$ cancels out with $\frac{\partial L}{\partial v^j} \frac{\partial \nu^j}{\partial q^i}$, so that we obtain:

$$\frac{\partial H_0}{\partial q^i} = -\frac{\partial \phi_\alpha}{\partial q^i} \nu^\alpha - \frac{\partial L}{\partial q^i} \quad (4.35)$$

Notice that we had replaced ψ_α by $-\phi_\alpha$ since by construction their derivative with respect to q^i coincide. Now, assume that we restrict our study to a smooth curve $\gamma : \mathbb{R} \rightarrow M$ so that $q = \gamma(t)$ and the vector field corresponding to the velocity at time t is tangent to the curve at every time t and lives in the image of the section ν , i.e. $\dot{q}(t) = \dot{\gamma}(t) = \nu(q(t), p(t))$. The image through \mathcal{L} of the path $t \mapsto (q(t), \dot{q}(t))$ defines a path in the phase space $t \mapsto (q(t), p(t))$. Then, one may add $\frac{d}{dt} \left(\frac{\partial L}{\partial v^i} \right)$ to Equation (4.35) and subtract \dot{p}_i (since they compensate one another on Γ by Equation (4.12)), to obtain:

$$\frac{\partial H_0}{\partial q^i} = -\dot{p}_i - \frac{\partial \phi_\alpha}{\partial q^i} \nu^\alpha(t) + E_i(q(t), \nu(t), \dot{\nu}(t)) \quad (4.36)$$

where $E_i(q(t), \nu(t), \dot{\nu}(t))$ is the smooth function defined in Formula (4.5), which vanishes precisely when the path is a solution of the Euler-Lagrange equations (4.6). Now, assume that the path γ is a solution of the Euler-Lagrange equations (4.6), and that we have $\nu(t) = \nu(q(t), p(t)) = \dot{\gamma}(t) = \dot{q}(t)$. Then, the image of such a path through the Legendre transform \mathcal{L} defines a path $t \mapsto (q(t), p(t))$ staying in the primary constraint surface Γ , and whose time derivative gives the infamous *Hamilton equations of motion* satisfied by q^i and p_i :

$$\dot{q}^i = \frac{\partial H_0}{\partial p_i} + \frac{\partial \phi_\alpha}{\partial p_i} \nu^\alpha \quad (4.37)$$

$$\dot{p}_i = -\frac{\partial H_0}{\partial q^i} - \frac{\partial \phi_\alpha}{\partial q^i} \nu^\alpha \quad (4.38)$$

We obtained these equations by gathering Equations (4.33), (4.34) with Equations (4.36) and reordering the terms. Notice that, due to the constraints, they do not precisely respect the usual form of Hamilton's equations of motions. We will soon see how one can recast these in this form.

Recall that, although the first Hamilton equations of motion (4.37) are mere consequences of the Legendre transform (and are valid without assuming that ν is of the form $\dot{\gamma}(t)$), the second ones (4.38) are satisfied if the Euler-Lagrange equations (4.6) are satisfied (this is a consequence, and not an equivalence). Moreover, in both case we see that, for points of T^*Q to be considered as potential candidates for physical states of the system – or equivalently, for paths to be considered physical trajectories in the phase space – they at least need to live on Γ , where the Hamiltonian is defined. It does not mean however that every point of the primary constraint surface Γ is an admissible physical state – and we will see that in general they do not. Finally, notice that Equations (4.37) and (4.38) can be recasted in a system of Equations which resembles more Hamilton equations of motions, at the cost of enforcing the constraint equations:

$$\begin{cases} \dot{q}^i = \frac{\partial}{\partial p_i} (H_0 + \phi_\alpha \nu^\alpha) \\ \dot{p}_i = -\frac{\partial}{\partial q^i} (H_0 + \phi_\alpha \nu^\alpha) \\ \phi_\alpha = 0 \end{cases}$$

where here ϕ_α is evaluated on the smooth path $(q(t), p(t))$. This set of equations is consistent with the set of equations (4.15): indeed, if one adds $\phi_\alpha \nu^\alpha$ to Equation (4.26), one obtains Equation (4.9) for $v = \nu(q, p)$. Then Equations $\frac{\partial H_c}{\partial v^i} = 0$ imply that $H_c(q, \nu(q, p), p)$ does not

depend on the section ν , or equivalently, that we are working on $N_{\mathcal{L}}$, which is alternatively said by imposing $\phi_\alpha = 0$ and $p_a = \mathcal{L}_a$.

Unfortunately, since H_0 is a priori not defined outside the primary constraint surface Γ , we cannot write the above set of equations with the help of the Poisson bracket. For this, a function defined all over the phase space would be necessary. However, the presence of the terms $H_0 + \phi_\alpha \nu^\alpha$ reminds us of the discussion surrounding Equation (4.28) where we said that replacing H_0 by any smooth function $H \in \mathcal{C}^\infty(T^*Q)$ such that $H|_\Gamma = H_0$ would lead to the same physics and, more importantly, would open the use of the Poisson bracket on the phase space. Indeed, the corollary of Proposition 3 in [Rothe and Rothe, 2010] shows that for any smooth function $H \in \mathcal{C}^\infty(T^*Q)$ such that $H|_\Gamma = H_0$, there exists smooth functions $u^\alpha \in \mathcal{C}^\infty(T^*Q)$ such that the Hamiltonian equations of motions can be recasted as:

$$\begin{cases} \dot{q}^i = \frac{\partial}{\partial p_i} (H + \phi_\alpha u^\alpha) \\ \dot{p}_i = -\frac{\partial}{\partial q^i} (H + \phi_\alpha u^\alpha) \\ \phi_\alpha = 0 \end{cases}$$

The justification comes from the fact that, since H coincides with H_0 on the constraint surface Γ , it may be written (at least locally) as $H_0 + \phi_\alpha \lambda^\alpha$ (see the proof of Proposition 4.24), and the smooth functions λ^α are so that on the constraint surface, one has $\lambda^\alpha = \frac{\partial H}{\partial p_\alpha}$. It then implies that $u^\alpha = \nu^\alpha - \lambda^\alpha$. The latter hamiltonian equations of motions are quite convenient because they are defined outside of Γ , if not on the whole phase space (when the constraints are so defined).

The fact that the primary constraints appear explicitly in the above set of equations also justifies that the correct Hamiltonian is not H , but the total Hamiltonian $H_T = H + u^\alpha \phi_\alpha$, as postulated in Equation (4.28). Indeed, denoting $\{.,.\}$ the canonical Poisson bracket on the cotangent bundle, associated to the canonical symplectic form on T^*Q , one can then recast Hamilton's equations of motions as:

$$\begin{cases} \dot{q}^i = +\frac{\partial H_T}{\partial p_i} = \{q^i, H_T\} \\ \dot{p}_i = -\frac{\partial H_T}{\partial q^i} = \{p_i, H_T\} \\ \phi_\alpha = 0 \end{cases} \quad (4.39)$$

These equations are a consequence of the extended Euler-Lagrange equations (4.15) and thus, of the original ones as well. Thus, although $H_T|_\Gamma = H_0$, the presence of the primary constraints in its definition are of utter importance. We will see later that we can find a set of equations extending (4.39) which is equivalent to the Euler-Lagrange equations. The discussion appearing in section 2.1 of [Gitman and Tyutin, 1990] is quite interesting although a bit obscure, because it justifies that although the splitting into independent conjugate momenta p_a and dependent ones p_α on Γ is not unique (one could have chosen another set of independent coordinates p_a), the hamiltonian H_0 is uniquely defined and the total hamiltonian forms a class of function 'equivalent' to that of H_0 .

4.3 The Bergmann-Dirac algorithm

The importance of the primary constraint surface in the Hamiltonian formalism of singular Lagrangian theories can be best shown after introducing some adapted notation:

Definition 4.31. *We say that two functions $f, g \in \mathcal{C}^\infty(T^*Q)$ are weakly equivalent if they coincide on Γ , and we note:*

$$f \approx g$$

For clarity, we say that they are strongly equivalent if they coincide on the whole of phase space T^*Q .

Being weakly equivalent is an equivalence relation, and this notion will be thoroughly used in the text. Since Γ is an embedded submanifold of Q defined as a level set of a set of smooth functions – the primary constraints –, it turns out that any smooth function vanishing on Γ is functionally locally dependent on the primary constraints, as Proposition 4.24 showed. Then, using the notation of Definition 4.31, one can recast equations (4.39) as:

$$\dot{q}^i \approx \{q^i, H_T\} \quad (4.40)$$

$$\dot{p}_i \approx \{p_i, H_T\} \quad (4.41)$$

The Poisson bracket has to be evaluated on Γ *after* it has been computed – i.e. we compute $\{q^i, H_0\}$ but $\{q^i, H_T\}$ and then we apply $\phi_\alpha = 0$. The total Hamiltonian thus defines (minus) the flow of time when we restrict ourselves to the primary constraint surface Γ . Equations (4.40) and (4.41) imply in turn that the total Hamiltonian computes the dynamics of any smooth function which is evaluated on any physical path sitting in Γ . More precisely, let $f \in \mathcal{C}^\infty(T^*Q)$ be any smooth function, and let (\tilde{q}, \tilde{p}) be any point of Γ . Then for small times t , and under the assumption that the undefined parameters u^α are fixed, there is a unique path $t \mapsto (q(t), p(t))$ such that:

$$\begin{cases} q(0) = \tilde{q} \\ p(0) = \tilde{p} \end{cases} \quad \text{and} \quad \begin{cases} \dot{q}^i(0) = \{q^i, H_T\}(\tilde{q}, \tilde{p}) \\ \dot{p}_i(0) = \{p_i, H_T\}(\tilde{q}, \tilde{p}) \end{cases} \quad (4.42)$$

which is contained in Γ , i.e. such that $\phi_\alpha(q(t), p(t)) = 0$ for all times t . We then define the following real numbers:

$$\dot{f}(q(t), p(t)) = \frac{\partial f}{\partial q^i}(q(t), p(t))\dot{q}^i(t) + \frac{\partial f}{\partial p_i}(q(t), p(t))\dot{p}_i(t)$$

By unicity of the Cauchy problem the value of \dot{f} only depends on the point, and not on the path. The right-hand side is not only a smooth function of the time t , but also of the base point (\tilde{q}, \tilde{p}) . Then, we can define a smooth assignment:

$$\begin{aligned} \dot{f} : \Gamma &\longrightarrow \mathbb{R} \\ (\tilde{q}, \tilde{p}) &\longmapsto \frac{\partial f}{\partial q^i}(\tilde{q}, \tilde{p})\dot{q}^i(0) + \frac{\partial f}{\partial p_i}(\tilde{q}, \tilde{p})\dot{p}_i(0) \end{aligned} \quad (4.43)$$

where $\dot{q}^i(0)$ and $\dot{p}_i(0)$ are uniquely defined by Conditions (4.42). Since the primary constraint surface is an embedded submanifold of T^*Q , the assignment (4.43) admits at least locally a smooth extension, and two such extensions coincide on Γ . Then we have the following important result about dynamics:

Lemma 4.32. *Let $f \in \mathcal{C}^\infty(T^*Q)$ be any smooth function, and let \dot{f} be any smooth extension of the associated smooth assignment (4.43); then:*

$$\dot{f} \approx \{f, H_T\}$$

Since every physical solution of the Hamilton equations (4.39) should be contained in the constraint surface Γ , it means that if one evaluates the primary constraints ϕ_α on any such physical path $t \mapsto (q(t), p(t))$, one has $\dot{\phi}_\alpha(q(t), p(t)) = 0$ because $\phi_\alpha(q(t), p(t)) = 0$ for all t . Using Lemma 4.32, this necessary condition reads:

$$\{\phi_\alpha, H_T\} \approx 0 \quad (4.44)$$

We call this equation the *persistence* of the primary constraints ϕ_α . Alternatively, by Lemma 2.58, it can be geometrically interpreted under the following condition, assuming that the primary constraints generate the ideal of functions vanishing on Γ :

for any physical path to sit in Γ , the Hamiltonian vector field X_{H_T} has to be tangent to Γ

Thus, we have to make a good choice of parameters u^α so that this condition is satisfied. Developing the total hamiltonian, computing the bracket, and eventually evaluating the result on Γ allows to rewrite Equation (4.44) under the form:

$$\{\phi_\alpha, H\} + u^\beta \{\phi_\alpha, \phi_\beta\} \approx 0 \quad (4.45)$$

Recall that H is any smooth function on T^*Q that coincides with H_0 on the constraint surface: $H|_\Gamma = H_0$. Then, solving Equation (4.45) amounts to specify the parameters so that the hamiltonian vector field X_{H_T} is tangent to Γ .

Assume that we have a set of primary constraints defined as in Equation (4.25), so that their number is $n - R_{\mathcal{L}}$, where $R_{\mathcal{L}}$ is the rank of the Legendre transform $\mathcal{L}(q, -) : T_q Q \rightarrow T_q^* Q$, assumed to be constant over Q . We let M be the $(n - R_{\mathcal{L}}) \times (n - R_{\mathcal{L}})$ square matrix whose coefficients are the following smooth functions:

$$M_{\alpha\beta} = \{\phi_\alpha, \phi_\beta\} \quad (4.46)$$

Since M is a skew-symmetric matrix, it cannot be diagonalized over \mathbb{R} , but it is possible to bring it to a block diagonal form by a special orthogonal transformation:

Proposition 4.33. *If M is a $p \times p$ antisymmetric square matrix, there is an orthogonal $p \times p$ square matrix O and a $p \times p$ square matrix Λ of the form:*

$$\Lambda = \begin{pmatrix} 0 & \lambda_1 & & & & & & & \\ -\lambda_1 & & & & & & & & \\ & & 0 & \lambda_2 & & & & & \\ & & -\lambda_2 & 0 & & & & & \\ & & & & \ddots & & & & \\ & & & & & 0 & \lambda_d & & \\ & & & & & -\lambda_d & 0 & & \\ & & & & & & & 0 & \\ & & & & & & & & \ddots & \\ & & & & & & & & & 0 \end{pmatrix}$$

such that $\Lambda = O^T M O$ and the λ_i 's are strictly positive real numbers.

Proof. An antisymmetric matrix is diagonalizable over \mathbb{C} with purely complex eigenvalues. On the field of real numbers, it would then correspond to the above description. \square

From now on, we consider that we perform the above block diagonalization at a point $x \in \Gamma$. Then, the number of zero eigenvalues depend on the rank of M at x in Γ . In particular, if $p = n - R_{\mathcal{L}}$ is an odd integer, there is at least one zero eigenvalue. We assume moreover that the rank of the matrix M – and hence of Λ – is constant over the primary constraint surface Γ . As a consequence, the number of zero eigenvalues is constant over Γ . Both the coefficients of O and of Λ are in fact smooth functions over T^*Q . While the matrix O is always invertible, it may happen that some eigenvalues λ_i may vanish in some region of the phase space. We write $O_{\alpha\beta}$ the (α, β) -th coefficient of the matrix O , and we let:

$$\phi_\beta^{(1)} = \sum_{\alpha} O_{\alpha\beta} \phi_\alpha \quad (4.47)$$

It is as if we had performed a ‘rotation’ in the space of primary constraints. Notice that we do not lose information by performing this transformation because we can always come back to the original constraints by applying O^T (we do not need to invert the matrix O). Then the linear combinations $\phi_\alpha^{(1)}$ are still constraints: they generate the same primary constraint surface Γ and the same ideal of functions in $\mathcal{C}^\infty(T^*Q)$, but they are more adapted to the problem, in the sense that the matrix Λ is the matrix of the Poisson bracket of the new primary constraints:

$$\Lambda_{\alpha\beta} = \{\phi_\alpha^{(1)}, \phi_\beta^{(1)}\}$$

One can pass from the original set of constraints to the new one by using the matrix O or its transpose. Being orthogonal, these matrices are always invertible. From now on we will mostly use the latter set of constraints and we call them *first-stage constraints* (although the denomination is not standard, but merely practical). We can refine their description by assuming that the index α is split into two families; the first $2d$ constraints satisfy only one non-trivial bracket on Γ :

$$\{\phi_{2i-1}^{(1)}, \phi_{2i}^{(1)}\} = \lambda_i \neq 0 \quad (4.48)$$

for $1 \leq i \leq d$, while the last $n - R_{\mathcal{L}} - 2d$ constraints have their Poisson bracket vanishing with *every* other constraint $\phi_\alpha^{(1)}$. For simplicity, and contrary to our previous convention, here and from now on we consider that the index α runs from 1 to $n - R_{\mathcal{L}}$.

The persistence Equation (4.44) should now apply to the new constraints, and is then equivalent to the following necessary condition:

$$\{\phi_\alpha^{(1)}, H_T\} \approx 0 \quad (4.49)$$

Since the coefficients u^β in Equation (4.45) are still unspecified, and since the set of constraints ϕ_α is equivalent to the set $\phi_\alpha^{(1)}$, we can then rewrite Equation (4.49) as:

$$\{\phi_\alpha^{(1)}, H\} + v^\beta \{\phi_\alpha^{(1)}, \phi_\beta^{(1)}\} \approx 0 \quad (4.50)$$

where the coefficients v^β are yet unspecified smooth functions obtained from the u^α via the identity $v^\alpha = \sum_\alpha (O^T)_{\alpha\beta} u^\beta$ because they appear as such in the total Hamiltonian, when written with the first-stage constraints. Indeed, if one multiplies the primary constraints ϕ_α by O , one should multiply the coefficients u^α by $O^{-1} = O^T$. Then we have the following possible cases:

1. either the matrix M – and hence Λ – is invertible on the primary constraint surface (or at least locally where the primary constraints are defined). Then one can uniquely determine *all* parameters v^β in Equation (4.50), by setting (strongly):

$$v^\beta = -(\Lambda^{-1})^{\beta\alpha} \{\phi_\alpha^{(1)}, H\} \quad (4.51)$$

These are the unique coefficients making the persistence equation (4.49) valid. The total Hamiltonian H_T is then uniquely specified by plugging these coefficients in the following equation:

$$H_T = H + v^\beta \phi_\beta^{(1)} = H + \sum_\alpha v^\beta O_{\alpha\beta} \phi_\alpha$$

where summation on contracted indices is always implied. So the parameters u^α in Equation (4.28) are uniquely defined as $u^\alpha = v^\beta O_{\alpha\beta}$, and they are such that the hamiltonian vector field X_{H_T} is tangent to Γ .

2. or the matrix M – and hence Λ – is not invertible on Γ , thus it admits 0 as an eigenvalue and a number of corresponding eigenvectors. We let $2d \leq n - R_{\mathcal{L}}$ be the rank of Λ , and we assume that the first $2d$ constraints $\phi_{\alpha}^{(1)}$ have vanishing Poisson brackets with every other constraints except in the situation described in Equation (4.48). More precisely, for $1 \leq i \leq d$, Equation (4.50) becomes:

$$\{\phi_{2i-1}^{(1)}, H\} + v^{2i}\lambda_i \approx 0 \quad \text{and} \quad \{\phi_{2i}^{(1)}, H\} - v^{2i-1}\lambda_i \approx 0$$

Then, only the first $2d$ coefficients in Equation (4.50) are uniquely determined, while the last $n - R_{\mathcal{L}} - 2d$ are still unspecified. On the contrary, when $2d + 1 < \alpha \leq n - R_{\mathcal{L}}$, Equation (4.50) becomes:

$$\{\phi_{\alpha}^{(1)}, H\} \approx 0$$

If this equation is independent of the primary constraints, then it defines a new *second-stage* constraint¹⁹ $\phi_{\alpha}^{(2)} = \{\phi_{\alpha}^{(1)}, H\} \in \mathcal{C}^{\infty}(T^*Q)$.

The vanishing of this second-stage constraint $\phi_{\alpha}^{(2)}$ (for this particular α) on the smooth path $t \mapsto (q(t), p(t))$ is then interpreted as a necessary condition for the equations $\dot{\phi}_{\alpha}^{(1)} \approx 0$ to be satisfied. Notice that one cannot just replace the identities

$$\phi_{\alpha}^{(1)}(q(t), p(t)) = 0 \quad \text{for every } t \tag{4.52}$$

by the vanishing of the second-stage constraints because we originally used Equations (4.52) to define the second-stage constraints. Geometrically, it would be equivalent to dropping the condition that X_{H_T} is tangent to Γ . Rather, we deduce that physical solutions of Hamilton equations (4.39) should in fact be contained in the intersection of the constraint surface Γ and of the zero level set of all the second-stage constraints $\phi_{\alpha}^{(2)}$ (for every α for which they exist). Thus, let us define $\Gamma^{(2)}$ to be the submanifold of T^*Q corresponding to the zero level set of the primary (equivalently, first-stage) and second-stage constraints:

$$\Gamma^{(2)} = \Gamma \cap \bigcap_{\alpha} (\phi_{\alpha}^{(2)})^{-1}(0)$$

This definition is consistent because if $\phi_{\alpha}^{(2)} = 0$ – i.e. there is no second-stage constraint associated to $\phi_{\alpha}^{(1)}$ – then $(\phi_{\alpha}^{(2)})^{-1}(0) = T^*Q$. We can then use the same index α to label the second class constraints, although we know that there are maximum $n - R_{\mathcal{L}} - 2d$ of them.

Since the persistence at all time of the primary constraints is conditioned to the persistence of the second-stage constraints (for all time), one then should necessarily have $\dot{\phi}_{\alpha}^{(2)}(q(t), p(t)) = 0$ for any physical path $t \mapsto (q(t), p(t))$ (in addition to the condition that $\phi_{\alpha}^{(1)}(q(t), p(t)) = 0$ for any such path). Thus, we deduce that any solution of the Hamilton equations should sit in $\Gamma^{(2)}$ for all time t . Equivalently, this means that the hamiltonian vector field X_{H_T} should be tangent to $\Gamma^{(2)}$. Then, assuming that $\Gamma^{(2)}$ is an embedded submanifold of T^*Q , we deduce that the latter condition is satisfied when:

$$\{\phi_{\alpha}^{(1)}, H_T\}|_{\Gamma^{(2)}} = 0 \quad \text{and} \quad \{\phi_{\alpha}^{(2)}, H_T\}|_{\Gamma^{(2)}} = 0 \tag{4.53}$$

By definition of the second-stage constraints, the first equation is satisfied only if the second holds. For brevity and clarity of the statement (and to stick to the usage), until the next step

¹⁹The terminology "second-stage", "third-stage", etc. is taken from [Gitman and Tyutin, 1990], not to be confused with the notion of *L-th stage reducibility* of gauge theories introduced in [Gomis et al., 1995] and presented in Section 5.4.

of the algorithm, we will drop the restriction symbol $|_{\Gamma^{(2)}}$ and rather extend the meaning of the weak equivalence sign \approx by interpreting it as defined relatively to the submanifold $\Gamma^{(2)}$ defined by *all* the constraints generated up to this point: all the primary (equivalently, first-stage) and second-stage constraints, and not only the primary ones.

Then Equations (4.53) translate as:

$$\{\phi_\alpha^{(1)}, H_T\} \approx 0 \quad \text{and} \quad \{\phi_\alpha^{(2)}, H_T\} \approx 0 \quad (4.54)$$

where the \approx sign is now evaluated with respect to $\Gamma^{(2)}$. In other words, Equations (4.54) are a condition for any physical path to sit in the constraints surface $\Gamma^{(2)}$. Expanding the total Hamiltonian in the second equation, we obtain:

$$\{\phi_\alpha^{(2)}, H\} + v^\beta \{\phi_\alpha^{(2)}, \phi_\beta^{(1)}\} \approx 0 \quad (4.55)$$

If these equations are not trivial (of the form $0 = 0$), they will either provide a new relationship between the undefined parameters v^β 's, or a set of relationships between dynamical variables that we interpret as *third-stage* constraints $\phi_\alpha^{(3)} = \{\phi_\alpha^{(2)}, H\}$. This analysis can be performed in details by looking at the rank of the rectangular matrix whose coefficients are $\{\phi_\gamma^{(2)}, \phi_\delta^{(1)}\}$. At the cost of redefining the first-stage and second-stage constraints, we can put the rectangular matrix in a convenient form where only the eigenvalues appear and then proceed as we did above for Λ . We impose the weak equivalence in Equation (4.55) relatively to $\Gamma^{(2)}$ because there may happen that the third-stage constraints could be redundant with the primary or second-stage constraints. Putting the latter – first and second stage constraints – to zero would then enforce the former to be automatically zero as well, and we could then avoid any redundancy.

We then define $\Gamma^{(3)}$ to be the surface defined by all the constraints found up to this point: primary (first-stage), second-stage, and third-stage constraints:

$$\Gamma^{(3)} = \Gamma^{(2)} \cap \bigcap_{\alpha} (\phi_\alpha^{(3)})^{-1}(0)$$

Persistence of the second-stage constraints requires the third-stage constraints to vanish over any physical path $t \mapsto (q(t), p(t))$. Any physical path satisfying Hamilton equations (4.39) should then belong to this third-stage constraint surface. Equivalently, the hamiltonian vector field X_{H_T} should be tangent to $\Gamma^{(3)}$. As for the second step, until the next step of the algorithm, the weak equivalence sign is now interpreted to be defined relatively to the submanifold $\Gamma^{(3)}$ defined by *all* the constraints generated up to this point: all the primary, second-stage *and* third-stage constraints. And then, the algorithm goes on with $\phi_\alpha^{(3)}$ whose time derivative should vanish on $\Gamma^{(3)}$ as a necessary condition for $\phi_\alpha^{(2)}$ and thus $\phi_\alpha^{(1)}$ to stay invariant through time. The vanishing of $\phi_\alpha^{(3)}$ translates as:

$$\{\phi_\alpha^{(3)}, H\} + v^\beta \{\phi_\alpha^{(3)}, \phi_\beta^{(1)}\} \approx 0$$

where the weak equivalence sign is now understood to be computed with respect to $\Gamma^{(3)}$. If these equations are not trivial, we may find four-stage constraints, and then fifth-stage constraints and so on, but the algorithm terminates because the dimension of the phase space T^*Q is finite. We end up, for each α , with a sequences of k -th stage constraints $\phi_\alpha^{(k)}$ (the $\phi_\alpha^{(k)}$ are considered to be smooth functions, at least on some local neighborhood W of a fixed point (\tilde{q}, \tilde{p}) on Γ), and the sequence terminates, for each α , at some integer $k_\alpha \geq 1$. In other words, $\phi_\alpha^{(k_\alpha)} \neq 0$ while $\phi_\alpha^{(k_\alpha+1)} = 0$ (as smooth functions defined over W or T^*Q). See section 3.4 in [Rothe and Rothe, 2010], pp. 98-107 of [Sudarshan and Mukunda, 1974] or sections 1.1.5-1.1.7 in [Henneaux and Teitelboim, 1992], and section 2.2 of [Gitman and Tyutin, 1990] for various explanations on the Bergmann-Dirac algorithm.

Example 4.34. We have seen in Example 4.29 that the Hamiltonian H_0 could be straightforwardly extended to the whole phase space as a function H . Since there is only one primary constraint $\phi = p_y$, the second term in Equation (4.45) vanishes and persistence of the primary constraint then reads:

$$\{p_y, H\} \approx 0 \quad (4.56)$$

The parameter u is thus left undetermined, and using Equation (4.29) one obtains that Equation (4.56) is equivalent to:

$$x - p_x \approx 0$$

This is a necessary condition so that $\dot{\phi} \approx 0$ for all times. Then it is promoted to a second-stage constraint $\phi^{(2)} = x - p_x$. Since $\{x - p_x, p_y\} = 0$, persistence of this constraint does not give rise to any new constraint, and the algorithm stops there.

Example 4.35. Let us proceed in the same way for Example 4.30. There was only one primary constraint $\phi = p_x$. Persistence of this constraint gives the following condition:

$$\{p_x, H\} \approx 0$$

where H is the straightforward extension of the function H_0 defined in Equation (4.31). This gives the following condition:

$$p_y e^x \approx 0$$

which in turn implies that we have a second-stage constraint $\phi^{(2)} = p_y$ (the dependence on x does not appear because $e^x \neq 0$). Persistence of this second stage constraint reads:

$$\{p_y, H\} \approx 0$$

which in turn implies that $p_z \approx 0$. This necessary condition for the persistence of $\phi^{(2)}$ – and then of ϕ altogether – gives rise to the following third-stage constraint $\phi^{(3)} = p_z$. Persistence of this function does not provide any new constraint so the algorithm stops here.

Example 4.36. On \mathbb{R}^3 with coordinates q^i, p_i , the following Lagrangian is a modification of the one presented in Equation (1.22) on page 5 of this [paper](#):

$$L = \frac{1}{2} (v_1^2 + v_2^2) - \frac{1}{2} v_3 (q_1^2 + q_2^2 - r^2)$$

We decided to change the variable q_3 into v_3 (and remove the mass m) to see what is changing at the Hamiltonian level. The Lagrangian is obviously singular because the velocity v_3 only appears linearly thus we expect constraints to show up. There is a primary constraint $\phi_1 = p_3 + \frac{1}{2} (q_1^2 + q_2^2 - r^2)$, so that the total Hamiltonian reads:

$$H_T = \frac{1}{2} (p_1^2 + p_2^2) + u \phi_1$$

where u is (yet) an arbitrary smooth parameter (originally played by the velocity v_3).

The persistence equation for ϕ_1 (also possibly denoted $\phi_1^{(1)}$) reads:

$$\{\phi_1, H_T\} \approx q_1 p_1 + q_2 p_2$$

The right-hand side is promoted to be a new – hence second-stage – constraint ϕ_2 or $\phi_1^{(2)}$. The persistence equation for ϕ_2 gives:

$$\{\phi_2, H_T\} \approx p_1^2 + p_2^2 - u (q_1^2 + q_2^2) \approx p_1^2 + p_2^2 - u r^2$$

We have indeed added and subtracted $u r^2$ to make the first constraint appear. Then we deduce that there is no third-stage constraint (contrary to the original example in Matschull's paper), but that the parameter u can be fixed to the following value: $u = \frac{p_1^2 + p_2^2}{r^2}$.

All k -th stage constraints are called *secondary constraints*, in order to emphasize that they come after imposing some condition on the primary constraints. The subset of T^*Q consisting (at least locally) of the zero level set of the set of all the constraints (primary *and* secondary) is called the (*secondary*) *constraint surface* and is denoted Σ :

$$\Sigma = \bigcap_{k \geq 1} \Gamma^{(k)}$$

This surface is independent on the choice of primary (and then secondary) constraints that is originally made (that was already implicit in the discussion following Remark 4.16). From now on, we will consider now that the weak equivalence is now defined with respect to the secondary constraint surface Σ , and not Γ anymore, so that we will use the notation \approx of Definition 4.31 to indicate equivalence of functions on the constraint surface Σ . As for the primary constraint surface $\Gamma^{(1)} = \Gamma$, the secondary constraint surface Σ is assumed to satisfy a regularity condition similar to that of Scholie 4.19.

Scholie 4.37. Regularity condition on secondary constraints. *Every point $(q, p) \in \Sigma$ admits a neighborhood U on which there exists $\text{codim}(\Sigma)$ constraints which play the role of local transverse coordinates. In particular, the secondary constraint surface Σ is an embedded submanifold of T^*Q .*

Let us explain the importance of the regularity condition on secondary constraints. The regularity condition for the primary constraint surface – Scholie 4.19 – was a statement about Γ , and not about the primary constraints, which by construction were natural local transverse coordinates of the image of the Legendre transform. In Scholie 4.18, we merely asked that this image did not self intersect itself, leading to the fact that the primary constraint surface was an embedded submanifold. The primary constraints were by construction then automatically playing the role of local transverse coordinates. Scholie 4.19 can be interpreted also by saying that the primary constraints form a *regular sequence*²⁰. However, it may well happen that secondary constraints cannot be used as transverse coordinates to Σ because they emerged from the equations of motions and not from the rank theorem as primary constraints did. The regularity condition for the secondary constraint surface however solves this issue by specifically requiring that the constraints play that role. For example, it allows us to extend the validity of Proposition 4.24 to Σ .

If this regularity condition was not satisfied, but the secondary surface Σ still assumed to be embedded, by Proposition 2.43, we would deduce that locally it is a closed embedded submanifold which is a level set of a submersion from some open set U to $\mathbb{R}^{\text{codim}(\Sigma)}$. Then, the components of this submersion play the role of transverse coordinates to Σ and thus locally generate $\mathcal{I}_\Sigma \cap \mathcal{C}^\infty(U)$ (and the constraints). Then, we would always be able to find a local replacement of the constraints that locally generate \mathcal{I}_Σ but such functions only have a mathematical meaning and do not have the same physical meaning as the constraints since they did not emerge from the formalism. The role of Scholie 4.37 is precisely to have the best of both worlds: to keep the constraints for their physical importance, but also use them for their mathematical relevance (as local generators of the ideal \mathcal{I}_Σ on the open set U , say). If the regularity condition is not satisfied directly, it is possible to modify some (secondary) constraints without changing Σ so that we obtain a new set of constraints that satisfy Scholie 4.37.

Example 4.38. Take $Q = \mathbb{R}^n$ and assume that one of the secondary constraint is $\phi = (p_1)^2$. Then the secondary constraint surface is included in the plane of equation $p_1 = 0$, but the very

²⁰At the cost of slightly changing the Bergmann-Dirac algorithm – by not computing the constraints stage by stage, each stage at a time, but rather by implementing every new secondary constraint directly in the algorithm by requiring that it is weakly zero – one can ensure that this new set of secondary constraints (possibly smaller than the one defined in the present section) forms a regular sequence.

form of the constraint prevents to use it as a transverse coordinate because it is never negative. Thus, a proper secondary satisfying the regularity condition for Σ would be $\phi' = p_1$, which is of course physically equivalent to ϕ , but mathematically quite different (as it is only linear and generates ϕ).

The Bergmann-Dirac algorithm tells us that any physical path – i.e. a solution of Hamilton equations (4.39) – should then be sitting in Σ (thus modifying the original statement of Lemma 4.32). This is a consequence of the fact that secondary constraints, which are hidden in the persistence equations of the primary constraints, are actually needed to draw an equivalence with Euler-Lagrange equations:

Proposition 4.39. *The Euler-Lagrange equations (4.6) are equivalent to the following set of Hamilton equations:*

$$\begin{cases} \dot{q}^i &= \{q^i, H_T\} \\ \dot{p}_i &= \{p_i, H_T\} \\ \dot{\phi}_\alpha^{(k)} &= 0 \quad \text{for all } k \geq 1 \end{cases}$$

where we choose the notation $\{\phi_\alpha^{(1)}, \phi_\alpha^{(2)}, \phi_\alpha^{(3)}, \dots, \phi_\alpha^{(k)}, \dots\}_{\alpha,k}$ to denote the set of k -th stage constraints.

Proof. See page 29 of [Rothe and Rothe, 2010] and the subsequent section. \square

The Bergmann-Dirac algorithm then amounts to finding the parameters u^α (equivalently, v^α) in the total hamiltonian H_T so that the hamiltonian vector field X_{H_T} is tangent to the primary constraint surface $\Gamma = \Gamma^{(1)}$. This requirement leads to a chain of new conditions: that the secondary constraints hold at all time, i.e. that the persistence equations $\dot{\phi}_\alpha^{(k)} \approx 0$ hold for every constraint. Then, assuming the regularity condition for Σ , Lemma 2.58 tells us that this set of conditions can be summarized in a simple geometrical statement:

the hamiltonian vector field X_{H_T} should be tangent to the secondary constraint surface Σ

The solutions of Hamilton's equations forming the integral curves of $-X_{H_T}$, we deduce that the physical paths are then necessarily constrained to Σ , and the image through the Legendre transform of a solution of the Euler-Lagrange equations consequently also sits in Σ .

Notice that the difference between primary and secondary constraints is not so clear because the k -th stage secondary constraint $\phi_\alpha^{(k)}$ (when it exists) often involves primary constraints in its expression, which then vanish when evaluated over Γ . Even the choice of primary constraint is not unique, since the choice of a minor in the Hessian matrix of the Lagrangian determines the primary constraints, and we have later even performed a 'rotation' in the space of primary constraints by Equation (4.47). Eventually, a given system of Lagrangian equations of motion can sometimes be derived from more than one Lagrangian, but which constraints are primary and which are secondary depends on the functional form of the Lagrangian. Then, any other choice of first-stage constraints $\phi_\alpha'^{(1)}$ would give rise to secondary constraints $\phi_\alpha'^{(k)}$, so that the new constraints could be obtained from the set of original constraints through a linear transformation:

$$\phi_\alpha'^{(k)} \approx \sum_{l \geq 1} C_{(l)\alpha}^\beta \phi_\beta^{(l)} \quad (4.57)$$

where summation over repeated indices is implicit and the $(C_{(l)})_{l \geq 1}$ is a family of square matrices whose coefficients are smooth functions. Although secondary constraints are often undistinguishable from primary constraints, some author value primary constraints as carrying noticeable

information: see e.g. page 39 and page 72 of [Gitman and Tyutin, 1990], page 10 of [Henneaux and Teitelboim, 1992], page 148 of [Earman, 2003] or subsection 3.3.2 in [Rothe and Rothe, 2010]. In the latter reference, it is postulated that the parameters u^α appearing in the total Hamiltonian may be considered as the projections of the velocities on the zero eigenspace of the Hessian of the Lagrangian. The primary constraints in this context simply state that these velocities stay finite.

Example 4.40. This example is taken from [Rothe and Rothe, 2010]. In order to study the meaning of primary constraints, let us modify Example 4.14 so that its Lagrangian is obtained as a limit $\alpha \rightarrow 0$ of the following (non-singular) Lagrangian:

$$L(q, v) = \frac{1}{2}v_x^2 + v_x y + \frac{\alpha}{2}v_y^2 + \frac{1}{2}(x - y)^2$$

As much as $\alpha \neq 0$ the Hessian of L is non-singular:

$$\mathcal{H}(q, v) = \begin{pmatrix} 1 & 0 \\ 0 & \alpha \end{pmatrix}$$

The Legendre transform is then bijective and given by:

$$\mathcal{L}_x(q, v) = \mathcal{L}(q, v)(\partial_x) = v_x + y \quad \text{and} \quad \mathcal{L}_y(q, v) = \mathcal{L}(q, v)(\partial_y) = \alpha v_y$$

In particular, the relationship between velocities and momenta is given by $v_x = p_x - y$ and $v_y = \frac{p_y}{\alpha}$. Thus, evaluating the canonical Hamiltonian H_c on $N_{\mathcal{L}}$ as in Example 4.29, gives H_0 :

$$\begin{aligned} H_0(q, p; \alpha) &= p_x(p_x - y) + p_y \frac{p_y}{\alpha} - \left(\frac{1}{2}(p_x - y)^2 + (p_x - y)y + \frac{1}{2\alpha}p_y^2 + \frac{1}{2}(x - y)^2 \right) \\ &= \frac{1}{2}p_x^2 + \frac{1}{2\alpha}p_y^2 - \frac{1}{2}x^2 + xy - yp_x \end{aligned} \quad (4.58)$$

Let us rewrite the second term $\frac{1}{2\alpha}p_y^2$ as $\frac{1}{2\alpha}\psi^2$ where $\psi = \alpha v_y$, because this is how the y -velocity and the y -momentum are related to one another via the Legendre transform. This rewriting emphasizes that, although the denominator makes the fraction $\frac{1}{2\alpha}$ diverge when $\alpha \rightarrow 0$, at the same time the numerator ψ^2 will converge to 0 twice quicker, making the overall term to vanish. Then, we have $H_0(q, p; \alpha) \xrightarrow{\alpha \rightarrow 0} H_0(q, p)$, where $H_0(q, p)$ is the Hamiltonian defined in Equation (4.29). One can recast Equation (4.58) as:

$$H_0(q, p; \alpha) = H_0(q, p) + \frac{1}{2}v_y p_y \quad (4.59)$$

Under this form, one is reminded of the general form of the total Hamiltonian of Equation (4.30). Comparing the latter equation with Equation (4.59) one realizes that the parameter u associated to the primary constraint in Equation (4.30) is related to the velocities (they may be interpreted as coordinates on the preimage of p_x , as is explained on page 10 of [Henneaux and Teitelboim, 1992]). The example of the free electromagnetic field in Section B.4 also shows that the coefficients associated with the primary constraints are related to velocities.

Remark 4.41. Notice that there exists a similar algorithm for the Lagrangian picture, see Section 2 of [Rothe and Rothe, 2010]. The counterpart of the constraint surface in the Lagrangian picture is an embedded submanifold of the tangent bundle TQ , on which the Euler-Lagrange equations are deterministically solvable. It has been shown that this perspective is equivalent to that obtained through the Bergmann-Dirac algorithm [Batlle et al., 1986, Rothe and Rothe, 2003]. As a concluding remark, on the mathematical side, Gotay and collaborators [Gotay et al., 1978] have generalized the Bergmann-Dirac algorithm to the more general case where the phase space is not a cotangent bundle but a presymplectic manifold.

4.4 First-class and second-class constraints, gauge transformations

We have seen that the true object of physical interest may not be the constraints themselves but the surface induced by the constraints, which can actually be generated by many different functions as their zero level set. In particular, a distinguished choice of such functions would be those generating the multiplicative ideal of all functions vanishing on the secondary constraint surface $\mathcal{I}_\Sigma \subset \mathcal{C}^\infty(T^*Q)$. These functions, if they do not coincide with the constraints, would however generate them, as the latter vanish on Σ . Let us provide a more rigorous statement to this claim.

Locally, the choice of primary constraints made in Equation (4.25) is such that they are ‘trivially’ functionally independent on the open subset W , in the following sense:

$$\forall f^\alpha \in \mathcal{C}^\infty(W) \text{ such that } \sum_\alpha f^\alpha \phi_\alpha = 0 \quad \implies \quad f^\alpha = \sum_\beta \sigma^{\alpha\beta} \phi_\beta \text{ with } \sigma^{\alpha\beta} = -\sigma^{\beta\alpha}$$

Thus, the primary constraints chosen in such a way are not exactly functionally independent on the open set W , but the dependence functions are minimal – or ‘trivial’ – in the sense that they intervene only because of their antisymmetry property. However, when performing the Bergmann-Dirac algorithm, the k -th stage constraints have been defined at the condition that all other lower stage constraints only be kept to zero. It might happen that some of the k -th stage secondary constraints be functionally dependent, via some functions f^α which are not trivial in the above sense. We thus end up with a total set of constraints which may or may not be functionally dependent. This justifies the following terminology:

Definition 4.42. *We say that the (primary and secondary) constraints ϕ_α are irreducible if the only functional dependence between them is minimal in the following sense:*

$$\forall f^\alpha \in \mathcal{C}^\infty(W) \text{ such that } \sum_\alpha f^\alpha \phi_\alpha = 0 \quad \implies \quad f^\alpha = \sum_\beta \sigma^{\alpha\beta} \phi_\beta \text{ with } \sigma^{\alpha\beta} = -\sigma^{\beta\alpha}$$

*We say that they are reducible if their functional dependence is non trivial, i.e. if there exists a set of smooth functions $\{Z_I^\alpha\}_{\alpha,I}$ on T^*Q which do not vanish everywhere on Σ and such that we have, for every I :*

$$Z_I^\alpha \phi_\alpha = 0 \tag{4.60}$$

Irreducible constraints are virtually functionally independent; the functional dependence of irreducible constraints is thus sometimes said to be ‘trivial’, in accordance with the denomination of *trivial gauge transformations* (see subsection 3.1.5 in [Henneaux and Teitelboim, 1992]). Notice that the functions f^α in the former part of the definition vanish on the secondary constraint surface Σ , while the *reducibility functions* Z_I^α precisely do not (everywhere). Obviously, the choice of such functions is not unique as one can always add to Z_I^α a contribution $\sigma^{\alpha\beta} \phi_\beta$ which does not change Equation (4.60). The functions Z_I^α might additionally be functionally dependent, opening the possibility of having higher reducibility functions (see Section 5.4).

The notion of (ir)reducibility will have large consequences in the geometrical treatment of constrained Hamiltonian systems. As shown in Section 4.2, one can always find an irreducible set of primary constraints. The role of the regularity condition for the secondary constraint surface is to assume that this is true for secondary constraints as well. Indeed, Scholie 4.37 establishes that the secondary constraint surface Σ is an embedded submanifold, and as such, there exists locally a minimal number of generators of the ideal \mathcal{I}_Σ of functions vanishing on Σ , to which the primary and secondary constraints belong. Moreover, one can always find these generators within the set of constraints. This is straightforward if the constraints are irreducible as they have no non-trivial functional dependency so they generate \mathcal{I}_Σ . For reducible

constraints, the regularity condition establishes that, locally on an open set W of a point of Σ , there exists a subset of irreducible constraints which generate all the others (this is the first regularity condition appearing in subsection 1.1.2 of [Henneaux and Teitelboim, 1992]). On another open neighborhood W' of a point of Σ , this subset might change, hence the impossibility in the reducible case to pick up a consistent, globally defined, subset of irreducible constraints generating all the others on T^*Q . While it is always possible to locally work with an irreducible set of constraints, keeping a reducible system of constraints might be preferable not only because it enables to work globally, in a coordinate free description of the problem, but also because it may not be so easy to reduce the full set of reducible constraints to an irreducible subset without spoiling amenable physical properties such as Lorentz covariance, locality, unitarity, etc.

This discussion shows that the set Ω_0 of (primary and secondary) constraints defined in Section 4.3 is geometrically not particularly relevant, and that we can find another, alternative set of constraints which is in some sense equivalent to the former, at the condition that the regularity condition for the latter is still satisfied.

Definition 4.43. *A set Ω of functions vanishing on the secondary constraint surface Σ is said to be equivalent to the set of primary and secondary constraints Ω_0 and we note:*

$$\Omega \sim \Omega_0$$

if elements of Ω can be used as local transverse coordinates as in Scholie 4.37. We will also call constraints the elements of such a set.

Remark 4.44. An alternative, equivalent definition (which always works locally) is as follows. Assume that the primary and secondary constraints ϕ_α are irreducible. Denote by $\tilde{\phi}_\beta$ the elements of Ω and suppose that they are irreducible as well. Then Ω is equivalent to Ω_0 if and only if there is a square matrix O , invertible on Σ , such that $\tilde{\phi}_\beta = O_\beta^\alpha \phi_\alpha$, as in Equation (4.47).

We see that the present notion of equivalence is an equivalence relation between sets of smooth functions. All such sets generate the ideal \mathcal{I}_Σ of vanishing functions on Σ . The elements of an equivalent set of constraints Ω are obtained by linear combinations of primary and secondary constraints, while the converse is true as well (see top of page 25 in [Gitman and Tyutin, 1990] for a small discussion about this topic). This notion offers some flexibility in the way of treating hamiltonian system. We observed at the end of Section 4.3 that the splitting into primary and secondary constraints is neither mathematically nor really physically relevant. A better distinction is that of *first-class* and *second-class* constraints, originally proposed by Dirac and which has deep relationship with gauge transformations and the *Dirac conjecture*. This section is dedicated to study these kind of constraints and we will first introduce the following central notion:

Definition 4.45. *Let $\Omega = \{\phi_i\}$ be a set of constraints equivalent to that of primary and secondary constraints Ω_0 . A smooth function $f \in \mathcal{C}^\infty(T^*Q)$ is said to be first-class (relatively to Ω) if its Poisson bracket with every constraint $\phi_i \in \Omega$ vanishes everywhere on Σ , i.e. we have:*

$$\{f, \phi_i\} \approx 0 \tag{4.61}$$

for every constraint ϕ_i . The function f is said to be second-class (relatively to Ω) if, for every $x \in \Sigma$, there exists a constraint $\phi_i \in \Omega$ such that:

$$\{f, \phi_i\}(x) \neq 0$$

Example 4.46. If one considers that the weak equivalence is now defined with respect to the secondary constraint surface – as it should be, the total Hamiltonian H_T is a first class function, with respect to the set of primary and secondary constraints Ω_0 . Indeed, by construction the persistence equation $\{\phi_\alpha^{(k)}, H_T\} \approx 0$ is satisfied for both primary and secondary constraints.

While we can easily understand the notion of first-class functions, we need to discuss the latter notion of second-class functions. Physicists often define the former as the negation of the latter by saying that second-class functions are those which are *not first-class* (even in Dirac's lectures, top of page 18 in [Dirac, 1964]). This poses a problem of interpretation because mathematically, the converse of Equation (4.61) would be the following statement: there exists a point $x \in \Sigma$ and a constraint $\phi_i \in \Omega$ such that $\{f, \phi_i\}(x) \neq 0$, but nothing is said about elsewhere. In particular the bracket of f with any constraint could possibly vanish entirely on Σ sufficiently far away from x . To avoid such misunderstanding, physicists sometimes precise that f is second-class if there exists a constraint ϕ_i such that the Poisson bracket $\{f, \phi_i\}$ *never* vanishes on Σ (see e.g. subsection 1.1.10 in [Henneaux and Teitelboim, 1992]). Notice that not being a first-class constraint does not imply this property. In other words, in this sense, the non-triviality condition $\{f, \phi_i\} \neq 0$ is satisfied globally over Σ , so f is 'maximally non first-class'. This is still not the same assumption as in Definition 4.45, and we will see in Example 4.52 that this assumption is sometimes not satisfied. The correct definition is that of Definition 4.45, that we will reformulate later when we apply the notion of first-class and second-class to constraints themselves (see Definition 4.49 and Proposition 4.51).

Although we defined first-class and second-class functions relatively to a set of constraints, their dependence on this particular set of constraints is artifactual: we did it to facilitate further discussions on the topic. We can then state the following result:

Proposition 4.47. *If a smooth function f is first-class (resp. second-class) relatively to a set of constraints, it is first-class (resp. second-class) with respect to any other set of constraint which is equivalent to the former.*

Proof. Let Ω and Ω' be two equivalent sets of constraints, then it means that the latter is obtained using linear combinations of elements of the former, e.g. $\phi'_i = \lambda_i^j \phi_j$ where $\lambda = (\lambda_i^j)$ is an invertible matrix at each point. Then, by using the Leibniz property of the Poisson bracket applied to $\lambda_i^j \phi_j$, if f is first-class relatively to Ω , it will also be first-class relatively to Ω' . The same argument is used in order to show that the statement holds as well for second-class functions: suppose that f is second-class with respect to the set Ω , but that it is not anymore relatively to the set of constraints Ω' . That is to say: there exists a point $x \in \Sigma$ such that $\{f, \phi'_i\}(x) = 0$ for every constraint $\phi'_i \in \Omega'$. Since Ω and Ω' are equivalent, every constraint of Ω can be written as a linear combination of the constraints ϕ'_i . But then by the same Leibniz property as before, it means that, at x , we have that $\lambda_i^j(x) \{f, \phi_j\}(x) = 0$ for every i . By inverting the matrix λ we deduce that $\{f, \phi_j\}(x) = 0$ for every constraint $\phi_j \in \Omega$, but this is a contradiction with the assumption that f is a second-class function. \square

By Proposition 4.47, we do not need to specify with respect to which set of constraints a first-class (resp. second-class) function is first-class (resp. second-class). This fact has the following nice consequence regarding the bracket of two first-class functions:

Proposition 4.48. *The Poisson bracket of two first-class functions is first class, so they form a Lie subalgebra of $(\mathcal{C}^\infty(T^*Q), \{.,.\})$.*

Proof. Let $\Omega = \{\phi_i\}$ be a set of constraint equivalent to the set Ω_0 of primary and secondary constraints. Let $f, g \in \mathcal{C}^\infty(T^*Q)$ be first-class functions relatively to Ω . Although f and g need not be vanishing on the constraint surface Σ , their Poisson brackets with any constraint ϕ_i is vanishing. Since the regularity condition for the secondary constraint surface Σ is satisfied by the elements of Ω , Proposition 4.24 applied to the constraints ϕ_i implies that these Poisson brackets are strongly equivalent to a linear combination of the constraints:

$$\{f, \phi_i\} = \sum_j F_{ij} \phi_j \quad \text{and} \quad \{g, \phi_i\} = \sum_j G_{ij} \phi_j$$

To evaluate if the Poisson bracket of f and g is first-class we compute:

$$\begin{aligned}\{\{f, g\}, \phi_i\} &= \{\{f, \phi_i\}, g\} + \{f, \{g, \phi_i\}\} \\ &= \{F_{ij}\phi_j, g\} + \{f, G_{ik}\phi_k\} \\ &= F_{ij}\{\phi_j, g\} + \{F_{ij}, g\}\phi_j + G_{ik}\{f, \phi_k\} + \{f, G_{ik}\}\phi_k\end{aligned}$$

which identically vanish on Σ because on the one hand $\phi_j, \phi_k \approx 0$ and on the other hand the functions f and g are first-class. \square

Definition 4.45 applies straightforwardly to constraints themselves, which are smooth functions as any other:

Definition 4.49. *A smooth function ϕ is said to be a first-class (resp. second-class) constraint if it vanishes on Σ and is a first-class (resp. second-class) function relatively to the set of primary and secondary constraints Ω_0 .*

Example 4.50. The primary and secondary constraints found in Example 4.34 and 4.35 are all first-class, with respect to the set of primary and secondary constraints. The primary and secondary constraints found in Example 4.36 are second-class because their Poisson bracket is equal to $q_1^2 + q_2^2$, equal to the positive constant r^2 on the constraint surface, i.e. non vanishing.

The set of primary and secondary constraints Ω_0 is not necessarily adapted to make the splitting into first-class and second-class constraints. That is to say: it is possible that no such constraint has a vanishing bracket with every other constraint (see Example 4.52). However, at the cost of redefining the constraints via functional linear combinations, it is possible to define a new set of constraints Ω which is equivalent Ω_0 , and which splits into two subsets of first-class and second-class constraints. By construction, this set of constraints satisfies the property that no linear combination of second-class constraints is first class.

In order to find an adapted set of constraints $\Omega = \{\phi_i\}$ which allows the splitting into first-class and second-class constraints, there is a straightforward technique. Let D be the matrix made of the Poisson brackets of the primary and secondary constraints: $D_{kl, \alpha\beta} = \{\phi_\alpha^{(k)}, \phi_\beta^{(l)}\}$. We assume from now on that the rank of this matrix is (at least locally) constant on the constraint surface Σ , and that its rank is $2r$. The rank is necessary an even integer because D is an antisymmetric matrix. By Proposition 4.33, the matrix D is diagonalizable by block and on the constraint surface Σ , it is semblable to the following matrix:

$$\Delta = \begin{pmatrix} 0 & \delta_1 & & & & & & & & & & \\ -\delta_1 & & & & & & & & & & & \\ & & 0 & \delta_2 & & & & & & & & \\ & & -\delta_2 & 0 & & & & & & & & \\ & & & & \ddots & & & & & & & \\ & & & & & & 0 & \delta_r & & & & \\ & & & & & & -\delta_r & 0 & & & & \\ & & & & & & & & 0 & & & \\ & & & & & & & & & \ddots & & \\ & & & & & & & & & & 0 & \\ & & & & & & & & & & & 0 \end{pmatrix}$$

such that the δ_i 's are strictly positive numbers. The columns of the orthogonal matrix diagonalizing the matrix D give us the coefficients appearing in Equation (4.57) in order to define the new constraints ϕ_i . They are divided in two sets: the first-class and the second-class constraints.

Usually first-class constraints are denoted φ_j and second class constraints are denoted χ_l . On the matrix Δ , we see that the number $2r$ corresponds to the number of independent second-class constraints. The latter can then be seen as labelled by indices from 1 to $2r$ and are such that $\{\chi_{2i-1}, \chi_{2i}\} = \delta_i$ for any $1 \leq i \leq r$, while any other bracket with χ_{2i-1} or χ_{2i} vanishes. If the primary and secondary constraints were irreducible, then the dimension of D is $n - R_{\mathcal{L}}$ and the number of first-class constraints φ_i is $n - R_{\mathcal{L}} - 2r$. Their Poisson bracket with every other constraints χ_l and φ_i is materialized by the last $n - R_{\mathcal{L}} - 2r$ lines (or columns) of the matrix Δ . These functions vanish on Σ and are irreducible if the original primary and secondary constraints were irreducible as well. The set $\Omega = \{\chi_1, \dots, \chi_{2r}, \varphi_1, \dots, \varphi_{n-R_{\mathcal{L}}-2r}\}$ is then the desired set of constraints split into first-class and second-class constraints. From this discussion we deduce the following property:

Proposition 4.51. *A smooth function χ is a second-class constraint if and only if there exists a set of constraints Ω equivalent to Ω_0 such that χ belongs to a subset of constraints $\{\chi_l\} \subset \Omega$ which has the property that the determinant of the matrix $\{\chi_l, \chi_k\}$ is nowhere vanishing on Σ .*

Proof. One direction is straightforward: if such a constraint χ satisfies the above condition then it is automatically second-class by Proposition 4.47. Conversely, assume that a function χ vanishing on Σ is second class (relatively to Ω_0). Assume that we had proceeded to the block diagonalization of the matrix D to obtain the matrix Δ as in the above discussion. This provides us with a set of constraints $\Omega = \{\chi_1, \dots, \chi_{2r}, \varphi_1, \dots, \varphi_{n-R_{\mathcal{L}}-2r}\}$, equivalent to Ω_0 , and which splits into first-class and second-class constraints. The number of second-class constraints is minimal and they generate all possible second-class constraint, in particular χ . Then, under the assumption that χ is regular in the vicinity of Σ , i.e. that it can be used as a local transverse coordinate, then one can always find a second-class constraint from Ω and replace it by χ , while keeping the set of second-class constraints independent. In that case, the determinant of the matrix of the Poisson brackets of the second class constraints would still be non vanishing on Σ . Indeed, if the determinant of this matrix would drop under the replacement of one of the second-class constraints of Ω by χ , then it would mean that there are less second-class constraints than expected, which is not possible. The assumption of Proposition 4.51 is then satisfied. \square

The condition stated in Proposition 4.51 is the definition appearing at the beginning of Section 2.3 of [Gitman and Tyutin, 1990]. This definition is often presented in physics textbooks as a consequence of the fact that second-class constraints are precisely those constraints which are *not* first-class. However, as we have seen when diagonalizing the matrix D , it is not at all certain that a given set of constraints Ω_0 coincides with the eigenvectors of Δ . This is illustrated in Example 4.52, which shows that the original set of constraints Ω_0 may not even possess first-class constraints. The latter appear only one performs an invertible change in the set of constraints (using linear combinations) so that the resulting set Ω can be split into first-class and second-class constraints (in the sense of Definition 4.49). Then, yes, a constraint which is not first-class is second-class, but this definition would not hold regarding the original set of constraints. This is why the physicists' definition of second-class constraints is often too imprecise to make mathematical sense.

Example 4.52. Let $Q = \mathbb{R}^2$ so that T^*Q admits as coordinate functions q_1, q_2, p_1, p_2 . Let $\phi_1 = q_1 - q_2$, $\phi_2 = \frac{q_1^2 + q_2^2 + p_1^2 + p_2^2}{2} - 1$ and $\phi_3 = p_1 + p_2$ three smooth functions that we consider to be the constraints and we suppose that there are no other. The constraint surface Σ is the circle of radius 2 sitting in the plane defined by the equations:

$$q_1 = q_2 \quad \text{and} \quad p_1 = -p_2 \tag{4.62}$$

Let $\lambda = \sqrt{(q_1 + q_2)^2 + (p_1 - p_2)^2} \geq 0$; although this continuous function is not smooth everywhere on the phase space, it is at least smooth in a tubular neighborhood of the constraint

surface Σ since, by Equations (4.62) and the identity $\phi_2 = 0$, on Σ we have that:

$$\lambda^2 = q_1^2 + q_2^2 + p_1^2 + p_2^2 + \underbrace{q_1 q_2}_{q_1^2} + \underbrace{q_1 q_2}_{q_2^2} - \underbrace{p_1 p_2}_{-p_1^2} - \underbrace{p_1 p_2}_{-p_2^2} = 4 \quad (4.63)$$

Let us compute the matrix D obtained by computing the Poisson brackets of the constraints with themselves:

$$D = \begin{pmatrix} 0 & p_1 - p_2 & 0 \\ -p_1 + p_2 & 0 & q_1 + q_2 \\ 0 & -q_1 - q_2 & 0 \end{pmatrix}$$

The matrix D is skew symmetric so has even rank; it cannot be zero on the constraint surface Σ for then $q_1 = q_2 = p_1 = p_2 = 0$ by Equations (4.62), which cannot be allowed by the identity $\phi_2 = 0$. Then it has rank 2 on Σ so we expect to find two second class constraints χ_1, χ_2 and one first-class constraint φ . Notice that on the constraint surface Σ , we can already say that ϕ_2 is a second-class constraint (we cannot have at the same time both $p_1 = p_2$ and $q_1 = -q_2$), but none of the constraints ϕ_1, ϕ_3 is first-class or second-class. We then have to make a transformation of the set of constraints $\{\phi_1, \phi_2, \phi_3\}$ to obtain the set $\{\chi_1, \chi_2, \varphi\}$ which makes explicit the splitting of constraints into first-class and second-class constraints.

The matrix D admits the following eigenvectors and null vector:

$$u = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad v = \frac{1}{\lambda} \begin{pmatrix} -(p_1 - p_2) \\ 0 \\ q_1 + q_2 \end{pmatrix} \quad \text{and} \quad w = \frac{1}{\lambda} \begin{pmatrix} q_1 + q_2 \\ 0 \\ p_1 - p_2 \end{pmatrix}$$

in the sense that:

$$Du = -\lambda v, \quad Dv = \lambda u \quad \text{and} \quad Dw = 0$$

We can block-diagonalize D , using the orthogonal matrix $O = (u \ v \ w)$. One can indeed check that $O^T O = I_3$, and that $\Delta = O^T D O$ reads:

$$\Delta = \begin{pmatrix} 0 & \lambda & 0 \\ -\lambda & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

From Equation (4.47), we read on the columns of O (equivalently, on the lines of O^T) how to define the eigenvectors of Δ from the three constraints ϕ_1, ϕ_2, ϕ_3 . So we set:

$$\begin{aligned} \chi_1 &= \phi_2 \\ \chi_2 &= \frac{1}{\lambda} (-(p_1 - p_2)\phi_1 + (q_1 + q_2)\phi_3) \\ \varphi &= \frac{1}{\lambda} ((q_1 + q_2)\phi_1 + (p_1 - p_2)\phi_3) \end{aligned}$$

Notice that by performing this transformation we did not lose any geometric information: both the initial constraints ϕ_1, ϕ_2, ϕ_3 and the new constraints χ_1, χ_2, φ can be used as local transverse coordinates on Σ and thus generate the ideal \mathcal{I}_Σ of functions vanishing on Σ (at least in a tubular neighborhood of Σ , since λ is not smooth everywhere). Since $O^T O = I_3$, we can indeed generate the initial constraints from the first-class and second-class constraints (in a tubular neighborhood of Σ , where λ is not zero):

$$\begin{aligned} \phi_1 &= \frac{1}{\lambda} (-(p_1 - p_2)\chi_2 + (q_1 + q_2)\varphi) \\ \phi_2 &= \chi_1 \\ \phi_3 &= \frac{1}{\lambda} ((q_1 + q_2)\chi_2 + (p_1 - p_2)\varphi) \end{aligned}$$

Then, the two sets of constraints are equivalent in the sense of Definition 4.43.

The role of λ is the definition of the constraints χ_1, χ_2, φ might seem superfluous because it is constant on Σ and will play no role whatsoever, but then Δ would not have been the matrix of the Poisson bracket of these constraints anymore. Indeed, one can compute that:

$$\{\varphi, \chi_1\} \approx 0, \quad \{\varphi, \chi_2\} \approx 0, \quad \text{and} \quad \{\chi_1, \chi_2\} \approx \frac{2}{\lambda}(\lambda^2 - 2) = \lambda$$

The last equality is obtained from Equation (4.63), which shows that $\lambda = 2$ on the constraint surface Σ . Thus not only we have that Δ is the matrix of the Poisson brackets of the new set of constraints $\{\chi_1, \chi_2, \varphi\}$, but we also see that φ is first-class with respect to the set of constraints $\{\chi_1, \chi_2, \varphi\}$, and since $\lambda \neq 0$ on Σ , the two constraints χ_1, χ_2 are second-class relatively to the same set. We already knew that $\chi_1 = \phi_2$ was second-class and this confirmed here, because we found another constraint χ_2 such that the matrix $\{\chi_i, \chi_j\}$ has non vanishing determinant on Σ . Notice as well that in the original set of constraints neither ϕ_1 nor ϕ_3 were first-class or second-class. This is why we had to perform a transformation to obtain an equivalent set which allows such splitting into first-class and second-class constraints. Finally, since the initial constraints can be generated by the new ones, one can check that the constraint φ is not only first-class with respect to χ_1, χ_2, φ but also with respect to ϕ_1, ϕ_2, ϕ_3 , proving once again Proposition 4.47.

Since the matrix used to block diagonalize D and obtain the set Ω split into first-class and second-class constraints is orthogonal, its transpose is its own inverse so we can obtain the original primary and secondary constraints back from Ω by a formula similar to Equation (4.57). This has the consequence that the set of constraints Ω generates the same ideal of functions \mathcal{I}_Σ as the primary and secondary constraints. If the latter were irreducible, then the former will be as well. In particular, by Lemma 2.58, it means that the Hamiltonian vector fields of every first-class constraints $X_{\varphi_i} = \{\varphi_i, \cdot\}$ are tangent to the secondary constraint surface Σ , while the hamiltonian vector field of secondary constraints are nowhere tangent to Σ (see Proposition 4.64). This was not necessarily the case with the Hamiltonian vector fields of the primary and secondary constraint alone. Moreover, the $\mathcal{C}^\infty(T^*Q)$ -module of first-class constraints is the intersection of the ideal \mathcal{I}_Σ and the Lie algebra of first-class functions (as defined in Proposition 4.48). Then we see that by passing from the original set of constraints to the new one, we do not lose any mathematical information, but we gain some things which are quite valuable – and we will see that this has some physical relevance.

Notice also that there is some latitude in the choice of first-class and second-class constraint, since e.g. one cannot make the difference between $\{\chi_l, \cdot\}$ and $\{\chi_l + a_l^j \varphi_j, \cdot\}$, when evaluated on the constraint surface against another constraint. Conversely, adding a linear combination of squares of second class constraint to a first-class constraint defines another first class constraint: $\varphi_j \mapsto \varphi_j + b_j^{kl} \chi_k \chi_l$, which eventually would be decomposed on the first-class generators φ_i of course. Then, reinterpreting Proposition 4.48 in light of these ambiguities, saying that the Poisson bracket of two first-class constraints is again a first-class function – hence a first-class constraint – amounts to saying that it is linear in first class constraints and square in second class. Moreover, the result of the Poisson brackets between first class constraints with second class constraints must be linear in first-class constraints and quadratic in second-class constraints. This is a heuristic way of checking that first-class and second-class constraints are well-defined, see e.g. section 1.3.1 in [Henneaux and Teitelboim, 1992] or the examples treated by Alberto Escalante on ArXiv where he mentions several times this property.

Remark 4.53. Our procedure to obtain a set Ω admitting a splitting of first-class and second-class constraints is particular. See Chapter 2 of [Dirac, 1964] for a less straightforward but maybe more concrete explanation of how to isolate the first-class and second-class constraints without block diagonalizing D .

During the above process of finding independent first-class and second-class constraints, we have certainly mixed up primary and secondary constraints. There is a finer way of doing things where we can preserve the primary or secondary character of constraints. There exists linear combinations of primary (resp. secondary) constraints which splits this set into first-class and second-class functions, with respect to the entire set of constraints. We call these functions *first-class primary* (resp. *secondary*) *constraints*. Notice that although the Poisson bracket of two first-class constraints is first-class as explained in Proposition 4.48, it does not mean it is still a primary constraint and in general it will not be. Notice as well that, even if the splitting of primary constraints into first-class primary and second-class primary involves primary constraints only, we cannot make this splitting *before* having found every secondary constraints. A strategy to perform this splitting can be found in subsection 1.1.10 of [Henneaux and Teitelboim, 1992] and possibly pp. 25-27 of [Gitman and Tyutin, 1990].

Here is how we can proceed: let $\phi_\alpha^{(k)}$ denote the k -th stage constraints; in particular first-stage constraints are denoted $\phi_\alpha^{(1)}$, that we suppose to be irreducible. Then the persistence equations are:

$$\{\phi_\alpha^{(k)}, H\} + u^\beta \{\phi_\alpha^{(k)}, \phi_\beta^{(1)}\} \approx 0 \quad (4.64)$$

These are differential equations where the unknown variables are the u^β . The general solutions of these equations are given by:

$$u^\beta = U^\beta + V^\beta$$

where U^β are particular solutions of Equations (4.64) while V^β are solutions of the associated homogeneous equations:

$$V^\beta \{\phi_\alpha^{(k)}, \phi_\beta^{(1)}\} \approx 0 \quad (4.65)$$

These solutions can be decomposed over a set of independent solutions V_i^β of Equation (4.64), so that $V^\beta = w^i V_i^\beta$. Then we set $\varphi_i^{(1)} = V_i^\beta \phi_\beta^{(1)}$, and these constraints satisfy:

$$\{\varphi_i^{(1)}, \varphi_\alpha^{(k)}\} = V_i^\beta \{\phi_\beta^{(1)}, \varphi_\alpha^{(k)}\} + \phi_\beta^{(1)} \{V_i^\beta, \varphi_\alpha^{(k)}\} \approx 0$$

Thus the $\varphi_i^{(1)}$ are first class functions with respect to the initial set of primary and secondary constraints. Moreover they are independent as the functions V_i^β are independent.

Then by construction any first-class function of the constraints that is also a linear combination of the primary constraints – i.e. it can be written $f = f^\beta \phi_\beta^{(1)}$ and satisfies Equation (4.65) – is functionally dependent on the first-class primary constraints. We can complete the set $\{\varphi_i^{(1)}\}$ by a set of constraints $\chi_k^{(1)}$ obtained from linear combinations of primary constraints such that the entire set $\Omega^{(1)} = \{\varphi_i^{(1)}, \chi_k^{(1)}\}$ is irreducible, then equivalent to the initial set of primary constraints $\{\phi_\alpha^{(1)}\}$. Then the constraints $\chi_k^{(1)}$ are automatically second class (with respect to the entire set of all constraints), for otherwise the exception would be a linear combination of first-class primary constraints. In particular their number is minimal. See also [Gotay and Nester, 1984] for an interesting discussion about the interactions between the Bergmann-Dirac algorithm and the separation of first-class constraints from second-class constraints.

The use of first-class and second-class constraints has an enlightening consequence on the treatment of Hamilton's equations of motion. The total Hamiltonian (4.28) contains primary constraints which, when split into first-class and second-class constraints, can be recasted as:

$$H_T = H + w^i \varphi_i^{(1)} + w^k \chi_k^{(1)} \quad (4.66)$$

where for the sake of the argumentation we assume for now that the parameters w^i, w^k are not specified. Assume that the set of secondary constraints has also been split into independent

first-class constraints φ_i and second-class constraints χ_k , where the latter set is the smallest possible (this is possible, see page 27 in [Gitman and Tyutin, 1990]). Then for consistency both primary and secondary constraints should satisfy the persistence equation (4.45). By definition of first-class constraints φ_i and second-class constraints χ_k , Equation (4.50) becomes:

$$\begin{aligned}\{\varphi_i, H\} &\approx 0 \\ \{\chi_k, H\} + w^l \{\chi_k, \chi_l^{(1)}\} &\approx 0\end{aligned}\tag{4.67}$$

Since we assumed that the number of independent second-class constraints is minimal, the square matrix C whose coefficient is $C_{mn} = \{\chi_m, \chi_n\}$ is invertible and one can then compute explicitly the parameters w^l as functions of the canonical variables.

$$w^l = -(C^{-1})^{lk} \{\chi_k, H\}\tag{4.68}$$

where the sum is made over all k , and where l is a label associated to a primary second-class constraint only. The difference between this equation and Equation (4.69) is that we may have more than $2d$ specified parameters, where $2d$ is the rank of the matrix in Proposition 4.33, so that we can find as many parameters w^l as there are second-class primary constraints. We see that the splitting of primary constraints into first-class and second-class constraints allow us to identify which ‘velocities’ (equivalently represented by the coefficients u^α , v^β or w^i and w^k) can be fully determined. We see on Equation (4.67) that the parameters w^i associated to the primary first-class constraints in Equation (4.66) cannot be determined from Hamilton equations. They will remain undetermined so they are then free functions of time and will play the role of gauge parameters of the system.

Moreover notice that we have defined the parameters w^l by a strong equation, rather than as a weak equality. The classical equations of motion are insensitive to this, but it will turn out to be relevant for defining the corresponding quantum theory. Moreover, under the assumption that we have a minimal set of second-class constraints, the solution w^l of Equation (4.68) are unique for the following reason: indeed, the only freedom in the choice of w^l would come from the solutions of the homogeneous equation $w^l \{\chi_k, \chi_l^{(1)}\} \approx 0$, but if w^l is such an arbitrary solution, then it would mean that $w^l \chi_l^{(1)}$ is a first class constraint because it commutes with all second class constraints, which is impossible given that it would then be equal to a linear combination of first-class primary constraints. See subsection 3.5 in [Rothe and Rothe, 2010] for additional informations.

In the case where there are no primary first-class constraints, i.e. if the primary constraints are all second-class and the matrix M defined in Equation (4.46) has rank $2d = n - R_{\mathcal{L}}$, then it is invertible and *all* ‘velocities’ u^α associated to the primary constraints are fully determined via Equation (4.45). We can then define them strongly:

$$u^\beta = -(M^{-1})^{\beta\alpha} \{\phi_\alpha, H\}\tag{4.69}$$

so the equation $\dot{\phi}_\alpha = 0$ holds strongly. The fact that, in that case, there are no secondary constraints has two consequences. First, $\Gamma = \Gamma^{(1)}$ so for any choice of initial point $(q_0, p_0) \in \Gamma$, the physical path starting from this point and satisfying Hamilton’s equations (4.39) necessarily stays on the constraint surface. Second, in the present context, the total Hamiltonian is the sum of a global extension H of H_0 and a linear combination of second-class constraints and it additionally satisfies Equation (4.44) (by definition of the persistence equations). It is then a first-class function, and it is then legitimate to ask is such sum $-H +$ second class constraints – is always a first-class function, even in the case where the primary constraints are not all second-class (and thus when there exists first-class primary constraints):

Definition 4.54. Let H be the smooth function which coincides with H_0 on the primary constraint surface Γ used into Definition 4.26, and let w^l be the uniquely defined parameters associated to the primary second-class constraints $\chi_l^{(1)}$ in Equation (4.68). Then we define the first-class Hamiltonian as:

$$H' = H + w^l \chi_l^{(1)}$$

Authors often consider that it is permissible to add any linear combination of primary first-class constraints to H' (this would correspond to a rewriting of the free parameters w^i), but for clarity we would not consider this option here. This occurs in particular if one chose another particular solution U^α of Equation (4.64) (see subsection 1.1.10 of [Henneaux and Teitelboim, 1992]). While it was obvious that the total Hamiltonian was a first-class function because it satisfied Equation (4.49), it turns out that this property is also shared by the first-class Hamiltonian, justifying its name:

Lemma 4.55. *The first-class hamiltonian is a first-class function, i.e. the first class hamiltonian provides a splitting of H_T into two first-class functions:*

$$H_T = H' + w^i \varphi_i^{(1)} \tag{4.70}$$

Proof. The so-called first-class hamiltonian is indeed first-class: for any constraint ϕ_α (be it primary or secondary), we have by Equation (4.70):

$$\{H', \phi_\alpha\} = \{H_T, \phi_\alpha\} - \{w^i, \phi_\alpha\} \varphi_i^{(1)} - w^i \{\varphi_i^{(1)}, \phi_\alpha\} \approx 0$$

The first term vanishes on-shell – i.e. on the constraint surface Σ – by Example 4.46. The second term is proportional to a constraint while the third term vanishes on Σ because $\varphi_i^{(1)}$ is a first-class constraint. We see that the first-class hamiltonian is not uniquely defined because any other choice of free parameters w^i still gives a first-class function. \square

We will now study the relationship between first-class constraints and gauge transformations, and what does the latter mean in the Hamiltonian context. Ideally, the physical state of a system at any time t should be determined by a unique point $(q(t), p(t))$, if the path $t \mapsto (q(t), p(t))$ satisfies the Hamilton equations of motion. However, it may well happen that at each time t , the state of the system can be specified by various, equivalent points of the phase space. In other words, although the state of the system at time t is uniquely defined once given a point $(q(t), p(t))$, the converse is not true, i.e. there is more than one set of values of the canonical variables representing the same physical state. Ideally, we would expect the equations of motions to fully determine the time evolution of physical states. However, we have seen that some parameters in the total Hamiltonian – those w^i associated with the primary first-class constraints $\varphi_i^{(1)}$ – are still unspecified. This implies that, given a physical state at time t_1 , determined by a point $(q(t_1), p(t_1))$, the solution of the equations of motion corresponds to a path $t \mapsto (q(t), p(t))$ in phase space which depends on the value of the afore mentioned free parameters w^i , until a terminal state at time t_2 . Although different such parameters induce different endpoints, we consider that a physical observable shall not depend on such arbitrary smooth variation – *because* they are arbitrary. In other words, any ambiguity in the canonical variables at any time should be a physically irrelevant ambiguity: this precisely characterizes *gauge theories*. We will now reformulate Definition 4.4 to be more amenable to computations (see [Earman, 2003]):

Definition 4.56. *Two points $(q(t_2), p(t_2)), (q'(t_2), p'(t_2)) \in \Sigma$ are gauge equivalent if they are both obtained from a point $(q(t_1), p(t_1)) \in \Sigma$ as solutions to Hamilton's equations (4.39) in the same lapse time $\delta t = t_2 - t_1$.*

Remark 4.57. This definition concerns points in phase space. It can be straightforwardly adapted to *histories*, see e.g. Def. 2 in [Earman, 2003], while Def. 1 is Definition 4.56. The consequences of this slight difference of perspective has generated heated debates in the philosophy of physics [Pooley and Wallace, 2022, Pitts, 2022]. We will not address these questions here, but advise the interested reader to turn to the rich literature addressing the notion of gauge symmetries in the philosophy of physics; for an overview of the field see [Gomes, 2021].

We will now explore in what sense the primary first-class constraints generate (at least part of) the gauge transformations. The idea of gauge transformations, as any regular symmetry in physics, is to preserve the form of equations of motions. This is a sort of covariance principle, as the one from general relativity, and one can consider gauge transformations as a sort of arbitrary change of coordinates. An infinitesimal gauge transformation is thus defined on the configuration space Q , so its time derivative defines a small variation in velocities (i.e. it would correspond to a tangent vector over the tangent bundle TQ). It can be then transported to the phase space via the Legendre transform, where constraints are defined, and corresponds to a vector field on T^*Q . We will then study the notion of gauge transformations based on the behavior of smooth functions under such transformations. A gauge transformation would define a transformation of a smooth function at a given time, and then could indeed be seen as a vector field on T^*Q .

Let f be a smooth function on T^*Q and let us compute two different time evolutions of f depending of two different sets of arbitrary smooth parameters attached to the primary first-class constraints. Let $t_2 = t_1 + \delta t$, then by Lemma 4.32 and Equation (4.70) we have:

$$\begin{aligned} f(t_2) &= f(t_1) + \dot{f}(t_1)\delta t + \mathcal{O}((\delta t)^2) \\ &\approx f(t_1) + \{f, H_T\}(t_1)\delta t + \mathcal{O}((\delta t)^2) \\ &\approx f(t_1) + \{f, H'\}(t_1)\delta t + w^i \{f, \varphi_i^{(1)}\}(t_1)\delta t + \mathcal{O}((\delta t)^2) \end{aligned} \quad (4.71)$$

where the w^i are a set of unspecified smooth functions associated to the primary first-class constraints $\varphi_i^{(1)}$, possibly depending on time t . Now, if one takes up another set of parameters w'^i , the time evolution of f is now:

$$f(t_2) \approx f(t_1) + \{f, H'\}(t_1)\delta t + w'^i \{f, \varphi_i^{(1)}\}(t_1)\delta t + \mathcal{O}((\delta t)^2) \quad (4.72)$$

Notice that here, although we used the same notation, the value of $f(t_2)$ differ in both expressions (4.71) and (4.72) because the point $(q(t_2), p(t_2))$ is not the same in both cases. By construction, the first-class Hamiltonian H' is however the same in (4.71) and (4.72). Then they cancel out when we compute the difference between the two expressions of $f(t_2)$ and we have:

$$\delta f \approx -\delta\epsilon^i \{f, \varphi_i^{(1)}\}(t_1)\delta t + \mathcal{O}((\delta t)^2)$$

where $\delta\epsilon^i = (w'^i - w^i)\delta t$. Thus, the primary first-class constraints $\varphi_i^{(1)}$ are generators of local transformations with infinitesimal parameters $(w'^i - w^i)\delta t$. More precisely, denoting the hamiltonian vector field associated to $\varphi_i^{(1)}$ by $X_{\varphi_i^{(1)}}$, one has:

$$\delta f \approx \delta\epsilon^i X_{\varphi_i^{(1)}}(f) + \mathcal{O}((\delta t)^2) \quad (4.73)$$

Since the parameters w^i and w'^i – and hence ϵ^i – are arbitrary, these transformations can be considered as legitimate gauge transformations, see Figure 22. In other words, the primary first-class constraints generate (a subset of) gauge transformations in the sense that the flow of the Hamiltonian vector fields $X_{\varphi_i^{(1)}}$ are precisely gauge transformations.




Figure 22: A gauge theory is a physical theory where the solutions of the equations of motions contain arbitrary smooth functions of time, here the parameters w^i and w'^i . They are free smooth parameters that one can choose at any time, thus giving different time evolutions of a physical path $t \mapsto (q(t), p(t))$. The difference between the two different values of a function f at the same time t_2 characterizes the gauge transformations generated by the primary first-class constraints $\varphi_i^{(1)}$.

We will now show that the primary first-class constraint cannot be the only set of functions generating gauge transformations. Indeed, starting from the former discussion, assume that between $t_2 = t_1 + \delta t$ and $t_3 = t_2 + \delta t$ we decide to apply $H' + w^i \varphi_i^{(1)}$ to $f(t_2)$ as defined in Equation (4.71) and $H' + w'^i \varphi_i^{(1)}$ to $f(t_2)$ as defined in Equation (4.72). Then we compare $f(t_3)$ obtained via the first path, and $f(t_3)$ obtained via the second path. Since we expect that any ambiguity in the canonical variables at any time should be a physically irrelevant ambiguity, we deduce that the difference between the two values of f at time t_3 is the result of a gauge transformation.

Proposition 4.58. *Beyond the primary first-class constraints, the set of generators of gauge transformations contains the following smooth functions:*

1. the Poisson bracket of any two primary first-class constraints;
2. the Poisson bracket of any primary first-class constraint and the first-class Hamiltonian.

Proof. We first apply $H' + w^i \varphi_i^{(1)}$ to $f(t_2)$ as defined in (4.71) and we obtain:

$$\begin{aligned}
f(t_3) &= f(t_2) + \{f, H'\}(t_2)\delta t + w^j \{f, \varphi_j^{(1)}\}(t_2)\delta t + \mathcal{O}((\delta t)^2) \\
&\approx f(t_1) + \{f, H'\}(t_1)\delta t + w^i \{f, \varphi_i^{(1)}\}(t_1)\delta t \\
&\quad + \{f + \{f, H'\}\delta t + w^i \{f, \varphi_i^{(1)}\}\delta t, H'\}(t_1)\delta t \\
&\quad + w^j \{f + \{f, H'\}\delta t + w^i \{f, \varphi_i^{(1)}\}\delta t, \varphi_j^{(1)}\}(t_1)\delta t + \mathcal{O}((\delta t)^3) \\
&\approx f(t_1) + 2\{f, H'\}\delta t + w^i \{f, \varphi_i^{(1)}\}\delta t + w^j \{f, \varphi_j^{(1)}\}\delta t + \{\{f, H'\}, H'\}(\delta t)^2 \\
&\quad + \{f, \varphi_i^{(1)}\}\{w^i, H'\}(\delta t)^2 + w^i \{\{f, \varphi_i^{(1)}\}, H'\}(\delta t)^2 + w^j \{\{f, H'\}, \varphi_j^{(1)}\}(\delta t)^2 \\
&\quad + w^j w^i \{\{f, \varphi_i^{(1)}\}, \varphi_j^{(1)}\}(\delta t)^2 + w^j \{f, \varphi_i^{(1)}\}\{w^i, \varphi_j^{(1)}\}(\delta t)^2 + \mathcal{O}((\delta t)^3) \tag{4.74}
\end{aligned}$$

where the evaluation at time t_1 is implicit for each term. Then we apply $H' + w^i \varphi_i^{(1)}$ to $f(t_2)$ as defined in (4.72), and we obtain:

$$\begin{aligned}
f(t_3) &\approx f(t_1) + 2\{f, H'\}\delta t + w^i \{f, \varphi_i^{(1)}\}\delta t + w^j \{f, \varphi_j^{(1)}\}\delta t + \{\{f, H'\}, H'\}(\delta t)^2 \\
&\quad + \{f, \varphi_i^{(1)}\}\{w^i, H'\}(\delta t)^2 + w^i \{\{f, \varphi_i^{(1)}\}, H'\}(\delta t)^2 + w^j \{\{f, H'\}, \varphi_j^{(1)}\}(\delta t)^2 \\
&\quad + w^j w^i \{\{f, \varphi_i^{(1)}\}, \varphi_j^{(1)}\}(\delta t)^2 + w^j \{f, \varphi_i^{(1)}\}\{w^i, \varphi_j^{(1)}\}(\delta t)^2 + \mathcal{O}((\delta t)^3) \tag{4.75}
\end{aligned}$$

where again the evaluation at time t_1 is implicit. Computing the difference between Equations (4.74) and (4.75) one obtains, after reordering the terms and noticing that their respective first line cancel each other:

$$\begin{aligned}
\delta f &\approx \left(\overline{\{f, \varphi_i^{(1)}\}\{w^i, H'\}} - \overline{\{f, \varphi_i^{(1)}\}\{w^i, H'\}} \right) (\delta t)^2 \\
&\quad + w^i \left(\overline{\{\{f, \varphi_i^{(1)}\}, H'\}} - \overline{\{\{f, H'\}, \varphi_i^{(1)}\}} \right) (\delta t)^2 \\
&\quad + w^j \left(\overline{\{\{f, H'\}, \varphi_j^{(1)}\}} - \overline{\{\{f, \varphi_j^{(1)}\}, H'\}} \right) (\delta t)^2 \\
&\quad + w^j w^i \left(\overline{\{\{f, \varphi_i^{(1)}\}, \varphi_j^{(1)}\}} - \overline{\{\{f, \varphi_j^{(1)}\}, \varphi_i^{(1)}\}} \right) (\delta t)^2 \\
&\quad + \left(w^j \{f, \varphi_i^{(1)}\}\{w^i, \varphi_j^{(1)}\} - w^j \{f, \varphi_i^{(1)}\}\{w^i, \varphi_j^{(1)}\} \right) (\delta t)^2 + \mathcal{O}((\delta t)^3)
\end{aligned}$$

Let us show that the two underlined terms combine to give the term $\{f, \{w^i \varphi_i^{(1)}, H'\}\}$ on the constraint surface. Then it is straightforward to antisymmetrize the computation and deduce that the two overlined terms give the term $-\{f, \{w^j \varphi_j^{(1)}, H'\}\}$, again on Σ . Indeed, since the constraints vanish on Σ , we can rewrite the first term while the regarding the second we use the Jacobi identity for the Poisson bracket:

$$\{f, \varphi_i^{(1)}\}\{w^i, H'\} + w^i \left(\overline{\{\{f, \varphi_i^{(1)}\}, H'\}} - \overline{\{\{f, H'\}, \varphi_i^{(1)}\}} \right) \approx \{f, \varphi_i^{(1)}\}\{w^i, H'\} + w^i \{f, \{\varphi_i^{(1)}, H'\}\}$$

The last term is weakly equivalent to the following one: $\{f, w^i \{\varphi_i^{(1)}, H'\}\}$, because H' being first-class, its bracket with any linear combination of constraints such as $\{\varphi_i^{(1)}, H'\}$ vanishes on Σ . Thus, the sum $\{f, \varphi_i^{(1)}\}\{w^i, H'\} + w^i \{f, \{\varphi_i^{(1)}, H'\}\}$ is weakly equivalent to $\{f, \{w^i \varphi_i^{(1)}, H'\}\}$ as desired.

Now, let us compute the sum $w^j w^i \{\{f, \varphi_i^{(1)}\}, \varphi_j^{(1)}\} + w^j \{f, \varphi_i^{(1)}\}\{w^i, \varphi_j^{(1)}\}$, when it is restricted to the constraint surface. Start by factorizing out w^j and make w^i enter the bracket so that the sum is weakly equivalent to:

$$w^j w^i \{\{f, \varphi_i^{(1)}\}, \varphi_j^{(1)}\} + w^j \{f, \varphi_i^{(1)}\}\{w^i, \varphi_j^{(1)}\} \approx w^j \{\{f, \varphi_i^{(1)}\} w^i, \varphi_j^{(1)}\}$$

Then the bracket $\{\{f, \varphi_i^{(1)}\}w^i, \varphi_j^{(1)}\}$ is weakly equivalent to $\{\{f, \varphi_i^{(1)}w^i\}, \varphi_j^{(1)}\}$ because:

$$\begin{aligned} \{\{f, \varphi_i^{(1)}w^i\}, \varphi_j^{(1)}\} &= \{\{f, \varphi_i^{(1)}\}w^i, \varphi_j^{(1)}\} + \{\varphi_i^{(1)}\{f, w^i\}, \varphi_j^{(1)}\} \\ &= \{\{f, \varphi_i^{(1)}\}w^i, \varphi_j^{(1)}\} + \{\varphi_i^{(1)}, \varphi_j^{(1)}\}\{f, w^i\} + \varphi_i^{(1)}\{\{f, w^i\}, \varphi_j^{(1)}\} \end{aligned}$$

Both last terms vanish on the constraint surface; the middle one because the Poisson bracket of two first class constraints vanishes on Σ by definition. Then, it turns out that the term $w^j\{\{f, \varphi_i^{(1)}\}w^i, \varphi_j^{(1)}\}$ is weakly equivalent to $w^j\{\{f, \varphi_i^{(1)}w^i\}, \varphi_j^{(1)}\}$, which in turn is weakly equivalent to $\{\{f, \varphi_i^{(1)}w^i\}, \varphi_j^{(1)}w^j\}$. By antisymmetry, we straightforwardly deduce that the sum $-w^jw^i\{\{f, \varphi_i^{(1)}\}, \varphi_j^{(1)}\} - w^i\{f, \varphi_j^{(1)}\}\{w^j, \varphi_i^{(1)}\}$ is weakly equivalent to $-\{\{f, \varphi_j^{(1)}w^j\}, \varphi_i^{(1)}w^i\}$. By the Jacobi identity, the two terms combine and give:

$$\{\{f, \varphi_i^{(1)}w^i\}, \varphi_j^{(1)}w^j\} - \{\{f, \varphi_j^{(1)}w^j\}, \varphi_i^{(1)}w^i\} = \{f, \{\varphi_i^{(1)}w^i, \varphi_j^{(1)}w^j\}\}$$

Gathering all the simplifications we obtained, we have the following weak equivalence, which characterize the gauge transformation applied to f :

$$\delta f \approx \{f, \{w^i\varphi_i^{(1)}, H'\}\}(\delta t)^2 - \{f, \{w^j\varphi_j^{(1)}, H'\}\}(\delta t)^2 + \{f, \{w^i\varphi_i^{(1)}, w^j\varphi_j^{(1)}\}\}(\delta t)^2 + \mathcal{O}((\delta t)^3)$$

where the evaluation at time t_1 is implicit. Then we see that the gauge transformation δf is generated by the three terms $\{w^i\varphi_i^{(1)}, H'\}$, $\{w^j\varphi_j^{(1)}, H'\}$ and $\{w^i\varphi_i^{(1)}, w^j\varphi_j^{(1)}\}$. Since the parameters w^i and w^j are arbitrary, we have the result. \square

Proposition 4.58 shows us that primary first-class constraints are not the only functions acting as generators of gauge transformations. The Poisson brackets $\{w^i\varphi_i^{(1)}, H'\}$, $\{w^j\varphi_j^{(1)}, H'\}$ and $\{w^i\varphi_i^{(1)}, w^j\varphi_j^{(1)}\}$ should indeed generate gauge transformations as well. Since by definition of the Bergmann-Dirac algorithm the bracket $\{\varphi_i^{(1)}, H'\} \approx \{\varphi_i^{(1)}, H\}$ is a second-stage constraint, we deduce that the two first brackets involve secondary constraints. Moreover, the third bracket $\{w^i\varphi_i^{(1)}, w^j\varphi_j^{(1)}\}$ vanishes on Σ by definition of first-class constraints, hence it is strongly equivalent to a linear combination of primary and secondary constraints. Eventually, by Proposition 4.48 we know that the brackets $\{w^i\varphi_i^{(1)}, H'\}$, $\{w^j\varphi_j^{(1)}, H'\}$ and $\{w^i\varphi_i^{(1)}, w^j\varphi_j^{(1)}\}$ are first-class functions. Together with the above arguments, it implies that these brackets are not only strongly equivalent to linear combinations of primary and secondary constraints, but also that these constraints are all *first-class*.

Although it is in general not possible to infer from these observations alone that *every* secondary first-class constraint generates a gauge transformation, we will usually assume that it is the case (this is the Dirac conjecture, Scholie 4.59). Indeed, we have seen that the distinction between primary and secondary constraint is contingent because it heavily relies on the original choice of coordinates when we perform the Legendre transform or on the functional form of the Lagrangian. On the contrary, first-class and second-class constraints is a fundamental distinction, brought up by the Poisson structure of T^*Q . Additionally, first-class constraints form a Lie subalgebra of $\mathcal{C}^\infty(T^*Q)$ so they form an ideal candidate for generators of gauge transformations, but one has then to consider *all of them*. Moreover, we will see later that quantization methods put the first-class constraints on the same footings; there is no known quantization scheme if one does only consider part of them as gauge generators. For more details and further discussion, see the fruitful subsection 1.2.1 of [Henneaux and Teitelboim, 1992]. These observations led Dirac to formulate the following assumption:

Scholie 4.59. Dirac conjecture. *The generators of the gauge transformations are the first-class constraints, both primary and secondary.*

The status of the Dirac conjecture is debated. Its name first – a conjecture – would in general implicitly say that it has not yet been proven, however see subsection 3.3.2 of [Henneaux and Teitelboim, 1992] for a proof of Dirac conjecture under mild assumptions. Fifty years of heated discussion have shown that the well-grounded character of this statement seems to mostly depend on its interpretation. For example, Henneaux and Teitelboim use the following Lagrangian defined on \mathbb{R}^2 (subsection 1.2.2 of [Henneaux and Teitelboim, 1992]):

$$L = \frac{1}{2}e^y v_x^2$$

The first constraint is a primary first-class constraint $\varphi^{(1)} = p_y$, and it induces a unique secondary first-class constraint $\varphi^{(2)} = \frac{1}{2}e^{-y}p_x^2$ which turns out to coincide with the Hamiltonian. There are no other constraints. Then, if one considers that the true secondary first-class constraint is $\tilde{\varphi}^{(2)} = p_x$ – as Henneaux and Teitelboim did – one observes that it does not generate gauge transformations. These authors chose to pass from $\varphi^{(2)}$ to the mathematically equivalent constraint $\tilde{\varphi}^{(2)}$ because they considered as ‘true’ constraints those that can serve as coordinate transverse to the secondary constraint surface.

However, Rothe and Rothe have shown (subsection 6.4 of [Rothe and Rothe, 2010]) that if one sticks to the secondary first-class constraint $\varphi^{(2)}$ then it generates a gauge transformation. They say that the ambiguity in the Dirac conjecture comes from an ambiguity in the interpretation of what is a ‘true’ first-class constraint, and that the validity of Dirac conjecture depends crucially on the chosen form for the constraints. The replacement of constraints by a formally equivalent set of constraints – choosing $\tilde{\varphi}^{(2)}$ instead of $\varphi^{(2)}$ – in fact may obliterate the full symmetry of the total action and will lose some important physical informations. Hence, this example shows that mathematically equivalent constraints may not be physically equivalent. A way of avoiding such a trouble is to accept both $\tilde{\varphi}^{(2)}$ and $\varphi^{(2)}$ in the new set of constraints. There would be some redundancy but at least both the regularity condition and the fact that $\varphi^{(2)}$ is first-class would be satisfied. Discussions about the Dirac conjecture have been vivid in the 1980s and the literature on the topic is rich [Gotay, 1983, Gotay and Nester, 1984, Stefano, 1983, Cabo, 1986, Gràcia and Pons, 1988, Cabo and Louis-Martinez, 1990, Pons, 2005]. In particular, it seems admitted now that a first-class constraint generates a gauge transformation if the iterative procedure which generates it does not pass through an *ineffective constraint*, i.e. a constraint whose gradient vanishes weakly [Earman, 2003], as in the above example. From now, we will stick to the modern view that the conjecture holds.

Now that we have determined all the generators of gauge transformations, we soon realize that the total Hamiltonian H_T does not contain every first-class constraints, and thus cannot generate all the gauge transformations. Thus we are led to adding the remaining first-class constraints to H_T to obtain a proper, more general Hamiltonian:

Definition 4.60. *Assume that there are p first-class constraints φ_i in total (both primary and secondary) and let w^1, \dots, w^p be arbitrary smooth parameters on the canonical coordinates (possibly depending on time also). Then we define the extended Hamiltonian as the following smooth function:*

$$H_E = H' + w^i \varphi_i$$

where H' is the first-class Hamiltonian.

Remark 4.61. As was said in Remark 4.27, the extended Hamiltonian, even more so than the total Hamiltonian, can be considered as the equivalence class of all the first-class Hamiltonians equivalent to the first-class Hamiltonian by addition of first-class constraints. This should certainly correspond to being a cocycle in a particular cohomology.

Thus, the extended Hamiltonian contains the primary second-class constraints (hidden into H') and all the first-class constraints. Equivalently, H_E is the sum of the total hamiltonian H_T with all secondary first-class constraints. When Hamilton equations involve the extended Hamiltonian, all the gauge transformations are allowed to be performed. However, a physical observable, being by definition gauge invariant, should not depend on such gauge transformations. Hence the choice of Hamiltonian one picks up in the equations of motion – be it H' , H_T or H_E – will not have any consequence on the smooth functions that are physically relevant, but will impact any other smooth function. Notice however that, while the total Hamiltonian was directly obtained from the Lagrangian formalism and would give back the Euler-Lagrange equations (see Proposition 4.39), the extended Hamiltonian is a new feature of the Hamiltonian formalism that does not have a Lagrangian counterpart (see subsection 1.2.3 in [Henneaux and Teitelboim, 1992]). The extended Hamiltonian allows to set all the first-class constraints on the same footing, which is useful to apply canonical quantization. The difference between the total and the extended Hamiltonians are further explored in sections 3.2 and 3.3 of [Henneaux and Teitelboim, 1992], as well as section 5.4 of [Rothe and Rothe, 2010]. See [Brown, 2022] for a full treatment of a singular Lagrangian following Bergmann-Dirac algorithm, where the extended Hamiltonian appears.

4.5 The geometry of the constraint surface

Let us now address the geometrical meaning of first-class and second-class constraints. Fix once and for all some irreducible first-class and second-class constraints, and assume that the latter are minimal in number. Then the zero-level set of the second-class constraints is called the *second-class constraint manifold* and forms a *cosymplectic* submanifold Σ_0 of $(M, \{.,.\})$ (see Section 3.3). It is assumed to be an embedded submanifold, and that the rank of the matrix D has constant rank over it (and not only on Σ). We then know that in a tubular neighborhood W of Σ_0 (or at least locally) one can define a Poisson bracket – called the *Dirac bracket*, see Equation (3.42) – so that Σ_0 is a symplectic leaf of $(W, \{.,.\}_{Dirac})$. In particular, we have shown that the Poisson bracket on Σ_0 induced by the Poisson-Dirac reduction coincides with the Dirac bracket (see Equation (3.48)) [Śniatycki, 1974]. So the second-class manifold is a symplectic manifold and can be taken to be a replacement of the original phase space. Then the first-class constraint define a submanifold of Σ_0 , which turns out to be the constraint surface Σ . For simplicity in the following we will often assume that the constraints are defined globally over the entire phase space T^*Q , but remember that in full generality the results are only defined on a tubular neighborhood $W \subset T^*Q$ of Σ , or at least locally around each point:

Lemma 4.62. *The constraint surface Σ is a coisotropic submanifold of the second-class constraint manifold $(\Sigma_0, \{.,.\}_{\Sigma_0})$, where $\{.,.\}_{\Sigma_0}$ is the restriction of the Dirac bracket to Σ_0 . When the Dirac bracket is defined globally over T^*Q , then the constraint surface Σ is a coisotropic submanifold of $(T^*Q, \{.,.\}_{Dirac})$.*

Proof. The second-class constraint manifold Σ_0 is a symplectic leaf of the Dirac bracket, so the restriction $\{.,.\}_{\Sigma_0}$ is well-defined (see Proposition 3.83), and $(\Sigma_0, \{.,.\}_{\Sigma_0})$ is a symplectic manifold. Then, let us show the second point directly. Assuming that the constraints – both first-class *and* second-class – are globally defined, we have that the secondary constraint surface Σ is a closed embedded submanifold of T^*Q . Then by Proposition 3.92 it is sufficient to show that the ideal $\mathcal{I}_\Sigma = \text{Span}(\varphi_a, \chi_e)$ of vanishing functions on Σ generated in $\mathcal{C}^\infty(T^*Q)$ by the first class *and* the second class constraints is a Lie subalgebra of $(\mathcal{C}^\infty(T^*Q), \{.,.\}_{Dirac})$. The definition of the Dirac bracket, Equation (3.42), has the following consequences:

1. the second-class constraints χ_e are Casimir elements of the Dirac bracket, so the Dirac bracket with any of them vanish, so in particular on Σ ;
2. the first-class constraints φ_a are such that $\{\varphi_a, \mathcal{I}_\Sigma\}$ vanish on Σ by Definition 4.45, which implies that on Σ we have:

$$\{\varphi_a, \cdot\}_{Dirac} = \{\varphi_a, \cdot\} - \underbrace{\{\varphi_a, \chi_d\} C^{de}}_{\approx 0} \{\chi_e, \cdot\} = \{\varphi_a, \cdot\} \quad (4.76)$$

This implies in turn that $\{\varphi_a, \varphi_b\}_{Dirac} = \{\varphi_a, \varphi_b\}$ on Σ which, by Definition 4.45, vanish on Σ .

These observations show that we have that the smooth functions belonging to the set $\{\mathcal{I}_\Sigma, \mathcal{I}_\Sigma\}_{Dirac}$ vanish on the secondary constraint surface Σ , i.e. they belong to \mathcal{I}_Σ . This proves that \mathcal{I}_Σ is a subalgebra of $(C^\infty(T^*Q), \{\cdot, \cdot\}_{Dirac})$. \square

Remark 4.63. There exists however an alternative way of getting rid of the second class constraints: instead of using the Poisson-Dirac reduction, one extends the phase space so that the second class constraints become gauge fixing conditions of first-class constraints. This is called the BFT formalism – from Batalin, Fradkin and Tyutin who formalized it at the end of the 1980s – and is described in Section 1.4.3 of [Henneaux and Teitelboim, 1992], in Chapter 7 of [Rothe and Rothe, 2010], and in the references there-in. E As the constraint surface is a presymplectic submanifold of the second-class submanifold [Gotay et al., 1978], extending the phase space in such a way corresponds to solving the problem of coisotropic embedding of presymplectic manifolds into a bigger symplectic manifold [Gotay, 1982].

Although the secondary constraint surface Σ is a coisotropic submanifold of the symplectic manifold $(T^*Q, \{\cdot, \cdot\}_{Dirac})$, it is *not* a coisotropic submanifold of the symplectic manifold $(T^*Q, \{\cdot, \cdot\})$, with the standard Poisson/symplectic structure. Indeed, the ideal \mathcal{I}_Σ is not stable under Poisson bracket, for not all Poisson brackets of second-class constraints vanish on Σ , but it is stable under the Dirac bracket because the second-class constraints are Casimir elements of the latter. This situation has actually a great mathematical relevance: this is called the problem of coisotropic embedding. We say that a submanifold S of a symplectic manifold (M, ω) is *pre-symplectic* if the restriction of ω to TS has constant rank (but is not necessarily non-degenerate). In our case, one can check that the secondary constraint surface Σ is a pre-symplectic submanifold of T^*Q equipped with the canonical symplectic structure. Then we have two important results regarding pre-symplectic (sub)manifolds:

1. for any pre-symplectic submanifold S of a symplectic manifold M , there exists a symplectic submanifold C of M such that S is coisotropically embedded into C [Marle, 1983], and
2. conversely, for any pre-symplectic manifold P , there exists a symplectic manifold M such that P is coisotropically embedded into M [Gotay, 1982].

Thus, from these observations, one sees that knowing that the secondary constraint surface Σ is a pre-symplectic submanifold of T^*Q is sufficient to know that it embeds as coisotropic submanifold into a symplectic submanifold of T^*Q , that we can take to be the second-class submanifold Σ_0 . See this chapter for a clear presentation of this latter point approach. A generalization of the work of Marle and Gotay to Poisson geometry has been achieved by Cattaneo and Zambon in the 2000 [Cattaneo and Zambon, 2009].

The proof of Lemma 4.62 has a very interesting consequence: Equation (4.76) shows that on the (secondary) constraint surface Σ , the hamiltonian vector fields of any first-class function

f – either computed with respect to the original Poisson bracket $\{.,.\}$ or the Dirac bracket – coincide. Indeed, it is straightforward to see that on Σ , Equation (4.76) is actually valid for *any* first-class function. This has to do with the fact that a first-class function with respect to a set of constraint is first-class with respect to any other equivalent set (see Lemma 4.47). There is then no ambiguity of talking about hamiltonian vector fields of first-class functions, when we restrict ourself to the constraint surface Σ . In particular, that would be the case for the first-class hamiltonian H' and every first-class constraint φ_a . However for second-class constraints, it is another story because they form Casimir elements of the Dirac bracket. We then have the following important, geometric observation, with physical ramifications (Theorem 2.2 and 2.3 in [Henneaux and Teitelboim, 1992]):

Proposition 4.64. *The hamiltonian vector fields X_{χ_e} associated to the second class constraints $\{\chi_e\}$ – and computed with respect to the original Poisson bracket $\{.,.\}$ – are nowhere tangent to the second-class constraint manifold Σ_0 (and hence to Σ).*

The hamiltonian vector fields X_{φ_a} associated to the first-class constraints $\{\varphi_a\}$ are everywhere tangent to the secondary constraint surface Σ and moreover induce a regular foliation on Σ .

Proof. The hamiltonian vector fields associated to the second class constraints are nowhere tangent to the second-class constraint manifold Σ_0 because for any second-class constraint χ_d and any point of Σ_0 , there is at least one bracket $\{\chi_d, \chi_e\} = X_{\chi_d}(\chi_e)$ with another second-class constraint χ_e which does not vanish at this point (see Definition 4.45). So the action of the Hamiltonian vector field X_{χ_d} on the ideal of vanishing functions on Σ_0 never lands in this ideal, so X_{χ_d} is not tangent to Σ_0 , and hence to $\Sigma \subset \Sigma_0$.

On the contrary, by Definition 4.45 of first-class functions, $\{\varphi_a, \mathcal{I}_\Sigma\} \subset \mathcal{I}_\Sigma$, so the hamiltonian vector fields X_{φ_a} are tangent to Σ . Now let us show the last item: if the set of first class-constraints is not irreducible, the regularity condition 4.37 implies that there are at least k independent first-class constraints which generate all the others. This is a local condition because even if the constraints are defined over the whole of T^*Q , their generators may change. So, locally, the set of first-class constraints is generated by k constraints. Let D be the smooth distribution generated by the hamiltonian vector fields X_{φ_i} . It has constant, finite rank k .

Now let us show that it is involutive. Since, moreover, Σ is a coisotropic submanifold of $(T^*Q, \{.,.\}_{Dirac})$, the multiplicative ideal \mathcal{I}_Σ of vanishing functions on Σ is a subalgebra of $(\mathcal{C}^\infty(T^*Q), \{.,.\}_{Dirac})$, i.e. there exist smooth functions C_{ab}^c on T^*Q such that:

$$\{\varphi_a, \varphi_c\}_{Dirac} = C_{ab}^c \varphi_b \quad (4.77)$$

Then, by Equation (3.7) and Equation (3.42) we have:

$$[X_{\varphi_a}, X_{\varphi_b}] = X_{\{\varphi_a, \varphi_b\}} = X_{\{\varphi_a, \varphi_b\}_{Dirac}} + X_{\{\varphi_a, \chi_d\} C^{de} \{\chi_e, \varphi_b\}} \quad (4.78)$$

The second term on the right-hand side reads:

$$X_{\{\varphi_a, \chi_d\} C^{de} \{\chi_e, \varphi_b\}} = C^{de} \{\chi_d, \varphi_b\} X_{\{\varphi_a, \chi_d\}} + \{\varphi_a, \chi_d\} X_{C^{de} \{\chi_e, \varphi_b\}} \approx 0$$

It indeed vanishes on Σ because $\{\varphi_a, \chi_d\}$ and $\{\chi_e, \varphi_b\}$ vanish on Σ by definition of first-class functions. The first term on the right-hand side of Equation (4.78) reads:

$$X_{\{\varphi_a, \varphi_b\}_{Dirac}} = X_{C_{ab}^c \varphi_c} = \varphi_c X_{C_{ab}^c} + C_{ab}^c X_{\varphi_c} \approx C_{ab}^c X_{\varphi_c} \quad (4.79)$$

We then conclude that Equation (4.78) can be written on Σ as follows:

$$[X_{\varphi_a}, X_{\varphi_b}] \approx C_{ab}^c X_{\varphi_c} \quad (4.80)$$

Then, the regular distribution D is involutive on Σ (and a priori only on Σ). We say that the algebra of vector fields generated by the X_{φ_a} *closes on-shell*. By Frobenius theorem, it is integrable and induces a regular foliation (on Σ). \square

Remark 4.65. The structure functions C_{ab}^c appearing on the right-hand side of Equation (4.77) are not uniquely defined *outside* of the constraint surface, as one can always add a term $D_{ab}^{cd}\varphi_d$ to them, such that $D_{ab}^{cd} = -D_{ab}^{dc}$. But this additional term would vanish on Σ , so that, in the irreducible case, the structure functions are uniquely defined *on* the constraint surface. Moreover, Exercise 9.5 in [Henneaux and Teitelboim, 1992] establishes that such a modification of the structure functions C_{ab}^c could be absorbed by a change of coordinates on the phase space. In the reducible case, Equation (4.80) establishes that these structure functions C_{ab}^c may not be uniquely defined even on Σ , because in that case the Hamiltonian vector fields X_{φ_c} would not be independent on the constraint surface. This dependence prevents to extract the structure function from Equation (4.80), resulting in an ambiguity in their definition.

The definition of first-class functions tells us that if a hamiltonian vector field X_f is tangent to the secondary constraint surface Σ , then f is a first-class function. If additionally this function vanishes on Σ , it can be written as a combination of the first-class constraints φ_a . Then the vector field X_f decomposes on the basis X_{φ_a} , and it is tangent to the leaves of the regular foliation described in Proposition 4.64. By this observation, adding to X_{H_T} any hamiltonian vector field X_f tangent to the leaves of the regular foliation does not change Hamilton's equations of Proposition 4.39. This is why defining the extended Hamiltonian H_E by adding any combination of secondary first-class constraints to H_T does not change the physics contained in Hamilton's equation and is a very natural and most general thing to do. The integral curves of the hamiltonian vector fields $X_{H'}$, X_{H_T} and X_{H_E} will however be transversal to the leaves of the foliation if the hamiltonians H' , H_T and H_E do not themselves vanish on Σ . Indeed we will see later that the leaves of this foliation represent the gauge equivalent physical states of the model, so the Hamiltonian vector fields $X_{H'}$, X_{H_T} and X_{H_E} indeed indicate physical evolution of the system because their integral curves do not live in only one leaf (if H' , H_T and H_E are not weakly vanishing).

Remark 4.66. We understand now that the hamiltonian vector fields of the original primary and secondary constraints may not be tangent to Σ . Only particular combinations of them, giving the first-class constraints, have their Hamiltonian vector field tangent to Σ . These are the vector fields defining the regular foliation presented in Proposition 4.64. There is a priori no reason that the hamiltonian vector fields of the initial primary and secondary constraints generate this foliation.

By Lemma 4.62, the secondary constraint surface Σ is a coisotropic submanifold of $(\Sigma_0, \{ \cdot, \cdot \}_{\Sigma_0})$, and under the assumption that the leaf space corresponding to the foliation of Proposition 4.64 is a smooth manifold, one can proceed to Poisson reduction since the assumptions of Proposition 3.97 are satisfied. The leaf space is then called the *reduced phase space* and is denoted Σ_{ph} , because it corresponds to the true non-gauge equivalent physical states of the system. When the leaf space is a smooth manifold, Proposition 3.97 applies and the Dirac bracket descends from Σ to Σ_{ph} by Poisson reduction, so that the resulting Poisson bracket is non-degenerate on Σ_{ph} . On the reduced phase space the equations of motions are the usual Hamilton's ones (see Appendix 2.A in [Henneaux and Teitelboim, 1992] to obtain more informations on the symplectic structure on Σ_{ph}). The smooth functions on Σ_{ph} are the physical observables, and thus we should characterize the space $\mathcal{C}^\infty(\Sigma_{ph})$ explicitly before quantizing the theory. Although it would seem desirable to work on the reduced phase space, it turns out that one often loses desirable features of the physical model such as Lorentz manifest invariance or, in the case of field theory, polynomiality of fields and locality in space. Moreover, it is often impossible to reformulate the

theory in terms of gauge invariant quantities only and then to subsequently quantize it from the reduced phase space picture. It is thus often a better choice to carry along the dynamical variables and keep track of the first-class constraints, without using Poisson reduction, and then quantize the theory (see subsection 2.2.3 of [Henneaux and Teitelboim, 1992]). This is the object of BRST formalism, which provides an algebraic formulation of gauge invariant functions, i.e. physical observables.

Assume now that the constraints are globally defined, so that Σ is a closed embedded submanifold of T^*Q . The constraint surface is characterized by the ideal $\mathcal{I}_\Sigma = \text{Span}(\varphi_a, \chi^e)$ of vanishing functions on Σ generated in $\mathcal{C}^\infty(T^*Q)$ by the first class *and* the second class constraints. Then, by the proof of the second item of Lemma 3.72 we have:

$$\mathcal{C}^\infty(\Sigma) \simeq \mathcal{C}^\infty(T^*Q) / \mathcal{I}_\Sigma$$

By Lemma 4.62, Σ is a coisotropic submanifold of $(T^*Q, \{\cdot, \cdot\}_{Dirac})$ (or possibly only on a tubular neighborhood W of Σ). Indeed under the Dirac bracket the second class constraints behave as zero thus \mathcal{I}_Σ is a Lie subalgebra of $\mathcal{C}^\infty(T^*Q)$, i.e. $\{\mathcal{I}_\Sigma, \mathcal{I}_\Sigma\}_{Dirac} \subset \mathcal{I}_\Sigma$. Let us now make sense of the Poisson reduction to Σ_{ph} in light of the knowledge we have on gauge transformations.

One may consider the set of gauge transformations as a family of infinitesimal transformations on $\mathcal{C}^\infty(T^*Q)$, i.e. as vector fields on T^*Q . Each gauge transformation is proportional to a (set of) smooth parameters ϵ^i (which in Section 4.4 corresponds to the difference $w'^i - w^i$ for example), where a priori i ranged from 1 to p , the number of first-class constraints, both primary and secondary (see Dirac conjecture, Scholie 4.59). We can then consider the family of parameters ϵ^i as the respective components of a smooth section ϵ of the trivial vector bundle $E = \mathbb{R}^p \times T^*Q$. We then denote the corresponding gauge transformation as a vector field $\delta_\epsilon : \mathcal{C}^\infty(T^*Q) \rightarrow \mathcal{C}^\infty(T^*Q)$, acting on smooth function f as:

$$\delta_\epsilon(f) = \epsilon^i X_{\varphi_a}(f) \tag{4.81}$$

where the X_{φ_a} are the Hamiltonian vector fields associated to the first-class constraints φ_a (we know that on Σ they do not depend if we picked up the Poisson bracket or the Dirac bracket to define them). Recall that the vector fields are independent if the primary first-class constraints are irreducible. Moreover notice that the dependence in δt – which is explicit in Equation (4.73) – has been suppressed in Equation (4.81) because its role was only to emphasize the infinitesimal character of the transformations. Thus, we have defined a vector bundle map $\delta : E \rightarrow T(T^*Q)$ – corresponding to a $\mathcal{C}^\infty(T^*Q)$ -linear map $\delta : \Gamma(E) \rightarrow \mathfrak{X}(T^*Q)$ – sending a section $\epsilon \in \Gamma(E)$ to the corresponding vector field δ_ϵ defined in Equation (4.81).

The space $\mathfrak{X}(T^*Q)$ is a Lie algebra (of infinite dimension) of which the image of δ is an infinite dimensional subspace. The Lie algebra $\mathcal{G} \subset \mathfrak{X}(T^*Q)$ generated by the space of gauge transformations $\text{Im}(\delta)$ is abusively called the *algebra of gauge transformations*. It is generated by the Hamiltonian vector fields X_{φ_a} and their successive brackets:

$$[X_{\varphi_a}, X_{\varphi_b}], \quad [X_{\varphi_a}, [X_{\varphi_b}, X_{\varphi_c}]], \quad [X_{\varphi_a}, [X_{\varphi_b}, [X_{\varphi_c}, X_{\varphi_d}]]], \quad \text{etc.}$$

If the vector fields X_{φ_a} are involutive then we have the equality $\mathcal{G} = \text{Im}(\delta)$. If not, we nonetheless always have that \mathcal{G} defines an involutive –hence integrable – generalized distribution, to which corresponds a singular foliation [Hermann, 1962]. On the constraint surface, the leaves of this foliation coincide with that of the Hamiltonian vector fields X_{φ_a} , as they are involutive on Σ (see Equation (4.80)). Moreover, by Chow-Rashevskii theorem 2.72 and its Corollary 2.73, every two points belonging to the same leaf of \mathcal{G} outside of Σ are linkable by a smooth path, tangent to which are the Hamiltonian vector fields X_{φ_a} .

The notion of ‘algebra of gauge transformations’ is misleading in the Hamiltonian formalism. It is indeed originating from the Lagrangian formalism, in which the gauge transformations take a different form (this is the content of Chapter 3 in [Henneaux and Teitelboim, 1992] and most of Section 2 in [Henneaux, 1990]). There, one defines a gauge transformation δ_θ with infinitesimal parameter θ as a infinitesimal transformation leaving the action invariant: $\delta_\theta S = \delta_\theta y_i \frac{\delta S}{\delta y^i} = 0$, where δ in the fraction indicates functional derivation. Usually such a transformation $\delta_\theta y_i$ can be decomposed into a ‘true’ gauge transformation, and a ‘trivial’ one:

$$\delta_{\theta=\epsilon+\mu} y^i = R_a^i \epsilon^a + \mu^{ij} \frac{\delta S}{\delta y^j}$$

where the notation of the first term on the right-hand side contains more than what is written, and is taken from subsection 3.1.3 in [Henneaux and Teitelboim, 1992], while $\mu^{ij} = -\mu^{ji}$ is any kind of fully antisymmetric smooth function. Trivial gauge transformations of the form $\delta_\mu y^i$ are not gauge transformations, as they are artefactual and emerge from mere indices symmetries:

$$\delta_\mu S = \mu^{ij} \frac{\delta S}{\delta y^j} \frac{\delta S}{\delta y^i} = 0$$

The identity is automatically trivial because μ^{ij} is fully antisymmetric on the indices while $\frac{\delta S}{\delta y^j} \frac{\delta S}{\delta y^i}$ is fully symmetric. Every physical theory admit such trivial transformations, even if they are not gauge theories.

The Lie bracket of two gauge transformations in the Lagrangian formalism is still a gauge transformation, so they form an infinite dimensional Lie algebra \mathcal{F} . Computations show that the subspace \mathcal{T} of trivial gauge transformations forms a Lie ideal in the Lie algebra \mathcal{F} . The subspace $\overline{\mathcal{T}}$ of non-trivial gauge transformations is a complement of \mathcal{T} in \mathcal{F} , but may not form a Lie subalgebra of \mathcal{F} , as their bracket may involve a trivial one (Equation (3.17) in [Henneaux and Teitelboim, 1992]):

$$R_a^j \frac{\delta R_b^i}{\delta y^j} - R_b^j \frac{\delta R_a^i}{\delta y^j} = C_{ab}^\gamma R_c^i + M_{ab}^{ik} \frac{\delta S}{\delta y^k}$$

for some smooth function $M_{ab}^{ik} = -M_{ab}^{ki}$. Thus, we may not have a semi-direct sum $\mathcal{F} = \mathcal{T} \oplus_S \overline{\mathcal{T}}$ of Lie algebras as we have the following identities:

$$[\mathcal{T}, \mathcal{T}] \subset \mathcal{T}, \quad [\mathcal{T}, \overline{\mathcal{T}}] \subset \mathcal{T}, \quad [\overline{\mathcal{T}}, \overline{\mathcal{T}}] \subset \mathcal{F}$$

One can always choose another set of generators for the non-trivial gauge transformations, which is then equivalent to choosing another complement to \mathcal{T} in \mathcal{F} , which might, this turn, be a proper Lie subalgebra. The semi-direct sum is then dependent on the choice of gauge generators, as not every complement of \mathcal{T} defines a semi-direct sum.

Definition 4.67. *We call the subspace $\overline{\mathcal{T}} \subset \mathcal{F}$ the algebra of (non-trivial) gauge transformations; it is said to be closed if*

$$[\overline{\mathcal{T}}, \overline{\mathcal{T}}] \subset \overline{\mathcal{T}}$$

and open otherwise. It is said rigid if the structure functions are constant, and soft otherwise.

Remark 4.68. For historical reasons, we say that $\overline{\mathcal{T}}$ is an ‘algebra’ even if it may not be a Lie algebra *per se*.

Now, we would like to know what is the translation of this statement in the Hamiltonian formalism, where we work with a Poisson bracket on the phase space. Gauge transformations in the Hamiltonian formalism are more general than in the Lagrangian formalism, as can be seen from the fact that the extended Hamiltonian (Definition 4.60) contains *more* informations that

the original total Hamiltonian, and thus, the Lagrangian. In particular, trivial gauge transformations in the Lagrangian formalism do not possess associated constraints in the Hamiltonian formalism (see Section 2 in [Henneaux, 1990] for a good discussion), while gauge transformations in the Hamiltonian formalism have a geometrical meaning: they correspond to *canonical transformations* (subsection 3.2.4 in [Henneaux and Teitelboim, 1992]). The algebra of gauge transformations $\overline{\mathcal{T}}$ is not straightforwardly apparent in the Hamiltonian formalism as the Lagrangian gauge transformations and the Hamiltonian ones only coincide on the constraint surface. What occurs outside is specific to the Hamiltonian formalism. The condition of closedness or openness of the algebra $\overline{\mathcal{T}}$ then translates to a condition on the (first-class) constraints, as they are the generators of the gauge transformations. A result by Pons and Gracia establishes a sufficient condition for the algebra of gauge transformations to be closed: this happens if at least all constraints but one are (*at most*) *linear in the momenta* (see Section V in [García and Pons, 2000]). For geometrical reasons that will be clear later, we decide to take a slight restriction of this condition:

Definition 4.69. *In the Hamiltonian formalism, we say that the algebra of gauge transformations is closed if the first-class constraints are linear in the momenta; it is said open otherwise.*

Remark 4.70. Henneaux and Teitelboim seem to claim that in the Hamiltonian formalism the algebra of gauge transformations is closed if and only if the structure functions are *constant* (subsection 3.2.5 in [Henneaux and Teitelboim, 1992]), i.e. if and only if it is rigid. However, this seems a bit contradictory with their statement in subsection 3.1.8 where they implicitly say that not every closed algebra of transformations is a (finite dimensional) Lie algebra.

Let us deduce some geometrical consequences of the closedness of the algebra of gauge transformations. Every (first-class) constraint φ_a is of the following form in local coordinates:

$$\varphi_a = -\rho_a^i p_i \quad (4.82)$$

where $\rho_a^i \in C^\infty(Q)$ does not depend on the momenta. Then, closure of the algebra of the first-class constraints (Equation (4.77)) implies that the structure functions C_{ab}^c neither depend on the momenta, because the left-hand side is linear in those, and the constraint on the right-hand side is linear as well (here we use Definition 4.69). In other words:

Lemma 4.71. *If the algebra of gauge transformation is closed (in the sense of Definition 4.69), then ρ_a^i and C_{ab}^c are smooth functions on Q .*

Remark 4.72. If the constraints are irreducible, the smooth functions ρ_a^i, C_{ab}^c are uniquely defined because we cannot add a term proportional to a constraint as was explained in Remark 4.65, but if the constraints are reducible then there is some ambiguity in their choice.

The following argument is taken from [Ikeda and Strobl, 2019]. The Jacobi identity of the Poisson bracket applied to the constraints then reads:

$$\begin{aligned} \{\varphi_a, \{\varphi_b, \varphi_c\}\} + \{\varphi_b, \{\varphi_c, \varphi_a\}\} + \{\varphi_c, \{\varphi_a, \varphi_b\}\} = \\ \left(\{\varphi_a, C_{bc}^d\} + C_{ae}^d C_{bc}^e + \{\varphi_c, C_{ab}^d\} + C_{ce}^d C_{ab}^e + \{\varphi_b, C_{ca}^d\} + C_{be}^d C_{ca}^e \right) \varphi_d \end{aligned}$$

Since the left-hand side is zero everywhere, as a function, the right-hand side is zero as well. Then, assuming in full generality that the constraints are possibly reducible, Definition 4.42 and the exhaustion property (5.81) of reducibility functions (see also subsection 10.2.1 in [Henneaux and Teitelboim, 1992]) imply that the sum of terms on the parenthesis on the right-hand side is of the following form:

$$\{\varphi_a, C_{bc}^d\} + C_{ae}^d C_{bc}^e + \{\varphi_c, C_{ab}^d\} + C_{ce}^d C_{ab}^e + \{\varphi_b, C_{ca}^d\} + C_{be}^d C_{ca}^e = \tau_{abc}^I Z_I^d + \sigma_{abc}^{de} \varphi_e \quad (4.83)$$

for some smooth function $\tau_{abc}^I, \sigma_{abc}^{de}$ fully antisymmetric on the lower indices, and for the latter, on the upper indices also: $\sigma_{abc}^{de} = -\sigma_{abc}^{ed}$. The functions Z_I^d are the reducibility functions appearing in reducible systems of constraints. But since the structure functions C_{ab}^c do not depend on the momenta, while the term $\sigma_{abc}^{de}\varphi_e$ does, it means that the latter is identically zero (even for reducible systems). Even more precisely, it means that $\sigma_{abc}^{de} = 0$ because if it were not, it would be of the form $\rho_{abc}^{def}\varphi_f$ with antisymmetry on the indices e, f but this is still containing one momenta while the left-hand side of Equation (4.83) does not. Moreover, for the same reason, for reducible constraints, neither τ_{abc}^I nor Z_I^d depend on the momenta so they are smooth functions on Q .

Lemma 4.73. *If the algebra of gauge transformation is closed (in the sense of Definition 4.69), then*

$$\sigma_{abc}^{de} = 0 \quad (4.84)$$

and thus, in full generality :

$$\{\varphi_a, C_{bc}^d\} + C_{ae}^d C_{bc}^e + \{\varphi_c, C_{ab}^d\} + C_{ce}^d C_{ab}^e + \{\varphi_b, C_{ca}^d\} + C_{be}^d C_{ca}^e = \tau_{abc}^I Z_I^d \quad (4.85)$$

where $\tau_{abc}^I, Z_I^d \in C^\infty(Q)$, and $Z_I^d = 0$ for irreducible constraints.

Remark 4.74. The tensor σ_{abc}^{de} is the *higher order structure function* encoding the algebra of constraints. For more details, see subsection 3.2.5 in [Henneaux and Teitelboim, 1992] and [Browning and McMullan, 1987].

In order to make sense of Lemma 4.73, one needs to introduce the following geometric structure:

Definition 4.75. *Let M be a smooth manifold. An almost Lie algebroid over M is a smooth vector bundle E , together with:*

1. *a skew-symmetric bracket $[\cdot, \cdot]_E : \Gamma(E) \wedge \Gamma(E) \longrightarrow \Gamma(E)$ on the space of sections,*
2. *and a vector bundle morphism $\rho : E \longrightarrow TM$ called the anchor,*

such that the Leibniz rule holds:

$$[a, fb]_E = f[a, b]_E + \rho(a)(f) b \quad (4.86)$$

together with the morphism property:

$$\rho([a, b]_E) = [\rho(a), \rho(b)] \quad (4.87)$$

for every $a, b \in \Gamma(E)$, and $f \in C^\infty(M)$. An almost Lie algebroid is said regular if the anchor map has constant rank.

Remark 4.76. An almost Lie algebroid justifies its name because the bracket is *not* a Lie bracket, i.e. it may not satisfy the Jacobi identity, contrary to Lie algebroids. That is why we need to assume the morphism property (4.87), while in the Lie algebroid case, it was a consequence of the Leibniz rule (see Exercise 2.27).

We introduced the notion of almost Lie algebroid as it can very efficiently encode the constraints. We have the following extremely nice and useful result:

Proposition 4.77. *Let $\varphi_1, \dots, \varphi_p$ be a system of (possibly reducible) globally defined first-class constraints defining a closed algebra in the sense of Definition 4.69. There exists a unique regular almost Lie algebroid structure on the trivial vector bundle $E = Q \times \mathbb{R}^p$ over Q satisfying the following identities:*

$$\rho(e_a) = \rho_a^i \frac{\partial}{\partial q^i} \quad \text{and} \quad [e_a, e_b]_E = C_{ab}^c e_c$$

where e_1, \dots, e_p is the canonical global frame on E .

Proof. From Lemma 4.71 and 4.73, we know that all the structure functions are smooth functions on Q . The Leibniz rule (4.86) is enforced, and using it on the bracket $[e_a, e_b]$, one deduces that:

$$\left[\rho_a^i \frac{\partial}{\partial q^i}, \rho_b^j \frac{\partial}{\partial q^j} \right]_E = \left(\rho_a^i \frac{\partial \rho_b^k}{\partial q^i} - \rho_b^j \frac{\partial \rho_a^k}{\partial q^j} \right) \frac{\partial}{\partial q^k} \quad (4.88)$$

The coefficient of the right-hand side appears in the formula of the Poisson bracket $\{\varphi_a, \varphi_b\}$, as we have (using the definition of the constraints in Equation (4.82)):

$$\{\varphi_a, \varphi_b\} = - \left(\rho_a^i \frac{\partial \rho_b^k}{\partial q^i} - \rho_b^j \frac{\partial \rho_a^k}{\partial q^j} \right) p_k$$

The left-hand side being equal to $C_{ab}^c \varphi_c$, we deduce that:

$$C_{ab}^c (-\rho_c^k p_k) = - \left(\rho_a^i \frac{\partial \rho_b^k}{\partial q^i} - \rho_b^j \frac{\partial \rho_a^k}{\partial q^j} \right) p_k$$

implying that :

$$\rho_a^i \frac{\partial \rho_b^k}{\partial q^i} - \rho_b^j \frac{\partial \rho_a^k}{\partial q^j} = C_{ab}^c \rho_c^k \quad (4.89)$$

This shows that Equation (4.88) is equal to $\rho(C_{ab}^c e_c)$, which is Equation (4.86) for e_a, e_b .

Regularity of the almost Lie algebroid structure is a consequence of the fact that locally, the set of (possibly reducible) constraints is generated by a subset of irreducible ones, always the same number. This requires that locally, the image of the anchor map has constant rank. \square

Remark 4.78. The almost Lie algebroid E being regular, it implies that the image of the anchor map is a vector subbundle of TQ . It is involutive by Equation (4.87), hence integrable by Frobenius theorem 2.68. So every system of constraints – be they reducible or irreducible – which are linear in momenta and satisfying the regularity condition 4.37, defines a *regular* foliation on Q . The constraint surface Σ then coincides with the annihilator bundle $F^\circ \subset T^*Q$. Conversely, a singular foliation on Q , in the sense of [Laurent-Gengoux et al., 2020]: a (locally) finitely generated involutive subsheaf of the tangent sheaf of Q , induce constraints on T^*Q which are linear in momenta (by the identification between vector fields on Q and linear functions on T^*Q), but which do *not* satisfy the regularity condition 4.37. More precisely, the singular foliation $\mathcal{F} \subset \mathfrak{X}(Q)$ gives rise to a Lie subalgebra $\mathcal{I}_{\mathcal{F}} \subset \mathcal{C}^\infty(T^*Q)$ generated by functions which are linear in the momenta. Then, because of the singularities of the foliation, the zero level set of $\mathcal{I}_{\mathcal{F}}$ does *not* define an embedded submanifold of T^*Q , and a fortiori not an annihilator bundle as it was the case for regular foliations.

Notice that the reducibility condition on the constraints appears only when one computes the Jacobiator of the bracket $[\cdot, \cdot]_E$, which is not necessarily Lie. Then, Equation (4.85) implies that we have:

$$[e_a, [e_b, e_c]] + [e_b, [e_c, e_a]] + [e_c, [e_a, e_b]] = \tau_{abc}^I Z_I^d e_d \quad (4.90)$$

This hints to the existence of a higher Lie algebroid structure. Let A_1 be the number of reducibility equations (4.60), and define E' to be a trivial rank A_1 vector bundle over Q , equipped with a global frame $\{e_I\}_{1 \leq I \leq A_1}$. Then τ_{abc}^I (resp. Z_I^d) would be interpreted as the coefficients of a 3-bracket: $\tau : \wedge^2 E \rightarrow E'$, $\tau(e_a, e_b, e_c) = \tau_{abc}^I e_I$ (resp. a linear map $Z : E' \rightarrow E$, $Z(e_I) = Z_I^d e_d$). Then Equation (4.91) becomes:

$$[e_a, [e_b, e_c]] + [e_b, [e_c, e_a]] + [e_c, [e_a, e_b]] = Z(\tau(e_a, e_b, e_c)) \quad (4.91)$$

This equation is typical of higher Lie algebroid theory, where the (graded) Jacobi identity is only satisfied *up to homotopy*. This observation, together with Remark 4.78, shows that E encodes the regular foliation $F = \rho(E)$ in the following way:

Proposition 4.79. *Let E be the regular almost Lie algebroid associated by Proposition 4.77 to a system of (possibly reducible) globally defined first-class constraints $\varphi_1, \dots, \varphi_p$ defining a closed algebra in the sense of Definition 4.69. Let $F = \rho(E)$ be the regular foliation – regular involutive subbundle of TQ – induced by the anchor map. Then:*

1. *if the constraints are irreducible, E is a foliation Lie algebroid of F ;*
2. *if the constraints are reducible, E is the first term of a universal Lie ∞ -algebroid of F (see [Laurent-Gengoux et al., 2020] for a definition).*

Remark 4.80. Notice that in the discussion there are two kinds of regular foliations at hand. One on Q , defined when the constraints are linear in the momenta and one on the constraint surface $\Sigma \subset T^*Q$, defined for every constrained system via Proposition 4.64. The former is the restriction of the latter to the zero section $Z = \{(q, 0) \in T^*Q\} \simeq Q$ of the cotangent bundle, as $Z \subset \Sigma$.

Remark 4.81. When we have a closed algebra, the fiber of the almost Lie algebroid $E \rightarrow Q$ associated to the constraints via Proposition 4.77 can be used to define the vector bundle in which the gauge parameters live (see the discussion around Equation (4.81)), but the anchor map $\rho : E \rightarrow TQ$ is *not* the bundle map $\delta : E \rightarrow T(T^*Q)$. In particular, the latter map is *not* a Lie algebra morphism, since $\delta_{[e_a, e_b]} = \delta_{C_{ab}^c e_c} = C_{ab}^c \delta_{e_c}$ while, as can be read from Equation (4.79), $[X_{\varphi_a}, X_{\varphi_b}] = C_{ab}^c X_{\varphi_c} + \varphi_c X_{C_{ab}^c}$, so we have $\delta_{[e_a, e_b]} \neq [\delta_{e_a}, \delta_{e_b}]$.

Since the tangent spaces to the leaves of the foliation integrating the distribution D on Σ are generated by the hamiltonian vector fields X_{φ_a} associated to the first-class constraint φ_a , we deduce that the leaves of the regular foliation characterize gauge equivalent physical states. In other words, two points on the same leaf – this is a geometric equivalence relation – are ‘gauge equivalent’ in the sense that any physical observable $O \in \mathcal{C}^\infty(\Sigma)$ should take the same value in these two points. Physical observables are said to be the *gauge invariant functions on Σ* and, in the geometric picture, they correspond to those functions being constant *along* each leaf of the foliation. Then they are invariant with respect to the vector fields tangent to the leaves:

$$X_{\varphi_a}(O) = 0$$

for every first-class constraints φ_a , both primary *and* secondary. In the mathematical literature, functions which are invariant under the flow of the vector fields of a regular foliation are called *basic functions* [Moerdijk and Mrcun, 2003].

The gauge invariant functions on Σ are then constant along the leaves of the foliation induced by the vector fields X_{φ_a} . As we assume that the leaf space Σ_{ph} is a smooth manifold, we deduce that the gauge invariant functions on the constraint surface Σ pass to the quotient $\Sigma \twoheadrightarrow \Sigma_{ph}$ and define smooth functions on Σ_{ph} . Conversely, any smooth function on Σ_{ph} – the

putative physical observables – induce a smooth function on Σ which is constant along the leaves of the foliation, i.e. gauge invariant. Then we conclude that, as expected, there is a one to one correspondence between physical observables (on Σ_{ph}) and gauge invariant smooth functions (on Σ). Moreover, the properties of Poisson reduction, and in particular the fact that Proposition 3.97 applies to the current situation, shows that taking the Dirac bracket of any (global) extensions of two gauge invariant functions on Σ is equivalent to taking the reduced Poisson bracket of the corresponding physical observables on Σ_{ph} .

Definition 4.82. *By abuse of denomination, we will call gauge-invariant function or observable any smooth function $f \in \mathcal{C}^\infty(T^*Q)$ (or possibly only on a tubular neighborhood W of Σ) such that:*

$$X_{\varphi_a}(f) \approx 0 \tag{4.92}$$

for every first class constraint φ_a . Moreover, if $O \in \mathcal{C}^\infty(\Sigma)$ is constant along the gauge orbits, then any smooth function $f \in \mathcal{C}^\infty(T^*Q)$ restricting to O on Σ is a gauge-invariant function called a gauge-invariant extension of O .

Any such gauge invariant function, when restricted to the constraint surface Σ , is a proper gauge invariant function, i.e. a physical observable $O = f|_\Sigma \in \mathcal{C}^\infty(\Sigma)$, and any such latter function admits a (possibly global) extension satisfying the assumption Definition 4.82. We would now characterize the space of physical observables by using such extended notion of gauge invariant function. It is indeed easier to work on T^*Q as we do not work on a quotient, as would be the case if one worked with $\mathcal{C}^\infty(\Sigma)$. However, notice that while there was a one-to-one correspondence between physical observables – i.e. smooth functions on Σ_{ph} – and gauge invariant functions on Σ , there are much more gauge-invariant functions $f \in \mathcal{C}^\infty(T^*Q)$ as characterized in Definition 4.82. The latter are however quite useful to characterize the space of function $\mathcal{C}^\infty(\Sigma_{ph})$ because they are defined over the whole phase space T^*Q , which has a more regular smooth structure than the constraint surface Σ .

Condition (4.92) equivalently means that $X_f(\mathcal{I}_\Sigma) \approx 0$, where here the Hamiltonian vector field is computed with respect to the Dirac bracket (so that X_f vanishes anyway on second-class constraints). Since the constraints span the ideal \mathcal{I}_Σ of functions vanishing on Σ and since Σ is a (closed) embedded submanifold of T^*Q , by Lemma 2.58 it implies that the hamiltonian vector field of f (with respect to the Dirac bracket) is tangent to Σ . Hence, the smooth functions $f \in \mathcal{C}^\infty(T^*Q)$ inducing gauge invariant functions on Σ or, equivalently, physical observables, are precisely those smooth functions $f \in \mathcal{C}^\infty(T^*Q)$ such that:

$$\{f, \mathcal{I}_\Sigma\}_{Dirac} \approx 0 \tag{4.93}$$

These are precisely the gauge-invariant functions of Definition 4.82, because on the one-hand second-class constraints are Casimir elements of the Dirac bracket, and on the other hand first-class constraints are such that $\{\varphi_a, \cdot\}|_\Sigma = \{\varphi_a, \cdot\}_{Dirac}|_\Sigma$.

The maximal subalgebra of $(\mathcal{C}^\infty(T^*Q), \{\cdot, \cdot\}_{Dirac})$ generated by such functions is denoted \mathcal{N}_Σ . It contains \mathcal{I}_Σ because, although it is generated by both first-class and second-class constraints, the second-class constraints are Casimir elements of the Dirac bracket, while the bracket of \mathcal{I}_Σ with any first-class constraint vanishes on Σ . By construction, we can identify the physical observables $\mathcal{C}^\infty(\Sigma_{ph})$ with the following quotient:

$$\mathcal{C}^\infty(\Sigma_{ph}) \simeq \mathcal{N}_\Sigma / \mathcal{I}_\Sigma$$

By Equation (4.93), \mathcal{I}_Σ is a Lie ideal of $(\mathcal{N}_\Sigma, \{\cdot, \cdot\}_{Dirac})$, so this quotient inherits a canonical Lie algebra structure induced from the Dirac bracket. An alternative view on this approach,

using the original canonical Poisson bracket and not the Dirac bracket, hence replacing Condition (4.93) by (4.92), is given in [Kimura, 1993]. It has the benefits of not passing through the Dirac bracket and thus circumvent the difficulty of handling second-class constraints.

Another characterization – relying on Poisson reduction – is based on the following argument: by Lemma 4.62, we know that Σ is a coisotropic submanifold of the second-class constraint manifold $(\Sigma_0, \{ \cdot, \cdot \}_{\Sigma_0})$, where $\{ \cdot, \cdot \}_{\Sigma_0}$ is the restriction of the Dirac bracket to Σ_0 (see Proposition 3.83). The gauge orbits on Σ form a regular foliation generated by the Hamiltonian vector fields associated to the first-class constraints, and that the leaf space is the reduced phase space Σ_{ph} and is supposed to be a smooth manifold. Then, by Proposition 3.97 we can proceed to a Poisson reduction from $(\Sigma_0, \{ \cdot, \cdot \}_{\Sigma_0})$ and define a Poisson bracket on Σ_{ph} , that we denote $\{ \cdot, \cdot \}_{ph}$. We deduce that $(\mathcal{C}^\infty(\Sigma_{ph}), \{ \cdot, \cdot \}_{ph})$ is a Poisson algebra, and $(\Sigma_{ph}, \{ \cdot, \cdot \}_{ph})$ is a Poisson manifold (in fact symplectic). The relationship with the previous argument is that the Lie algebras $(\mathcal{C}^\infty(\Sigma_{ph}), \{ \cdot, \cdot \}_{ph})$ and $(\mathcal{N}_\Sigma / \mathcal{I}_\Sigma, \{ \cdot, \cdot \}_{Dirac})$ are isomorphic (where in the latter the bracket is the one induced from the Dirac bracket).

The reduced phase space Σ_{ph} corresponds, in constrained Hamiltonian systems, to the classical unconstrained picture in Hamiltonian mechanics: each and every point correspond to a different physical state and the Hamilton equations of motion take the usual form on it (i.e. as in the unconstrained formalism). Although the symplectic structure on Σ_{ph} is open to quantization, it is difficult to directly quantize the theory on the reduced phase space. Indeed, in practice, the explicit construction of the reduced phase space in terms of the given canonical variables may not be possible, since it requires solving the equations of motion for arbitrary initial data. The Dirac bracket, being constructed from inverting a matrix, has also certainly a quite complicated expression. Then, it is often advisable to work in the full $2n$ -dimensional phase space T^*Q with the original and more flexible form of Poisson bracket. This gives us the freedom to think of functions defined over all of phase space, compute their partial derivatives, for example, with respect to each of the q 's and p 's, and delay to the end the restriction of the variables to the constraint surface Σ , or even to the reduced phase space Σ_{ph} . Dirac achieved this strategy by quantizing the canonical coordinates on the total phase space T^*Q – which is a well known procedure – together with the constraints, by promoting them operators. This is necessary to keep track of which function of q and p are physical observables or not. However, even there some issues arise, when the constraints are for example rather complicated to express with the position and momentum operators (due to ordering for example). Expressing and quantizing the gauge invariant functions, i.e. the observables, could then become tedious. See Chapter 5 for a discussion about quantization of constrained systems and how to circumvent the difficulties.

Remark 4.83. There is an alternative procedure leading to the same reduced phase space, established by Faddeev and Jackiw in the late 80's [Faddeev and Jackiw, 1988] (or also Section 4.4 in [Rothe and Rothe, 2010]). It has been shown since then that this approach is equivalent to the Bergmann-Dirac algorithm [García and Pons, 1997].

The reduced phase space Σ_{ph} is obtained as a quotient of the secondary constraint surface Σ : it corresponds to the leaf space of the foliation on Σ induced by the gauge transformations – equivalently, the hamiltonian vector fields of the first-class constraints. In theory it does not live inside T^*Q then. However, it could be convenient to have a representent of Σ_{ph} inside Σ . Indeed, this would allow to use the usual variables q^i and p_i or those attached to Σ if they are well-adapted to the situation. In our case, we assumed that the leaf space is a smooth manifold (this is rarely the case) and in that case, if $\dim(\Sigma) = m$ and the dimension of the gauge leaves is p (so it is also the number of independent first-class constraints φ_a), we can expect that there exists – at least locally – an embedded $m - p$ dimensional submanifold N of Σ which is *transverse* to the gauge orbits: it intersects each of them in only one point and at that point x , we have

$T_x M = T_x N \oplus T_x L_x$, where L_x is the gauge orbit (leaf) to which x belongs. Then, by counting dimensions one realizes that this so-called *transversal* is diffeomorphic to the leaf space. Notice that there can be an infinite number of transversals to the leaves (because), and each of them could be considered as a legitimate representent of the leaf space.




Figure 23: On this representation of the secondary constraint surface Σ , the gauge orbits (represented by the plain lines) form the leaves of a regular foliation, generated by the hamiltonian vector fields associated to the first-class constraints. The dashed lines represent two possible transversals: N_1 and N_2 – they are defined by different gauge fixing conditions. They cut each leaf at only one point, and they satisfy $T_x M = T_x N_i \oplus T_x L_x$, where L_x is the leaf through x . Under sufficiently mild assumptions, any two transversals to the leaves are diffeomorphic, and they are diffeomorphic to the reduced phase space Σ_{ph} .

Usually, such a transversal N is materialized in the theory under the form of a zero-level set of p independent smooth functions C_b :

$$N = \bigcap_b C_b^{-1}(0)$$

The equations $C_b = 0$ characterizing N are called *gauge conditions* and the procedure of defining N by using such equations is called *gauge fixing*. For now on, we will assume that these functions are globally defined to facilitate our discussion (obstructions to this are called *global anomalies*, while gauge conditions that do not fix a gauge uniquely are called *Gribov ambiguities*). Due to the diffeomorphism between N and Σ_{ph} , there is a one-to-one correspondence between different attainable physical states of the system and points on N . As Henneaux and

Teitelboim point out, the gauge conditions $C_b = 0$ are therefore ad-hoc equations brought from the outside of the theory to avoid "multiple counting of states" (see Section 1.4 in [Henneaux and Teitelboim, 1992]). Since N is defined as the intersection of zero-level sets of p smooth functions C_b , Lemma 2.58 indicates that the tangent vectors to N should vanish on each C_b , when evaluated on N . Thus, since the choice of the functions C_b has been made so that it defines a transversal to the gauge foliation – i.e. such that $T_x N \cap TxL_x = 0$ for every gauge orbit – we deduce that no Hamiltonian vector field X_{φ_a} can be tangent to N (they span L_x and adjacent leaves). So for every $1 \leq a \leq p$, there exists at least one function C_b for which:

$$X_{\varphi_a}(C_b)|_N = \{\varphi_a, C_b\}_{\Sigma_0}|_N \neq 0$$

Alternatively, this result can be reformulated as saying that the matrix of coefficients $\{\varphi_a, C_b\}_{Dirac}$ is invertible on N . But this condition means that the set $\{\varphi_a, C_b\}$ forms a system of second-class constraints in the manifold $(\Sigma_0, \{.,.\}_{\Sigma_0})$. Thus, one can proceed to Poisson-Dirac reduction on N and restrict the Poisson bracket (in fact symplectic structure) $\{.,.\}_{\Sigma_0}$ to N and obtain there a Poisson (in fact symplectic) structure, denoted $\{.,.\}_N$. This structure would then be equivalent to the one on Σ_{ph} :

Proposition 4.84. *Given any transversal N of the gauge orbits on Σ obtained via gauge fixing, the symplectic manifold $(N, \{.,.\}_N)$ induced by Poisson-Dirac reduction from $(\Sigma_0, \{.,.\}_{\Sigma_0})$ is symplectomorphic to the symplectic structure defined on the reduced phase space Σ_{ph} via Poisson reduction from Σ .*

Hence, not only have we N representing Σ_{ph} as an embedded submanifold of Σ , but also as a Poisson (symplectic) submanifold. It might thus be sometimes more convenient to work on $(N, \{.,.\}_N)$ than on the reduced phase space $(\Sigma_{ph}, \{.,.\}_{\Sigma_{ph}})$. Then, we observe that the first-class constraints can be turned into second-class constraints under an appropriate choice of transversal, whose choice is based on the characteristics of the leaf space Σ_{ph} . So the first-class constraints φ_a are a consequence of the theory, while the gauge fixing functions C_b are brought from the outside, and can be arbitrarily chosen as soon as the transversal N that they define is diffeomorphic to Σ_{ph} . Notice as usual that here we assumed that both constraints and functions are globally defined but rigorously this discussion is only local: the transversals only exist locally in general. Indeed, in the more general case, the C_b might not be defined everywhere on Σ . In this context, the transversal N can be understood as a local section of the fibre bundle $\Sigma \rightarrow \Sigma_{ph}$ but it may not be a global section. Obstructions to the existence of global sections are called *global anomalies*. A global anomaly is different from the *Gribov ambiguity*, because in a global anomaly, there is no consistent definition of the gauge field, while a Gribov ambiguity is a lack of uniqueness of the determination of the physical state after gauge-fixing. A global anomaly is a barrier to defining a quantum gauge theory that was discovered by Witten in 1980.

Remark 4.85. Seeing Σ as a fiber bundle – not a *vector* one – over Σ_{ph} justifies the principal bundle perspective of gauge transformations, that was developed to manage Lie groups of symmetries as in Yang-Mills theory. Indeed, in that case $\Sigma \rightarrow \Sigma_{ph}$ is a principal bundle and the hamiltonian vector fields of the first-class constraints define a Lie algebra (the structure functions are in fact constant). See the book of Baez and Muniain [Baez and Muniain, 1994], together with the [lecture notes](#) of Figueroa-O'Farill for a thorough introduction to this material.

A final remark on this: we have shown that any set of first-class constraints can be turned into second-class constraints upon (non unique) gauge fixation. This leaves open the converse question: does any second-class system comes from gauge fixation of a purely first-class system? The answer is yes, in a non-unique way, see subsection 1.4.3 in [Henneaux and Teitelboim, 1992]. It implies in particular that, instead of proceeding to Poisson-Dirac reduction of the standard

Poisson bracket onto the second-class constraint submanifold Σ_0 , one can on the contrary drop half of the second-class constraints so that the remaining half is a first-class system. The second-class constraints indeed always come in pairs, so their total number is even. To perform this choice, it is sometimes necessary to enlarge the phase space by adding degrees of freedom. However, the price is low compared to that of proceeding a Poisson-Dirac reduction on Σ_0 . Indeed, we know that the Dirac bracket involves an inverse matrix and is polynomial in the constraints, hence it is very complicated and this would bring complications for quantizations. On the contrary, extending the phase space preserves the standard Poisson bracket (in fact the symplectic structure) and the Poisson reduction is much more amenable to perform on this extended phase space with standard bracket, that on the secondary constraint surface Σ . For more on the procedure of embedding a mixed system of first and second-class constraints into a purely first-class system by extending the phase space – this is called the *BFT procedure* – see Chapter 7 of [Rothe and Rothe, 2010].

5 BRST formalism and quantization of constrained systems

Inspired by the first developments of quantum mechanics, physicists proposed to find new quantum models by *quantizing* classical systems. What this process is about is still subject to hot debates, precisely because *quantization* appears to be more a heuristics than a proper and clearly stated mathematical method [Gotay et al., 1996]. For example, we observed that a unique classical system can be the limit of several quantum systems, hence drowning the hope of having a uniquely defined quantization scheme. In physics, there exist many such quantization procedures – canonical quantization, path integral formalism, stochastic quantization – the usefulness of each one of them being relative to the situation, and depends on the aim of the researcher. Historically, canonical quantization was the first to emerge, and is a good heuristics to find simple quantum systems from classical ones, such as e.g. the quantum harmonic oscillator. It has been popularized through the work of Dirac on constrained Hamiltonian systems, and as of now it has been mostly used to quantize field theories.

When it comes to quantization of a classical mechanical system however, there is a set of properties that most physicists agree on. A *quantization scheme* is a map $\mathcal{Q} : f \mapsto \mathcal{Q}_f$ sending physical observables to self-adjoint operators on a separable Hilbert space, which satisfies the following properties:

1. $\mathcal{Q}_{q^i} = q^i$. and $\mathcal{Q}_{p_i} = \frac{\hbar}{i} \frac{\partial}{\partial q^i}$ are the well-defined operators position and momenta,
2. the quantization scheme $f \mapsto \mathcal{Q}_f$ is \mathbb{R} -linear and satisfies $\mathcal{Q}(1) = \text{id}$,
3. the Poisson bracket-commutator correspondence, or *Dirac rule* $[\mathcal{Q}_f, \mathcal{Q}_g] = i\hbar \mathcal{Q}_{\{f,g\}}$,²¹ and
4. the *von Neumann rule* $\mathcal{Q}_{g \circ f} = g(\mathcal{Q}_f)$.

The latter rule is necessary for example when a physical observable f is written in terms of the coordinates functions q^i and p_i , such as $f(q^i, p_i)$. Then by the fourth item, the quantization of f would be $f(Q^i, P_i)$, where we do not know precisely how to interpret the order of the non-commuting operators $Q^i = \mathcal{Q}_{q^i}$ and $P_i = \mathcal{Q}_{p_i}$. From the end of the 1920's, we possess a prescription for quantizing polynomials of q^i and p_i – the *Weyl's transform* – representing its quantization as the totally symmetric ordering of the operators Q^i and P_i . Unfortunately, Groenewold and de Hove have shown that one cannot consistently quantize the Poisson algebra of all polynomials in the positions q^i and momenta p_i as symmetric operators on some Hilbert space [Gotay, 1999]. Thus it is in principle impossible to quantize every classical observable, or even every polynomial observable, in a way consistent with the Schrödinger picture. At most one can only consistently quantize certain Lie subalgebras of observables, for instance polynomials which are at most quadratic, or observables which are affine functions of the momenta. Even worse, it can be shown that not all four items – not even three of them! – can be satisfied at the same time [Ali and Engliš, 2005]. Notice also that there may be another set of rules to be chosen [Shewell, 1959, Gotay et al., 1996], illustrating the varieties of existing quantization schemes.

In light of this, canonical quantization should not be understood as a legitimate procedure to obtain quantum systems from classical systems, because quantum physics is not a consequence

²¹Under this natural correspondence, we have the canonical commutativity relation $[Q^j, P_k] = i\hbar \delta_k^j \text{Id}$, and the correct formula for the time evolution of operators in the Heisenberg picture: $\frac{d\mathcal{Q}_A}{dt} = \frac{i}{\hbar} [\mathcal{Q}_H, \mathcal{Q}_A]$. This observation provides a physical *a posteriori* justification to the choice of convention for the definition of Hamiltonian vector fields in Poisson geometry. Indeed, although the convention $X_H = \{H, \cdot\}$ has the unfortunate consequence that the flow of X_H is *minus* the flow of time (as can be seen from Hamilton's equations of motion), it eventually generates the well-known formulas of quantum mechanics under their most accepted form.

of classical physics. For example, physicists have been aware as early as the 1930s that there exist different quantum systems sharing the same classical limit. Then, canonical quantization should rather be seen as a recipe to get insights about the quantum theory, without putting too much emphasis on their mathematical correspondence: the primary role of the classical theory is not in approximating the quantum theory, but in providing a framework in its interpretation [Woodhouse, 1997]. The canonical quantization of non-singular physical systems (i.e. those for which the Legendre transform is invertible) is in general straightforward. In that case physical observables form (a subalgebra of) the algebra of smooth functions on the phase space T^*Q and the Hilbert space of quantum states is defined relatively to the original phase space and could in general be found by hand when it is obvious. Although promising, canonical quantization quickly runs into difficulties because it depends on the initial choice of coordinates and is not invariant under general canonical transformations; that is why it should only be taken as a probe of the quantum theory.


These problems about non-consistency of canonical quantization arise when the physical system is singular, i.e. when the Legendre transform is non-invertible because in this context the role of the coordinates is central. We have indeed seen that if the Lagrangian is singular – as is the case in gauge theories – constraints emerge in the Hamiltonian setup, and the physical observables consist of the algebra of smooth functions defined on the reduced phase space Σ_{ph} . In general the geometry of this submanifold can be quite complicated: it has a much more intricate structure than the secondary constraint surface Σ because the former is obtained as a quotient of the latter (as the leaf space of the foliation induced by the gauge transformations). Computing the algebra of physical observables $\mathcal{C}^\infty(\Sigma_{ph})$ is then not a trivial problem and is an huge obstacle to quantization, not even to mention that there might not be any "nice" choice of coordinates. However, quantizing the physical observables could be done through a little detour, by first identifying them as the gauge invariant functions on Σ , since we have a relatively transparent way of defining $\mathcal{C}^\infty(\Sigma)$ from the well-known algebra $\mathcal{C}^\infty(T^*Q)$, and then quantize the latter while keeping track somehow of the gauge invariance and the constraints.

An additional problem is that the induced Poisson bracket on the reduced phase space Σ_{ph} is not as simple as the standard one on T^*Q . This obstruction arises from both the definition of the Dirac bracket, and from Poisson reduction. For example, for a purely second class system of constraints, the definition of the Dirac bracket on $\Sigma_0 = \Sigma$ necessitates to invert a matrix and to multiply it by constraints. The terms on the right of Equation (3.42) are thus possibly highly complicated and do not facilitate quantization (see subsection 13.1.3 in [Henneaux and Teitelboim, 1992]). On the other hand, the Poisson structure of a purely first-class system of constraints relies on Poisson reduction, which is a quotient, hence inducing also a possibly complicated Poisson structure on $(\mathcal{C}^\infty(\Sigma_{ph}), \{.,.\}_{ph})$. However there exists an algebraic way of encoding the physical observables of a purely first-class system, which then allows to define a simpler Poisson structure (even symplectic) on a bigger phase space. This Poisson structure is in some sense equivalent to that on $(\mathcal{C}^\infty(\Sigma_{ph}), \{.,.\}_{ph})$ but has the enormous advantage of being standard. The process of constructing this complex and its associated Poisson structure from a purely first-class system of constraints is called the *BRST formalism*²².

This formalism relies on algebraic techniques which are much more insensitive to the geometric subtleties of the problem: instead of reducing the phase space in an intricate way to Σ and then to Σ_{ph} , the BRST formalism treats the constraints as what they are – generators of an algebra, admitting a resolution – and *extends* the phase space so that the geometric information

²²Notice that the BRST formalism corresponds to two different things in mathematics and in physics, although related. Its mathematical formulation – presented in this chapter – mostly refers to the Hamiltonian setup, while its physical one is most often understood in the Lagrangian setup (with Fadeev-Popov ghosts etc.). The correspondence between the two formalism is addressed in several places [Henneaux and Teitelboim, 1992, Nirov and Razumov, 1993, Rothe and Rothe, 2010].

contained in Σ and Σ_{ph} is exactly contained in the zero-th group of cohomology of this resolution. We have thus transported the complicated informations carried by Σ and Σ_{ph} from the geometric picture of the canonical hamiltonian formalism to the algebraic, more abstract but more linear picture of the BRST formalism. Then, the extended phase space developed in the BRST formalism can be equipped with a standard Poisson (even symplectic) structure, which is much more easy to quantize. Indeed, although the zero-th group of cohomology is defined as a quotient so in the end we still have this quotient procedure that is characteristic of the Poisson reduction on Σ_{ph} , it is in general much more transparent to use cohomological techniques than geometric reduction. As the following diagram shows, quantization of the theory after applying the BRST formalism should coincide with what would have been obtained from the canonical quantization perspective:



Eventually, there are two alternative quantization schemes provided by mathematicians: *geometric quantization* and *deformation quantization*. While the former is state-oriented and tries to provide a Hilbert space of quantum states, the second is operator-oriented and tries to deform the algebra of observables – i.e. gauge invariant smooth functions on the phase space – so that we obtain a non-commutative associative algebra resembling the operator algebra physicists look for. None of them give a definite answer to quantization because they have their own, respective, issues. Notice also that deformation quantization is the mathematization of an alternative interpretation of quantum mechanics, what physicists call the *phase space formulation of quantum mechanics* [Schroeck, 1996, Mariño, 2021]. This formulation of quantum mechanics places the position and momentum variables on equal footing in phase space, while the traditional Schrödinger picture uses either the position or the momentum representation. As a historical note, see [Curtright et al., 2013, Carosso, 2022] for historical surveys of quantization from the physics side.

5.1 Canonical quantization of a Hamiltonian system

Throughout this section, we will keep in mind that canonical quantization is more a heuristics than a rigorous procedure to find quantum physical models. We will also use the notation \widehat{f} to denote the image of a classical observable f under the quantization scheme $\mathcal{Q} : f \mapsto \mathcal{Q}_f$. Usually, such an operator is obtained by quantizing first the canonical coordinates q^i and p_i and then, based on the functional dependence of f on these variables, find its associated operator \widehat{f} . This process obviously implies to choose an ordering of the operators \widehat{q}^i and \widehat{p}_i , and this in general leads to major complications. For now on however, we will leave these details aside and concentrate on the general theory of canonical quantization.

First, assume that the Lagrangian is non-singular, meaning that the Legendre transform $\mathcal{L} : TQ \rightarrow T^*Q$ is bijective. In that case, there are no constraints, and no gauge transformations: there is a one-to-one correspondence between points of the phase space and physical states of the system. We can then use the Poisson bracket-commutator correspondence:

$$[\widehat{f}, \widehat{g}] = i\hbar \widehat{\{f, g\}} \quad (5.1)$$

in order to define the commutator of operators. The Poisson bracket here is the standard Poisson bracket on T^*Q , it is non degenerate and dual to the canonical symplectic form on T^*Q (see subsection B.3).

Equation (5.1) implies the well-known equation governing time-evolution of operators from quantum mechanics:

$$\frac{d\hat{A}}{dt} = \frac{i}{\hbar} [\hat{H}, \hat{A}] \quad (5.2)$$

where \hat{H} represents the Hamiltonian operator. Equation (5.2) tells us that canonical quantization gives a quantum model in the *Heisenberg picture*, where the operators are time dependent. The quantization scheme:

$$\mathcal{Q} : \text{classical observables} \longmapsto \text{quantum operators}$$

indeed provides us with a set of self-adjoint operators. The separable Hilbert space \mathcal{H} of quantum states would then correspond to an irreducible representation of this Lie algebra of operators, some insight can be obtained by studying the Lie algebra of classical observable $(\mathcal{C}^\infty(T^*Q), \{.,.\})$. In particular see this page explaining that the latter Lie algebra is the Lie algebra of the *quantomorphism group*, which is an important Lie group governing the corresponding quantum model. This discussion can also be useful.

Example 5.1. The quantization of the standard phase space $T^*\mathbb{R}^n$ would give the Hilbert space $\mathcal{H} = L^2(\mathbb{R}^n)$ of functions on \mathbb{R}^n whose norm is square integrable. Then the canonical quantization scheme is straightforward and the position and momenta operators are the usual ones. By using the Weyl's prescription for ordering of operators, we can quantize polynomials of the canonical variables q^i and p_i . For example:

$$\mathcal{Q} : q^i p_j \longmapsto \frac{\hat{q}^i \hat{p}_j + \hat{p}_j \hat{q}^i}{2}$$

However by Groenewold-Van Hove theorem [Groenewold, 1946, Lepage and Godeaux, 1951, Gotay, 1999], this prescription contradicts other axioms of canonical quantization, and there are no existing ordering prescription that would fit.

When the Lagrangian is singular, i.e. when there are constraints in the Hamiltonian formalism and that the physics is restricted to the constraint surface Σ , canonical quantization has to incorporate the constraints in some way or another. Interestingly, the way of quantizing first-class and second-class constraints are quite different. Depending on the situation, it may be advantageous to transform a system of second-class constraints into first-class constraints to quantize them, or the other way around. As a matter of references, see Chapter 9 of [Rothe and Rothe, 2010] and Chapter 13 of [Henneaux and Teitelboim, 1992] for a presentation of quantization of first-class and second-class constrained systems, as well as the original lectures of Dirac [Dirac, 1964]. The paper [Matschull, 1996] presents some more details on quantization of first-class constraints which are worth looking at. Section 8 in [Earman, 2003] quickly summarizes the four various alternatives that we have at hand when confronted with a system of first-class constraints: gauge-fixing, reduced phase space, Dirac quantization and BRST quantization.

Since in general, one usually proceed to the Poisson-Dirac reduction of the standard Poisson bracket to the Dirac bracket on the second-class constraint manifold Σ_0 before considering the first-class constraints (see Section 4.5), we will first look at the quantization procedure of second-class constraints. We will then study how to quantize a set of first-class constraints and discuss the properties of both procedures. Let assume that we have a pure second-class system, then the second-class constraint submanifold Σ_0 coincides with the secondary constraint surface Σ . The

Dirac bracket is a redefinition of the standard Poisson bracket of T^*Q and is defined on at least a tubular neighborhood of $\Sigma = \Sigma_0$, so that the latter is a symplectic leaf of the Dirac bracket. By Poisson-Dirac reduction, the Dirac bracket $\{.,.\}_{Dirac}$ restricts to a non-degenerate Poisson bracket $\{.,.\}_\Sigma$ on the second-class constraint submanifold Σ , which could then be considered as the ‘true’ physical phase space of the classical theory. In particular, there is a one-to-one correspondence between points on Σ and different physical states of the system; We could then quantize the theory as if it were an unconstrained hamiltonian system.

Ideally, one should find a set of coordinates (r^j, s_j) adapted to Σ ; there are an even number of them because Σ is a symplectic leaf of $(T^*Q, \{.,.\}_{Dirac})$. Because there are no gauge symmetries in this kind of theory, the classical observables are the smooth functions on Σ , which would then locally depend on the coordinates r^j, s_j . Then in the canonical quantization procedure these coordinates should replace the standard coordinates q^i, p_i and the Poisson bracket $\{.,.\}_\Sigma$ should replace the standard Poisson bracket in Equation (5.1):

$$\begin{aligned} \widehat{r}^j &= r^j. & \text{and} & & \widehat{s}_j &= \frac{\hbar}{i} \frac{\partial}{\partial r^j}, \\ [\widehat{f}, \widehat{g}] &= i\hbar \widehat{\{f, g\}}_\Sigma \end{aligned} \tag{5.3}$$

The expression of the operators \widehat{f}, \widehat{g} is based on their expression as functions of the variables r^j, s_j , and a choice of operator ordering. The Hilbert space of quantum states would then be found as an irreducible representation of the Lie algebra of commutators defined by Equation (5.3), which would be possibly related to irreducible representations of the Lie algebra $(\mathcal{C}^\infty(\Sigma), \{.,.\}_\Sigma)$, as this [explanation](#) shows.

There are two main problems arising when quantizing second-class constraints, see subsection 13.1.2 of [[Henneaux and Teitelboim, 1992](#)]. First, the fact that smooth functions on Σ – the classical observables in a pure second-class system – can be easily formally described by the quotient $\mathcal{C}^\infty(T^*Q)/\mathcal{I}_\Sigma$, while their explicit form in terms of the local coordinates r^j, s_j may admit a much more complicated description. At least, another choice of local coordinates give another quantized system which is equivalent to the former, see bottom of page 85 of [[Gitman and Tyutin, 1990](#)]. Second, it may be difficult to find an explicit representation of the Lie algebra of self-adjoint operators defined by Equation (5.3). This problem can be caused by the fact that the very definition of the Dirac bracket necessitates to invert the matrix $C_{kl} = \{\chi_k, \chi_l\}$ and then turn it into quantum operators. The result can be quite involved because of ordering problems, which sometimes generate additional term on the right-hand side of Equation (5.3) at order 2:

$$[\widehat{f}, \widehat{g}] = i\hbar \widehat{\{f, g\}}_\Sigma + O(\hbar^2)$$

This makes even more difficult to evaluate the irreducible representation of the Lie algebra of quantum operators, hence the Hilbert space of quantum states.

A solution that avoids the first problem is to work directly on the phase space T^*Q . Then the usual position and momenta operators of the canonical variables q^i, p_i are used, while Equation (5.1) becomes:

$$[\widehat{f}, \widehat{g}] = i\hbar \widehat{\{f, g\}}_{Dirac} \tag{5.4}$$

Since the second-class constraints χ_l are Casimir elements of the Dirac bracket, their associated quantum operators $\widehat{\chi}_l$ have vanishing commutator with any other operator obtained through quantization:

$$[\widehat{f}, \widehat{\chi}_l] = i\hbar \widehat{\{f, \chi_l\}}_{Dirac} = 0$$

The second-class constraints are then understood to generate a set of equations at the level of operators:

$$\widehat{\chi}_l = 0 \tag{5.5}$$

The interpretation of the set of Equations (5.5) is the following: as the self-adjoint operators $\widehat{\chi}_l$ are functional expressions of \widehat{q}^i and \widehat{p}_i , Equations (5.5) put some relationship – or constraint – between these operators \widehat{q}^i and \widehat{p}_i . This is how the second-class constraints could be enforced in the quantum world, and the two methods of quantization should give equivalent quantum systems. Although the latter strategy of quantization of second-class system seems to be more promising than the former one, we are still left with the necessity of defining the Dirac bracket, which implies to invert the matrix $C_{kl} = \{\chi_k, \chi_l\}$ and then turn it into quantum operators, which can be a costly operation. Moreover, the equivalence of this method with the reduced phase space quantization coincide when the gauge group is unimodular and seems to be always provable [Earman, 2003]; see also Reference [7] of [García and Pons, 1997].

Now let us turn our attention to pure first-class systems. Recall that in that case the constraint surface Σ is a coisotropic submanifold of Σ_0 , which in the present case is T^*Q . The first-class constraint generate gauge transformations and the reduced phase space Σ_{ph} is the quotient of Σ by the orbits of these gauge transformations. Poisson reduction then allows to define a Poisson bracket $\{f, g\}_{ph}$ – in fact a symplectic structure – on Σ_{ph} , which then become the ‘true’ phase space of the physical theory, where each point is in one-to-one correspondence with a different physical state. We can then proceed to quantization following the same lines as for unconstrained systems. The smooth functions on Σ_{ph} are the classical observables and quantizing them consists in finding an adapted set of coordinates on Σ_{ph} , and modifying Equation (5.1) as:

$$[\widehat{f}, \widehat{g}] = i\hbar\widehat{\{f, g\}_{ph}}$$

Obviously, this is often a much more difficult task than quantizing a second-class system. So sometimes, one uses the fact that Σ_{ph} is diffeomorphic to some adapted choice of transversal to the foliation of gauge orbits in Σ , obtained by gauge-fixing (see Section 4.5). Then, getting a number of gauge-fixing conditions transforms the first-class system of m constraints into a pure second-class systems of $2m$ constraints, and apply the above procedure. Quantizing this system as was done earlier is possibly much more attainable than quantizing the reduced phase space directly but is conditioned to the absence of global anomalies. These two procedures – reduced phase space and gauge fixing – a priori give the same quantum theory if the reduced phase space is a well-defined smooth manifold and if there are no global obstruction to gauge fixing, see [Earman, 2003].

The above strategy for quantizing first-class systems has the advantage of quantizing directly the true classical observables – i.e. the gauge invariant functions on Σ – but has the drawback that it breaks manifest gauge symmetries, which are sometimes a necessary feature of physical theories (see subsection 13.2.3 in [Henneaux and Teitelboim, 1992]). Dirac then proposed a way to circumvent it by quantizing the first-class constraints in such a way that they preserve their valuable property: being generators of gauge transformations. Let us assume that we work with a set of first-class constraints φ_i , and that we neither proceed to Poisson reduction $\Sigma \rightarrow \Sigma_{ph}$, nor will use gauge-fixing conditions. Rather we want to first quantize the classical model on T^*Q and then proceed to some reduction of the state space.

Quantizing the first-class system then begins by defining the Hilbert space \mathcal{H} obtained by quantizing $(T^*Q, \{\cdot, \cdot\})$. Then we say that an element $|\psi\rangle \in \mathcal{H}$ is an *admissible quantum state* if it obeys the following equation:

$$\widehat{\varphi}_i |\Psi\rangle = 0 \tag{5.6}$$

for every first-class constraint φ_i . More precisely, if \mathcal{H} is the Hilbert space associated to the symplectic manifold T^*Q , then the admissible quantum states form a subspace \mathcal{S} of \mathcal{H} called the *physical state space* [Matschull, 1996], defined as the null eigenspace of the constraint operators

$\widehat{\varphi}_i$. In other words, the physical state space is the following:

$$\mathcal{S} = \left\{ |\Psi\rangle \in \mathcal{H} \text{ such that } \widehat{\varphi}_i |\Psi\rangle = 0 \text{ for every } i \right\}$$

Condition (5.6) is considered as the way of implementing gauge invariance in the quantum world. There is an alternative, maybe more explicit way of seeing this interpretation, that is presented on pp. 20-21 of [Matschull, 1996].

Now, we know that a gauge-invariant function $f \in \mathcal{C}^\infty(\Sigma)$ by definition satisfies Equation (4.92). This should be translated at the level of operators, by the statement that *physical quantum observables map \mathcal{S} into \mathcal{S}* . Let us show this: let f be a gauge-invariant function on Σ – i.e. a classical observable – so by Equation (4.92) we know that in a pure first-class system the Poisson bracket of f with any first-class constraint φ_i vanishes weakly. Thus, it is strongly equal to a linear combination of constraints: $\{f, \varphi_i\} = f^j \varphi_j$. Now assume that we have found an operator representation \widehat{f} for f , and that the commutator of operators as given in Equation (5.4) satisfies:

$$[\widehat{f}, \widehat{\varphi}_j] = i\hbar \widehat{f}_j^k \circ \widehat{\varphi}_k \quad (5.7)$$

which is the direct quantization of the Poisson bracket $\{f, \varphi_j\} = f_j^k \varphi_k$. Then, for any admissible state $|\psi\rangle \in \mathcal{S}$, we have:

$$\widehat{\varphi}_j \circ \widehat{f} |\psi\rangle = \widehat{f} \circ \underbrace{\widehat{\varphi}_j |\psi\rangle}_{=0} - \underbrace{[\widehat{f}, \widehat{\varphi}_j] |\psi\rangle}_{=i\hbar \widehat{f}_j^k \circ \widehat{\varphi}_k |\psi\rangle = 0} = 0 \quad (5.8)$$

This proves that the state $\widehat{f} |\psi\rangle$ is admissible, i.e. it belongs to the subspace \mathcal{S} .

Remark 5.2. There is an alternative way of implementing the quantification of first-class constraints on the Hilbert space \mathcal{H} , without using Equation (5.6). This is done by identifying two states $|\Phi\rangle$ and $|\Psi\rangle$ if they differ by a linear combination of constraints:

$$|\Phi\rangle \sim |\Psi\rangle \quad \text{if} \quad |\Phi\rangle - |\Psi\rangle = \widehat{\varphi}_i |\Omega_i\rangle$$

for some states $|\Omega_i\rangle$. Then the quotient of \mathcal{H} by this equivalent relation is \mathcal{S} . See page 20 of [Matschull, 1996] and subsection 13.3.5 in [Henneaux and Teitelboim, 1992].

Notice however that the condition that quantum observables leave the subspace \mathcal{S} invariant is based on the assumption that Equation (5.7) holds. In some cases however, due to some choice of ordering, the right-hand side may contain some terms which are not proportional to the constraint operators:

$$[\widehat{f}, \widehat{\varphi}_j] = i\hbar \widehat{f}_j^k \circ \widehat{\varphi}_k + \hbar^2 \widehat{B}_j$$

Permuting the operators on the right-hand side could indeed generate a term at second order in Planck's constant. Condition (5.8) is then satisfied if $\widehat{B}_j |\psi\rangle = 0$ for every admissible state $|\psi\rangle$. More precisely, we shall reduce the subspace \mathcal{S} to the sub-subspace of states which have this property. This is an additional condition that one has to enforce to ensure that Dirac quantization of first-class constraints is sustainable. In particular this can occur for two important cases, when the other constraints and the Hamiltonian are involved.

It can indeed happen that if one wants the first-term on the right-hand side to be the one we want, some reordering of operators could be necessary and we have some terms of order 2 which emerge:

$$\begin{aligned} [\widehat{H}, \widehat{\varphi}_j] &= i\hbar \widehat{a}_j^k \circ \widehat{\varphi}_k + \hbar^2 \widehat{C}_j \\ [\widehat{\varphi}_l, \widehat{\varphi}_j] &= i\hbar \widehat{C}_{lj}^k \circ \widehat{\varphi}_k + \hbar^2 \widehat{D}_{lj} \end{aligned}$$

The first equation comes from the fact that the Hamiltonian is a first-class function so that the Poisson bracket $\{H_T, \varphi_j\}$ vanishes on Σ , while the structure functions in the second equation are characterized by the Lie algebroid structure of the first-class constraints: $\{\varphi_l, \varphi_j\} = C_{lj}^k \varphi_k$.

The operators \widehat{C}_j and \widehat{D}_{lj} are called *gauge anomalies*, because they break gauge-invariance by possibly sending admissible states outside the physical state space \mathcal{S} . Then, in both cases, for consistency of the quantization, one needs to enforce the two conditions:

$$\widehat{C}_j |\psi\rangle = 0 \quad \text{and} \quad \widehat{D}_{lj} |\psi\rangle = 0 \quad (5.9)$$

Then the admissible states would be those satisfying both Conditions (5.6) and (5.9). However it may happen that imposing the latter two conditions is too strong and we end up with a subspace of admissible states which is much smaller than \mathcal{S} , if not vanishing. This is an obstruction to Dirac quantization of first-class constraints, but it can be avoided when passing through BRST formalism. See subsection 13.3.2 in [Henneaux and Teitelboim, 1992] for a detailed discussion of this problem.

Example 5.3. One can draw on the canonical formalism of Maxwell's electromagnetism presented in Section B.4 to quantize the electromagnetic field. This is done in details in [Matschull, 1996], where the space of admissible quantum states is justified to be constructed as a Fock space. A complete and more advanced treatment of the electromagnetic field via the BRST and antifield formalisms can be found in Chapter 19 of [Henneaux and Teitelboim, 1992].

As a final remark, when we have a mixed system of first-class and second-class constraints, there are in general two ways of quantizing it: either by turning the system into a pure second-class system by adding gauge-fixing conditions (see e.g. Section 9.3 in [Rothe and Rothe, 2010]), or by turning it into a pure first-class system, by adding formal variables, and seeing the second class constraints as gauge-fixing conditions on this extended phase space (such a procedure is called the *BFT procedure* – see Chapter 7 of [Rothe and Rothe, 2010]). Both methods meet their own issues: in the former, it may not be possible to find globally defined gauge-fixing conditions because of the so-called Gribov obstructions (see subsection 1.4.1 and Appendix 2.A in [Henneaux and Teitelboim, 1992]), while the latter strategy may be a bit challenging. After one has reformulated the theory in terms of a pure first-class system, one can then avoid the problems of gauge anomalies discussed earlier by using the BRST formalism, which precisely aims at providing a unified framework to treat and quantize a pure first-class system.

5.2 First principles of the BRST formalism: the irreducible case

Suppose that the Poisson-Dirac reduction has been performed on Σ_0 , then we are left with the first-class constraints φ_i . Then the ideal \mathcal{I}_Σ , although generated by all the constraints (both first-class and second-class), is a Lie subalgebra of $(\mathcal{C}^\infty(\Sigma_0), \{.,.\}_{Dirac})$. Since the bracket of two first-class constraint is by definition vanishing on Σ , we deduce that the above issue is not met with such constraints. However we are confronted with the issue of quantizing the physical observables, equivalently the gauge-invariant functions on Σ . The reduced phase space Σ_{ph} is quite complicated to describe, as well as its induced Poisson bracket obtained through Poisson reduction from that on Σ . A better way to describe the space of functions on Σ_{ph} is to use the *BRST formalism*: by extending the phase space we allow more degree of freedom and the theory becomes easier to handle, but we recover the underlying physics through the cohomology of a particular differential on this space. The central idea of the BRST formalism is to substitute for the original local gauge symmetry a fermionic global symmetry acting on the extended phase space. That symmetry captures the original gauge invariance and leads to a simpler formulation of the theory. Then quantization may be easier on this extended phase space than on the

original reduced phase space Σ_{ph} . From the mathematical point of view, the (Hamiltonian) BRST formalism is an alternative to Symplectic reduction [Stasheff, 1991, Figueroa-O’Farrill and Kimura, 1991a, Kimura, 1993], and thus could have been introduced already in the XIXth century, had physicists then been interested in extending classical mechanics to Grassmann variables [Henneaux and Teitelboim, 1988, Henneaux and Teitelboim, 1992].

BRST formalism applies to purely first-class systems, so from now on we assume that our classical physical system only contains first-class constraints. In particular, the second-class constraint manifold Σ_0 coincides with T^*Q , and the Dirac bracket is the standard Poisson bracket on T^*Q . The secondary constraint surface Σ is the zero-level set of the first-class (and only) constraints φ_i , which are labelled from 1 to p say. We additionally suppose that they are globally defined and irreducible (see Definition 4.42). This has the consequence that Σ is a closed embedded submanifold in T^*Q , which by the proof of the second item of Lemma 3.72, implies that the smooth functions on Σ can be identified with the following quotient:

$$\mathcal{C}^\infty(\Sigma) \simeq \mathcal{C}^\infty(T^*Q) / \mathcal{I}_\Sigma$$

where \mathcal{I}_Σ is the free multiplicative ideal generated by the (first-class) constraints in $\mathcal{C}^\infty(T^*Q)$. We recognize a quotient. The idea behind BRST formalism is first to obtain this quotient as the zero-th (co)homology group of a particular differential δ acting on some abstract space. Then, gauge invariant functions, which form a subspace of this (co)homology group, would be obtained by an additional quotienting operation, using another differential d . Then, having the gauge-invariant functions on Σ is by definition equivalent to having the functions on Σ_{ph} , i.e. the physical observables. In other words, the BRST formalism involves two steps:

1. Restriction to the constraint surface Σ , using the differential δ ;
2. Implementation of the gauge invariance condition on Σ , using the differential d .

Luckily, mathematicians have already described the quotient 5.2 using chain complexes and (co)homology. This quotient can indeed be seen as the quotient of the ring $\mathcal{C}^\infty(T^*Q)$ by an ideal generated by p elements which form what we call a *regular sequence*. Then it is known that this kind of quotient admits a resolution, that we call *Koszul complex*:

$$0 \longrightarrow K_{-p} \xrightarrow{\delta} K_{-p+1} \xrightarrow{\delta} \dots \xrightarrow{\delta} K_{-1} \xrightarrow{\delta} \underbrace{\mathcal{C}^\infty(T^*Q)}_{K_0} \longrightarrow 0$$

The K_{-i} are negatively graded by convention, because we will introduce later some spaces which would be their dual, and positively graded. By resolution, we mean the following: it is a chain complex (see Definition 1.42) such that the homology of the differential δ is zero except possibly at level 0 (see Definition 1.53). We use the term *homology* and not *cohomology* because we will soon see that the map δ reduces the number of anticommuting variables, contrary to e.g. the de Rham differential (1.32) or the Poisson differential (3.18). We denote the homology groups with an index downstairs, to emphasize that we do not work with cohomology. Then, in the present case, saying that we have a resolution of the quotient $\mathcal{C}^\infty(T^*Q) / \mathcal{I}_\Sigma$ means that $H_{-i}(\delta) = 0$ for every $i \geq 1$, i.e. $\text{Im}(\delta : K_{-i-1} \rightarrow K_{-i}) = \text{Ker}(\delta : K_{-i} \rightarrow K_{-i+1})$, while $H_0(\delta) = \mathcal{C}^\infty(T^*Q) / \mathcal{I}_\Sigma$. We will show that such a resolution exists by giving an explicit expression of the K_{-i} ’s and the action of the map δ .

We know the mathematical expression of K_{-i} , it is the following:

$$K_{-i} = \mathcal{C}^\infty(T^*Q) \otimes \wedge^i \mathbb{R}^p$$

In other words, K_{-i} is the sheaf of sections of the trivial vector bundle over T^*Q whose fiber is $\wedge^i \mathbb{R}^p$. Let us denote by $\mathcal{P}_1, \dots, \mathcal{P}_p$ the standard basis of \mathbb{R}^p in the present context, and we call these variables *ghost momenta* and we assign to them an abstract degree of -1 , called the *ghost number*, and which will be discussed in Section 5.3 once we have introduced their dual entities: the *ghosts*. This abstract grading is denoted gh and we have:

$$\text{gh}(\mathcal{P}_i) = -1$$

The grading can be extended to any homogeneous alternating powers of ghost momenta by the following convention: the monomial of ghost momenta $\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_k} \in K_{-k}$ of k ghost momenta is considered to have ghost number $-k$ (the number of ghost momenta in the monomial):

$$\text{gh}(\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_k}) = -k$$

So we see that so far the (negative) ghost number measures the polynomial degree of alternating powers of ghost momenta.

Then, we define the map δ on the basis $\mathcal{P}_1, \dots, \mathcal{P}_p$ so that each element \mathcal{P}_i is sent to the first-class constraint φ_i :

$$\begin{aligned} \delta : K_{-1} &\longrightarrow \mathcal{C}^\infty(T^*Q) \\ \lambda^i \mathcal{P}_i &\longmapsto \lambda^i \varphi_i \end{aligned}$$

where the λ^i are smooth functions. The linear map δ is then extended to a $\mathcal{C}^\infty(T^*Q)$ -linear graded derivation on $K_\bullet = \bigoplus_{i=1}^p K_{-i}$ by the following formula:

$$\delta(f \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_k}) = \sum_{i=1}^k (-1)^{i-1} \varphi_{j_i} f \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_{i-1}} \wedge \mathcal{P}_{j_{i+1}} \wedge \dots \wedge \mathcal{P}_{j_k} \quad (5.10)$$

We say that the map δ has ghost number $+1$ because there are $k-1$ ghost momenta in each monomial on the right-hand side so the ghost number of the right-hand side is $-k+1$, while the ghost number of the monomial $\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_k}$ is $-k$.

We can quickly see that δ is a differential by applying it again on Equation (5.10), when $k \geq 2$:

$$\begin{aligned} \delta^2(f \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_k}) &= \sum_{i=1}^k \sum_{l < i} (-1)^{i-1} (-1)^{l-1} \varphi_{j_l} \varphi_{j_i} f \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \widehat{\mathcal{P}_{j_l}} \wedge \dots \wedge \widehat{\mathcal{P}_{j_i}} \wedge \dots \wedge \mathcal{P}_{j_k} \\ &+ \sum_{i=1}^k \sum_{l > i} (-1)^{i-1} (-1)^{l-2} \varphi_{j_l} \varphi_{j_i} f \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \widehat{\mathcal{P}_{j_i}} \wedge \dots \wedge \widehat{\mathcal{P}_{j_l}} \wedge \dots \wedge \mathcal{P}_{j_k} \end{aligned}$$

One can check that the second line is minus the first one, so $\delta^2 = 0$. This identity also holds at the lowest level, when $k = 1$, because the image of $\delta : K_{-1} \rightarrow K_0$ is by construction $\mathcal{I}_\Sigma \subset K_0$, and $\delta|_{K_0} = 0$. Then, the zero-th homology group is the quotient we desire:

$$H_0(\delta) = \mathcal{C}^\infty(T^*Q) / \mathcal{I}_\Sigma \simeq \mathcal{C}^\infty(\Sigma)$$

The fact that the homology of the chain complex (K_\bullet, δ) is zero except at the lowest level is stated in Theorem 9.1 and proven in Appendix 9.A of [Henneaux and Teitelboim, 1992]. So we have managed to express $\mathcal{C}^\infty(\Sigma)$ in terms of a homology group. The Koszul complex can be

understood as a ‘de-quotienting’ of $\mathcal{C}^\infty(T^*Q)/\mathcal{I}_\Sigma$, and in general it is much more easier to use such non singular structures, than working with the quotient directly.

What we want now is to describe the functions on Σ which are gauge invariant; we will denote them $\mathcal{C}^\infty(\Sigma)_{inv}$. Since gauge transformations are generated by the Hamiltonian vector fields associated to the (first-class) constraints, the gauge invariant functions satisfy Equation (4.92) and are moreover in one-to-one correspondence with the smooth functions on the reduced phase space Σ_{ph} . We will describe the gauge-invariant functions in cohomological terms – with respect to a differential d – so that the algebra of functions $\mathcal{C}^\infty(\Sigma_{ph})$ could be obtained by purely (co)homological methods. To each first-class constraint φ_i corresponds its hamiltonian vector field X_i (and denoted X_i in Section 4.5) and, dually, a differential one-form η^i on T^*Q . More precisely, since the constraints are irreducible (the reducible case may be treated slightly differently, see Section 5.4), no vector field X_i ever vanishes, otherwise they would be dependent. They define a rank p regular distribution on T^*Q – i.e. an assignment, for every $(q, p) \in T^*Q$, of a p dimensional subspace $D_{(q,p)}$ of $T_{(q,p)}T^*Q$. This distribution is involutive on the constraint surface Σ because $[X_i, X_j] = X_{\{\varphi_i, \varphi_j\}} = X_{C_{ij}^k \varphi_k} \approx C_{ij}^k X_k$ (see Proposition 4.64).

Then it admits a dual bundle D^* which is isomorphic to a rank p regular *codistribution*, i.e. the assignment, for every $(q, p) \in T^*Q$, of a subspace $D'_{(q,p)}$ of $T_{(q,p)}^*T^*Q$. The corresponding one-forms generating this codistribution are some nowhere vanishing differential one-forms η^i such that:

$$\eta^i(X_j) = \delta_j^i \quad (5.11)$$

The one-forms η^i are smooth sections of the codistribution D' , and since they are nowhere vanishing and independent at every point of the phase space T^*Q , D' can be identified with a subbundle of T^*T^*Q and the η^i are a frame for it. While the dual bundle D^* is canonical, this codistribution is not, although isomorphic to the former. Indeed, there is some ambiguity in the choice of one forms η^i because one can always add a component which sits in the annihilator bundle of D without changing Equation (5.11). Another choice of differential one-forms would define an alternative codistribution D' . However the number of free generators is always p , equal to the number of vector fields X_i , the codimension of the constraint surface in the irreducible case.

Remark 5.4. More generally one can elaborate on Equation (5.11) by defining *longitudinal vector fields* – the vector fields on T^*Q which are parallel to Σ – and subsequent *longitudinal differential forms*. While the longitudinal vector fields are *not* generated by the X_i – i.e. they form a bigger $\mathcal{C}^\infty(T^*Q)$ -module – the longitudinal differential forms are generated by the η^i . See Section 5.3 of [Henneaux and Teitelboim, 1992] for further details on first-class constraint surface.

We call the differential one-forms η^i dual to the X_i the *ghosts*²³. We attribute to them an arbitrary ghost number of $+1$, so that any monomial $\eta^{i_1} \wedge \dots \wedge \eta^{i_k} \in \Gamma(\wedge^k D')$ of k ghosts has degree $+k$:

$$\text{gh}(\eta^i) = 1 \quad \text{and} \quad \text{gh}(\eta^{i_1} \wedge \dots \wedge \eta^{i_k}) = k$$

As its name suggests then, the positive ghost number then measures the polynomial degree of alternating powers of ghosts. We can define a map $d : \wedge^\bullet D' \rightarrow \wedge^{\bullet+1} D'$ of ghost number $+1$ – and abusively called the *longitudinal differential* – by the following identities:

$$df = X_i(f) \eta^i \quad \text{for every } f \in \mathcal{C}^\infty(T^*Q) \quad (5.12)$$

$$d\eta^k = -\frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \quad (5.13)$$

²³See the introduction of [Henneaux and Teitelboim, 1988] for a short survey of why the ghosts were introduced and gained prominence in theoretical physics, both at the quantum and at the classical level.

where the smooth functions $C_{ij}^k \in \mathcal{C}^\infty(T^*Q)$ are the structure functions appearing in the formula $\{\varphi_i, \varphi_j\} = C_{ij}^k \varphi_k$. The map d is defined to satisfy the usual Leibniz rule:

$$d(f\eta^{j_1} \wedge \dots \wedge \eta^{j_k}) = df \wedge \eta^{j_1} \wedge \dots \wedge \eta^{j_k} + f d(\eta^{j_1} \wedge \dots \wedge \eta^{j_k}) \quad (5.14)$$

Be aware that in general we *do not* have $df(X) = X(f)$ for any random vector field, because here d is not the de Rham differential. These data – Equation (5.12) and (5.13) – *do not* make the map d a differential on the entire phase space, but only on the constraint surface, as the following discussion will show.

One recognizes in Equation (5.13) something reminiscent of the Cartan-Eilenberg formula characteristic of the differential of Lie algebras. Then we expect that the action of d^2 on η^k corresponds to some Jacobi identity. And indeed we observe that Equations (5.12) and (5.13) imply:

$$d^2\eta^k = -\frac{1}{2} \left(X_i(C_{mn}^k) + C_{ij}^k C_{mn}^j \right) \eta^i \wedge \eta^m \wedge \eta^n \quad (5.15)$$

where full antisymmetry on the i, m, n indices in the coefficient is brought by the product of the three ghosts. The right-hand side of Equation (5.15) does not automatically vanish on T^*Q but instead satisfies Equation (4.83), without reducibility functions and rewritten in the present (irreducible) context as:

$$X_{[i}(C_{mn}^k) + C_{[ij}^k C_{|mn]}^j = \frac{1}{3} \sigma_{imn}^{kl} \varphi_l \quad (5.16)$$

The factor $\frac{1}{3}$ on the right-hand side comes from the fact that the left-hand side possesses a factor $\frac{1}{3}$ contained in the full antisymmetrization of the indices i, m, n :

$$X_{[i}(C_{mn}^k) + C_{[ij}^k C_{|mn]}^j = \frac{1}{3} \left(X_i(C_{mn}^k) + C_{ij}^k C_{mn}^j + X_m(C_{ni}^k) + C_{mj}^k C_{ni}^j + X_n(C_{im}^k) + C_{nj}^k C_{im}^j \right)$$

Thus, Equation (5.15) becomes:

$$d^2\eta^k = -\frac{1}{6} \sigma_{imn}^{kl} \varphi_l \eta^i \wedge \eta^m \wedge \eta^n \quad (5.17)$$

This expression does not vanish on the phase space, but when restricted to the constraint surface, it does, because of the presence of φ_l . The differential d on the ghosts η^k thus measures the closure of the Jacobi identity (4.83).

The action of the differential d on the smooth function is easy to compute:

$$\begin{aligned} d^2f &= d(\eta^k X_k(f)) = -\frac{1}{2} C_{ij}^k X_k(f) \eta^i \wedge \eta^j + \eta^l \eta^k X_l(X_k(f)) \\ &= \frac{1}{2} \left([X_i, X_j](f) - C_{ij}^k X_k(f) \right) \eta^i \wedge \eta^j = \frac{1}{2} \varphi_k X_{C_{ij}^k}(f) \eta^i \wedge \eta^j \end{aligned} \quad (5.18)$$

where we passed from the first line to the second by using the following identity: $\eta^l \wedge \eta^k X_l(X_k(f)) = \frac{1}{2} [X_i, X_j](f) \eta^i \wedge \eta^j$. We used Equation (4.79) to obtain the very last term of the last line. The last line then shows that $d^2f \approx 0$ on the constraint surface only, and that there, we indeed have the already known result (Equation (4.80)): $[X_i, X_j] \approx C_{ij}^k X_k$. So, using Equations (5.17), (5.18) and the Leibniz rule (5.14), one deduce that:

$$d^2 \approx 0 \quad (5.19)$$

on any section of $\Gamma(\wedge^\bullet D')$.

The fact that d^2 does not square to zero outside Σ implies that the following diagram is *not* a chain complex:

$$0 \longrightarrow \mathcal{C}^\infty(T^*Q) \xrightarrow{d} \Gamma(D') \xrightarrow{d} \dots \xrightarrow{d} \Gamma(\wedge^{p-1}D') \xrightarrow{d} \Gamma(\wedge^p D') \longrightarrow 0$$

However, the introduction of the linear map d is necessary as we will now show. Since the vector fields X_i are tangent to Σ (see Proposition 4.64), their action on the smooth functions vanishing on Σ is zero. Thus, any function $f \in \mathcal{I}_\Sigma$ is sent to zero under the action of the differential of Equation (5.12). Then, by the Leibniz rule (5.14), we see that restricting each vector bundle $\wedge^k D'$ to the embedded submanifold Σ does not have any impact on the definition of the differential in Equations (5.12) and (5.13). More precisely, since $d\varphi_i = \{\varphi_j, \varphi_i\} \eta^j = C_{ij}^k \varphi_k \eta^j$ is weakly vanishing, we deduce that $d(\mathcal{I}_\Sigma) \subset \Gamma_\Sigma(D')$, where Γ_Σ means the section of D' vanishing on Σ . More generally, we observe that:

$$d(\Gamma_\Sigma(\wedge^k D')) \subset \Gamma_\Sigma(\wedge^{k+1} D') \quad (5.20)$$

Then, let W^k be the space of sections (with respect to Σ) of the vector bundle $\wedge^k D'$. As $\mathcal{C}^\infty(\Sigma) \simeq \mathcal{C}^\infty(T^*Q)/\mathcal{I}_\Sigma$, the quotient of $\Gamma(\wedge^k D')$ by \mathcal{I}_Σ gives W^k , because when quotienting we only consider the sections that differ on Σ , and their values outside the surface is of no importance for us whatsoever:

$$W^k \simeq \Gamma(\wedge^k D')/\mathcal{I}_\Sigma$$

So, the space W^k is the sheaf of sections *over* Σ of the vector bundle $\wedge^k D'$, and by Equation (5.20), the map d passes to the quotient and canonically induces a map $d : W^k \rightarrow W^{k+1}$.

Thus, we can pass each term of the sequence $\Gamma(\wedge^\bullet D')$ to the quotient by \mathcal{I}_Σ and obtain, this time, a chain complex:

$$0 \longrightarrow \underbrace{\mathcal{C}^\infty(T^*Q)/\mathcal{I}_\Sigma}_{\simeq \mathcal{C}^\infty(\Sigma)} \xrightarrow{d} W^1 \xrightarrow{d} \dots \xrightarrow{d} W^{p-1} \xrightarrow{d} W^p \longrightarrow 0$$

This is a chain complex because each term is a space of sections over Σ and $d^2 = 0$ on Σ . Notice however that the chain complex is not necessarily exact, i.e. a priori there is no reason for it to be a resolution. We call the corresponding cohomology groups the *cohomology of d relatively to δ* and we denote it by $H^k(d|H^\bullet(\delta))$ in full generality, or $H^k(d)$ for simplicity.

Our main group of interest is the zero-th cohomology group of d . Since $\mathcal{C}^\infty(\Sigma)$ is isomorphic to the quotient $\mathcal{C}^\infty(T^*Q)/\mathcal{I}_\Sigma$, the equivalence classes in $H^0(d)$ are in one-to-one correspondence with the gauge-invariant functions on Σ :

$$\begin{aligned} H^0(d) &= \left\{ \text{equivalence classes of } \mathcal{C}^\infty(T^*Q)/\mathcal{I}_\Sigma \text{ invariant under the } X_i \text{'s} \right\} \\ &\simeq \left\{ \text{gauge invariant functions on } \Sigma \right\} \simeq \mathcal{C}^\infty(\Sigma_{ph}) \end{aligned}$$

We emphasize that the present cohomology is defined with respect to the chain complex (W^\bullet, d) and not $(\Gamma(\wedge^\bullet D'), d)$. Thus, we managed to express the physical observables using only cohomological techniques, that we will push a bit further later. Notice that the higher cohomology groups of the chain complex (W^\bullet, d) are related with anomalies in quantum field theories (i.e. the failure to finding a quantum action that is gauge invariant), see subsection 11.1.2 in [Henneaux and Teitelboim, 1992] and [Fuster et al., 2005].

As the ghosts η^i are differential one-forms, we see that they can be understood as coordinate functions on some p -dimensional vector space \mathbb{R}^p , identified with the fiber $D_{(q,p)}$ of the regular foliation D at each point $(q,p) \in T^*Q$. We will let μ_i be the corresponding free generators of this vector space \mathbb{R}^p , and consider that they have ghost number -1 . They are then in one to

one correspondence with the vector fields X_i which generate the foliation made of the gauge orbits. Similarly, we will consider that the ghost momenta are also linear coordinate functions, conjugate to the ghosts, and then dual to p abstract elements u^i that we consider carrying a ghost number $+1$. More precisely, the ghost momenta \mathcal{P}_j can be identified with the linear fiber coordinates p_i on the cotangent bundle $T^*\mathbb{R}^p$, and their dual abstract elements u^i can be identified with the generators dq^i of the fiber of $T^*\mathbb{R}^p$.

Let us now define a new phase space containing the variables that we introduced. The *extended phase space* is the following vector space:

$$\mathfrak{P} = T^*Q \otimes \wedge^1(\mu_1, \dots, \mu_p) \otimes \wedge^1(u^1, \dots, u^p)$$

The letter on the left is a gothic capital P . Several remarks should be made: first, notice that $\mathfrak{P} \simeq T^*Q \otimes \mathbb{R}^p \otimes \mathbb{R}^p$; we used the notation \wedge^1 to emphasize that the elements on the right-hand side have ghost number -1 (for the μ_i) or $+1$ (for the u^i). We could have used for example the notation $\text{Span}(\mu_1, \dots, \mu_p)$ or just \mathbb{R}^p but this would have obliterated the ghost number of the variables. This degree will turn out to be quite important later on. So in the end, we could formally write that $\mathfrak{P} \simeq T^*Q \otimes T^*\mathbb{R}^p$, where we should keep in mind that the generators of the last cotangent bundle actually carry a ghost number.

Now, for a reason that will become clear in Section 5.3, we will formally define the ‘functions’ on \mathfrak{P} as the following space:

$$\mathcal{C}^\infty(\mathfrak{P}) = \mathcal{C}^\infty(T^*Q) \otimes \wedge^\bullet(\eta^1, \dots, \eta^p) \otimes \wedge^\bullet(\mathcal{P}_1, \dots, \mathcal{P}_p) \quad (5.21)$$

Notice that, as usual, the ‘functions’ on the space \mathfrak{P} depend on the coordinate functions q^i, p_i, η^i and \mathcal{P}_i . The product $\wedge^\bullet(\eta^1, \dots, \eta^p) \otimes \wedge^\bullet(\mathcal{P}_1, \dots, \mathcal{P}_p)$ plays the role of $\mathcal{C}^\infty(T^*\mathbb{R}^p)$ where the \wedge symbolizes that ghost numbers have been taken into account (see Section 5.3). As in the classical cotangent bundle situation – see Example 3.3 – the space \mathfrak{P} comes equipped with a standard graded Poisson bracket that in the present context can be defined on alternating powers of the variables η^a, \mathcal{P}_b as follows (the formula will be explained in Proposition 5.46):

$$\{F, G\} = \sum_{i=1}^n \frac{\partial F}{\partial q^i} \frac{\partial G}{\partial p_i} - \frac{\partial F}{\partial p_i} \frac{\partial G}{\partial q^i} - (-1)^{\text{gh}(F)} \sum_{a=1}^p \frac{\partial F}{\partial \eta^a} \frac{\partial G}{\partial \mathcal{P}_a} + \frac{\partial F}{\partial \mathcal{P}_a} \frac{\partial G}{\partial \eta^a} \quad (5.22)$$

In the above formula, the polynomial functions F and G are both homogeneous functions of $\mathcal{C}^\infty(\mathfrak{P})$. The plus sign between the two last terms comes from the fact that η^a and \mathcal{P}_b have odd ghost number, so that their derivatives anti-commute:

$$\frac{\partial}{\partial \eta^a} \wedge \frac{\partial}{\partial \mathcal{P}_b} = - \frac{\partial}{\partial \mathcal{P}_b} \wedge \frac{\partial}{\partial \eta^a} \quad (5.23)$$

For more details see the discussion around Equation (5.68).

The bracket (5.22) is compatible with the original Poisson bracket on the cotangent space T^*Q by setting that the Poisson bracket with a smooth function $f \in \mathcal{C}^\infty(T^*Q)$ and a ghost, or a ghost momentum, is zero. On $\mathcal{C}^\infty(T^*Q)$, we decide that the Poisson bracket is the canonical standard one. These rules then extends the Poisson bracket (5.22) to the whole of $\mathcal{C}^\infty(\mathfrak{P})$. It should be emphasized that this bracket is *graded skew-symmetric*, i.e. we have:

$$\{F, G\} = -(-1)^{\text{gh}(F)\text{gh}(G)} \{G, F\} \quad (5.24)$$

In particular, if F and G are ‘functions’ on \mathfrak{P} of odd ghost number, the commutator $\{F, G\}$ is symmetric. In such a case, if $F = G$ then $\{F, F\}$ does not automatically vanish, so if it does it

is not a trivial identity (see e.g. Equation (5.72)). Moreover, the fact that the ghosts and their ghost momenta are *conjugate variables* translates as:

$$\{\eta^a, \mathcal{P}_b\} = \{\mathcal{P}_b, \eta^a\} = \delta_b^a \quad (5.25)$$

where the symmetry of the bracket comes directly from Equation (5.24) and the fact that $\text{gh}(\eta^a)\text{gh}(\mathcal{P}_b) = -1$.

The space of ‘functions’ on \mathfrak{P} (see Section 5.3) – denoted $\mathcal{C}^\infty(\mathfrak{P})$ – can be understood a bi-graded vector space $M^{\bullet,\bullet}$ such that for any $0 \leq m, n \leq p$, one has:

$$M^{m,n} = \mathcal{C}^\infty(T^*Q) \otimes \wedge^m(\eta^1, \dots, \eta^p) \otimes \wedge^n(\mathcal{P}_1, \dots, \mathcal{P}_p)$$

The elements of $M^{m,n}$ have ghost number $m - n$. This bi-graded vector space can be equipped with an ‘almost’ bi-complex structure. More precisely, we can extend the maps δ and d to $M^{\bullet,\bullet}$ by defining their action on ghosts and ghost momenta, respectively, so that these two maps d and δ can be graphically represented as follows:

$$\begin{array}{ccccccc} & \dots & & \dots & & \dots & \\ & \downarrow \delta & & \downarrow \delta & & \downarrow \delta & \\ M^{0,2} & \xrightarrow{d} & M^{1,2} & \xrightarrow{d} & M^{2,2} & \xrightarrow{d} & \dots \\ & \downarrow \delta & & \downarrow \delta & & \downarrow \delta & \\ M^{0,1} & \xrightarrow{d} & M^{1,1} & \xrightarrow{d} & M^{2,1} & \xrightarrow{d} & \dots \\ & \downarrow \delta & & \downarrow \delta & & \downarrow \delta & \\ M^{0,0} & \xrightarrow{d} & M^{1,0} & \xrightarrow{d} & M^{2,0} & \xrightarrow{d} & \dots \end{array}$$

The map d is not a differential on the lower line so there is a priori no reason that it would square to zero on the whole bigraded space, preventing it to be a *bicomplex*.

The bi-graded vector space $M^{\bullet,\bullet}$ is equal to the tensor product of K_\bullet with $\wedge^\bullet(\eta^1, \dots, \eta^p)$. We let N_k be the graded vector space defined as:

$$N_k = K_{-k} \otimes \wedge^\bullet(\eta^1, \dots, \eta^p)$$

It corresponds to the horizontal lines in the above diagram. The properties of the Koszul resolution – in particular Lemmas 15.30.2 and 15.30.3 for such a statement in this [stack](#) where one should identify M with $\wedge^\bullet(\eta^1, \dots, \eta^p)$ – imply that N_\bullet , when equipped with the differential δ only, is a resolution of $\mathcal{C}^\infty(\Sigma) \otimes \wedge^\bullet(\eta^1, \dots, \eta^p) \simeq W^\bullet$. More precisely, the zero-th homology group of $\delta : N_k \rightarrow N_{k-1}$ is isomorphic to the chain complex W^\bullet :

$$H_0(N, \delta) \simeq W^\bullet$$

Since (W^\bullet, d) is a resolution of $\mathcal{C}^\infty(\Sigma_{ph})$, we deduce that the latter space can be obtained as the zero-th cohomology group of the restriction of d to $H_0(N, \delta)$. The procedure can be illustrated as first passing to the δ -homology, and then to the d -cohomology:

$$\begin{array}{cccccccc}
\dots & \dots & \dots & & \dots & \dots & \dots & \\
0 & 0 & 0 & \dots & 0 & 0 & 0 & \dots \\
0 & 0 & 0 & \dots & 0 & 0 & 0 & \dots \\
\mathcal{C}^\infty(\Sigma) & \xrightarrow{d} & W^1 & \xrightarrow{d} & W^2 & \xrightarrow{d} & \dots & \\
& & & & \mathcal{C}^\infty(\Sigma_{ph}) & & H^1(d) & & H^2(d) & & \dots
\end{array}$$

The above procedure allows to unify homological properties of both ghosts and ghost momenta. However, the problem with this procedure is that it is not implemented in the formalism: one has to take one homology first, and then take another cohomology. We would like to obtain the final data – smooth functions on the reduced phase space Σ_{ph} – as the zero-th cohomology group of only one differential. Working with a bi-graded vector space and two differentials δ and d in general allows to take their sum $s = \delta + d$ as a total differential on $M^{\bullet,\bullet}$. However in the present case, since $d^2 \neq 0$ on $M^{\bullet,\bullet}$ – it is only weakly vanishing on the lower line – we will not have $s^2 \neq 0$. As we have seen so far, it seems not sensible to require d to be a differential, but we can require it to satisfy a weaker condition:

Definition 5.5. *We say that the linear map $d : M^{\bullet,\bullet} \rightarrow M^{\bullet+1,\bullet}$ is a δ -exact differential if there exists a derivation Δ of $M^{\bullet,\bullet}$ satisfying:*

$$d^2 = -[\delta, \Delta] \quad (5.26)$$

Since d^2 increases the number of ghosts by 2, and since δ decreases the number of ghost momenta by one, one deduces that, if it exists, the derivation Δ sends $M^{\bullet,\bullet}$ to $M^{\bullet+2,\bullet+1}$. Moreover, it can be shown that Equation (5.19) can be recasted into the form of Equation (5.26), if one assumes that $\delta(\eta^i) = 0$. Indeed, since $\delta(\mathcal{C}^\infty(T^*Q)) = 0$ we conclude that the condition $\delta(\eta^i) = 0$ implies that $\delta(\Gamma(\wedge^\bullet D')) = 0$. Then, knowing how δ acts on ghost momenta, one can find a derivation $\Delta : M^{\bullet,0} \rightarrow M^{\bullet,1}$ so that $d^2 f = -\delta \circ \Delta(f)$ (the missing term $-\Delta \circ \delta$ from Equation (5.26) being identically zero on $\Gamma(\wedge^\bullet D')$):

$$\Delta(f) = -\frac{1}{2} X_{C_{ij}^k}(f) \eta^i \wedge \eta^j \otimes \mathcal{P}_k \quad (5.27)$$

Regarding the ghosts η^k , by looking at Equation (5.17) we deduce that:

$$\Delta(\eta^k) = -\frac{1}{6} \sigma_{imn}^{kl} \eta^i \wedge \eta^m \wedge \eta^n \otimes \mathcal{P}_l \quad (5.28)$$

so that $d^2(\eta^k) = -\delta \circ \Delta(\eta^k)$. There is no other choice, because $\delta(\eta^k) = 0$ so we cannot have any contribution to $d^2(\eta^k)$ of the form $-\Delta \circ \delta(\eta^k)$. We have a minus sign on the right-hand side of Equation (5.28) because when $-\delta$ acts on it, it has to jump over three ghosts before acting on \mathcal{P}_l , hence not changing the overall sign and giving the desired formula (5.17).

These observations show that the lower line of the bigraded vector space $M^{\bullet,\bullet}$ satisfies Equation (5.26). The aim of the BRST formalism is then to:

1. extend the action of d to the ghost momenta \mathcal{P}_a so that it becomes δ -exact everywhere on $M^{\bullet,\bullet}$, and

2. find additional terms to be added to the map:

$$s = \delta + d + \text{'more'} \quad (5.29)$$

so that it squares to zero.

As a consequence, one can show that the (now well-defined) zero cohomology group $H^0(s)$ is precisely isomorphic to $C^\infty(\Sigma_{ph})$ (this is the content of Theorem 5.12).

Notice that the first term in the list of additional terms symbolized by ‘more’ in Equation (5.29) precisely corresponds to the derivation Δ defined on the right-hand side of Equation (5.26). Its action on ghost momenta can be determined through the following proposition:

Proposition 5.6. *It is possible to extend the maps δ and d on the whole bigraded vector space $M^{\bullet,\bullet}$ so that they satisfy their respective properties:*

$$\delta^2 = 0 \quad \text{and} \quad d^2 \text{ is } \delta\text{-exact} \quad (5.30)$$

and so that they commute as well:

$$\delta \circ d + d \circ \delta = 0 \quad (5.31)$$

Remark 5.7. Henneaux and Teitelboim call a linear derivation d on $M^{\bullet,\bullet}$ satisfying Equations (5.30) and (5.31) a *differential modulo δ* (see subsection 8.2.9 in [Henneaux and Teitelboim, 1992]).

Proof. In order to satisfy Equations (5.30), one could naively say that $\delta(\eta^i) = 0$ and $d(\mathcal{P}_j) = 0$, but this would not be consistent with Equation (5.31). Indeed, on the ghost momenta, one would have:

$$(\delta \circ d + d \circ \delta)(\mathcal{P}_j) = d \circ \delta(\mathcal{P}_j) = d(\varphi_j) = X_i(\varphi_j) \eta^i = \{\varphi_i, \varphi_j\} \eta^i = C_{ij}^k \varphi_k \eta^i \quad (5.32)$$

which is in general not zero except on the constraint surface Σ . A way out of this problem is to first drop the assumption that $d(\mathcal{P}_j) = 0$, and then, knowing that $d \circ \delta(\mathcal{P}_j) = C_{ij}^k \varphi_k \eta^i$, to observe that Equation (5.31) applied to \mathcal{P}_j translates as:

$$\delta \circ d(\mathcal{P}_j) = -C_{ij}^k \varphi_k \eta^i = \delta(C_{ij}^k \mathcal{P}_k \otimes \eta^j) = \delta(-C_{ij}^k \eta^j \otimes \mathcal{P}_k)$$

The last equalities can be understood by recalling that permuting a ghost and a ghost momentum brings a minus sign. So we set the following prescriptions:

$$\delta(\eta^i) = 0 \quad \text{and} \quad d(\mathcal{P}_i) = -C_{ij}^k \eta^j \otimes \mathcal{P}_k \quad (5.33)$$

Remark 5.8. As a remark, if one sees ghosts and ghost momenta as functions on the extended phase space \mathfrak{P} , then the Poisson bracket (5.22) applies to products of ghosts and ghost momenta, and in particular to terms such as the element $\Omega_{(1)} = -\frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \otimes \mathcal{P}_k \in M^{2,1}$, which is equivalent to $d\eta^k \otimes \mathcal{P}_k$. Equation (5.13) and the second equation in (5.33) can be recasted in the following way:

$$d\eta^k = \{\Omega_{(1)}, \eta^k\} \quad \text{and} \quad d\mathcal{P}_i = \{\Omega_{(1)}, \mathcal{P}_i\}$$

So the longitudinal differential is Poisson-exact; this observation will become a key argument in the BRST formalism.

In order to show that d^2 is δ -exact, we need to show that there exists a derivation Δ of $\mathcal{C}^\infty(\mathfrak{P})$ of ghost number +1 and such that Equation (5.26) is satisfied. We already know how this derivation acts on $\Gamma(\wedge^\bullet D')$, as is presented in the discussion below Definition 5.5. On smooth functions and on the ghosts, we have defined the map Δ via Equations (5.27) and (5.28), in such a way that $d^2(f) = -\delta \circ \Delta(f)$ and $d^2(\eta^k) = -\delta \circ \Delta(\eta^k)$. These two latter equations correspond to Equation (5.26) on f and η^k , as $\delta(f) = 0$ and $\delta(\eta^k) = 0$.

We now need to show that such a derivation Δ is well defined on the ghost momenta \mathcal{P}_j . In order to do that we will use some algebraic property of the Koszul complex, which will allow us to deduce that if d^2 is δ -closed on the ghost momenta, then it is δ -exact. When d^2 acts on the bigraded vector space $M^{\bullet, \bullet}$, it preserves the number of ghost momenta, while it adds two ghosts. Then, if one denotes the p -dimensional vector space of ghost momenta by the letter $E = \wedge^1(\mathcal{P}_1, \dots, \mathcal{P}_p)$, and the full exterior algebra of ghosts by the letter $F = \wedge^\bullet(\eta^1, \dots, \eta^2)$, the map d^2 can be understood as an element of the vector space $E^* \otimes E \otimes \mathcal{C}^\infty(T^*Q) \otimes F$ (more precisely, as an element of $E^* \otimes E \otimes \mathcal{C}^\infty(T^*Q) \otimes \wedge^2(\eta^1, \dots, \eta^2)$ as it has ghost number 2 and does not increase the number of ghost momenta). But $E \otimes \mathcal{C}^\infty(T^*Q)$ is the first space of the Koszul resolution, denoted earlier by the letter K_{-1} . So we can imagine that there is a chain complex:

$$0 \longrightarrow E^* \otimes F \otimes K_{-p} \xrightarrow{\delta} \dots \xrightarrow{\delta} E^* \otimes F \otimes K_{-1} \xrightarrow{\delta} E^* \otimes F \otimes \mathcal{C}^\infty(T^*Q) \longrightarrow 0$$

without a priori knowing (yet) if this is a resolution.

The map δ acts only on the ghost momenta, and acts neither on E^* nor on F . One can easily check that the action of δ on an element $\alpha \in E^* \otimes F \otimes K_{-k}$ is realized via the graded commutator of operators:

$$\delta(\alpha) = [\delta, \alpha] = \delta \circ \alpha - (-1)^{\text{gh}(\alpha)} \alpha \circ \delta$$

Then, since d^2 is an element of $E^* \otimes F \otimes K_{-1}$, if the differential were exact at $E^* \otimes F \otimes K_{-1}$, i.e. if $H_1(E^* \otimes F \otimes K_\bullet) = 0$ and if we have that $[\delta, d^2](\mathcal{P}_j) = 0$ then d^2 would be δ -exact on the ghost momenta. Thanks to very properties of the Koszul complex, this is precisely the case. Both E^* and F are $\mathcal{C}^\infty(T^*Q)$ -modules, and the first-class constraints φ_i are a $E^* \otimes F$ -regular sequence. Then, the new chain complex $(E^* \otimes F \otimes K_\bullet, \delta)$ is still acyclic, i.e. it is exact with respect to the differential δ (see Lemmas 15.30.2 and 15.30.3 for such a statement in this [stack](#) where one should identify M with $E^* \otimes F$). Then, we deduce from this discussion that we only need to show that d^2 is δ -closed, i.e. that $[\delta, d^2](\mathcal{P}_j) = 0$ or, equivalently, that the following identity holds (d^2 has ghost number 2):

$$\delta(d^2(\mathcal{P}_i)) = d^2(\delta(\mathcal{P}_i)) \tag{5.34}$$

We will first compute the left-hand side, and then we will show that it is equal to the right-hand side, which have already been computed in Equations (5.18) for a general function. By the second equation in (5.33) we observe that we have:

$$\begin{aligned} d^2(\mathcal{P}_i) &= -d(C_{ij}^k \eta^j \otimes \mathcal{P}_k) \\ &= -d(C_{ij}^k) \wedge \eta^j \otimes \mathcal{P}_k - C_{ij}^k d\eta^j \otimes \mathcal{P}_k + C_{ij}^k \eta^j \otimes d\mathcal{P}_k \\ &= -\{\varphi_l, C_{ij}^k\} \eta^l \wedge \eta^j \otimes \mathcal{P}_k + \frac{1}{2} C_{ij}^k C_{mn}^j \eta^m \wedge \eta^n \otimes \mathcal{P}_k - C_{ij}^k C_{kr}^s \eta^j \wedge \eta^r \otimes \mathcal{P}_s \\ &= \left(-\{\varphi_m, C_{in}^k\} + \frac{1}{2} C_{ij}^k C_{mn}^j + C_{im}^j C_{nj}^k \right) \eta^m \wedge \eta^n \otimes \mathcal{P}_k \\ &= \frac{1}{2} \left(\{\varphi_n, C_{im}^k\} + \{\varphi_m, C_{ni}^k\} + C_{ij}^k C_{mn}^j + C_{im}^j C_{nj}^k + C_{ni}^j C_{mj}^k \right) \eta^m \wedge \eta^n \otimes \mathcal{P}_k \end{aligned} \tag{5.35}$$

The term in the parenthesis of the last line resembles a Jacobi identity. Indeed, it corresponds to the left-hand side of Equation (4.83), without reducibility functions and missing the term $\{\varphi_i, C_{mn}^k\}$:

$$\{\varphi_i, C_{mn}^k\} + C_{ij}^k C_{mn}^j + \{\varphi_n, C_{im}^k\} + C_{nj}^k C_{im}^j + \{\varphi_m, C_{ni}^k\} + C_{mj}^k C_{ni}^j = \sigma_{mni}^{kl} \varphi_l \quad (5.36)$$

Where $\sigma_{mni}^{kl} = -\sigma_{mni}^{lk}$ some smooth function fully antisymmetric on the lower and the upper indices. In particular we have:

$$C_{ij}^k C_{mn}^j + \{\varphi_n, C_{im}^k\} + C_{nj}^k C_{im}^j + \{\varphi_m, C_{ni}^k\} + C_{mj}^k C_{ni}^j = -\{\varphi_i, C_{mn}^k\} + \sigma_{mni}^{kl} \varphi_l \quad (5.37)$$

So, recalling that $X_{C_{mn}^k}(f) = -\{f, C_{mn}^k\}$, when Equation (5.37) is applied to the last term of Equation (5.18) for $f = \varphi_i$ we obtain:

$$d^2 \varphi_i = \frac{1}{2} \varphi_k \left(C_{ij}^k C_{mn}^j + \{\varphi_n, C_{im}^k\} + C_{nj}^k C_{im}^j + \{\varphi_m, C_{ni}^k\} + C_{mj}^k C_{ni}^j \right) \eta^m \wedge \eta^n \quad (5.38)$$

The contribution coming from $\sigma_{mni}^{kl} \varphi_l$ automatically vanishes because it is contracted with φ_k . As $\varphi_i = \delta(\mathcal{P}_i)$, we deduce that Equation (5.38) corresponds to $d^2(\delta(\mathcal{P}_i))$. On the other hand, the action of δ on the last line of Equations (5.35) precisely gives Equation (5.38). Thus, this implies that Equation (5.34) holds, meaning that d^2 is δ -exact on the ghost momenta (see Equation (5.40)).

We can actually give the explicit expression of the derivation Δ on the ghost momenta, because it resembles that of Equation (5.27). Reinjecting Equation (5.37) into the last line of Equations (5.35) gives the following expression:

$$d^2(\mathcal{P}_i) = \frac{1}{2} \left(X_{C_{mn}^k}(\varphi_i) + \sigma_{mni}^{kl} \varphi_l \right) \eta^m \wedge \eta^n \otimes \mathcal{P}_k \quad (5.39)$$

Here the contribution $\sigma_{mni}^{kl} \varphi_l$ does not vanish because there is no contraction with φ_k , but it still vanishes on Σ , together with the second term. So, if one sets (notice the position of the k, l indices):

$$\Delta(\mathcal{P}_k) = -\frac{1}{4} \sigma_{mni}^{kl} \eta^m \wedge \eta^n \otimes \mathcal{P}_l \wedge \mathcal{P}_k$$

then by Equation (5.27), together with the fact that $\delta(\mathcal{P}_i) = \varphi_i$, one has:

$$d^2(\mathcal{P}_i) = -[\delta, \Delta](\mathcal{P}_i) \quad (5.40)$$

which is Equation (5.26) on the ghost momenta.

Now one only needs to show Equation (5.31). Let $\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n} \in M^{m,n}$, for $n \geq 1$. Then, since the differential δ has ghost number +1 and does not act on ghosts, one has:

$$\begin{aligned} d \circ \delta(\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) &= (-1)^m d(\eta^{i_1} \wedge \dots \wedge \eta^{i_m}) \otimes \delta(\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \\ &\quad + \eta^{i_1} \wedge \dots \wedge \eta^{i_m} \otimes d \circ \delta(\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \end{aligned}$$

while on the other hand:

$$\begin{aligned} \delta \circ d(\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) &= (-1)^{m+1} d(\eta^{i_1} \wedge \dots \wedge \eta^{i_m}) \otimes \delta(\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \\ &\quad + \eta^{i_1} \wedge \dots \wedge \eta^{i_m} \otimes \delta \circ d(\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \end{aligned}$$

Summing the two identities, one finds that:

$$(\delta \circ d + d \circ \delta)(\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \otimes \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) = \eta^{i_1} \wedge \dots \wedge \eta^{i_m} \otimes (\delta \circ d + d \circ \delta)(\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \quad (5.41)$$

But the very definition of the action of d on the ghost momenta in Equation (5.33) has been chosen so that the right-hand side of Equation (5.32) vanishes. Knowing this fact, one can straightforwardly check that the right-hand side of Equation (5.41) is zero. For $n = 0$ the identity (5.31) is trivial because $\delta(M^{\bullet,0}) = 0$. This concludes the proof. \square

Remark 5.9. The role of Proposition 5.6 is to replace Equation (5.19) by a much more general condition symbolized by Equation (5.26) which is valid for d everywhere on the bigraded vector space $M^{\bullet,\bullet}$, and which reduces to Equation (5.19) on the bottom line $M^{\bullet,0} = \Gamma(\wedge^\bullet D')$. Interestingly, Henneaux and Teitelboim claim that it is also possible to show that $d^2(\mathcal{P}_j) \approx 0$ (see the end of subsection 9.2.3 and Exercice 9.14 [Henneaux and Teitelboim, 1992]) – meaning that Equation (5.19) could be extended to the entire bigraded space $M^{\bullet,\bullet}$ – but Equation (5.39) shows that a priori $d^2(\mathcal{P}_k)$ needs not vanish on Σ . Given the argument below Equation (5.39), it would indeed require that the structure functions are gauge invariant. Be aware that d^2 not being weakly zero on the ghost momenta does not imply that it is not zero on W^\bullet , as in the latter space there are no ghost momenta, hence no obstruction for d being a differential. Eventually, the BRST formulation of regular Lie algebroids seems to be a counter example to Henneaux and Teitelboim’s claim (see Example 5.52).

So we have now the knowledge of the action the derivations δ, d and Δ on the functions, ghosts and ghost momenta:

$$\delta(f) = 0, \quad d(f) = X_i(f)\eta^i, \quad \Delta(f) = -\frac{1}{2}X_{C_{ij}^k}(f)\eta^i \wedge \eta^j \otimes \mathcal{P}_k, \quad (5.42)$$

$$\delta(\eta^k) = 0, \quad d(\eta^k) = -\frac{1}{2}C_{ij}^k \eta^i \wedge \eta^j, \quad \Delta(\eta^k) = -\frac{1}{6}\sigma_{imn}^{kl} \eta^i \wedge \eta^m \wedge \eta^n \otimes \mathcal{P}_l, \quad (5.43)$$

$$\delta(\mathcal{P}_i) = \varphi_i, \quad d(\mathcal{P}_i) = -C_{ij}^k \eta^j \otimes \mathcal{P}_k, \quad \Delta(\mathcal{P}_i) = \frac{1}{4}\sigma_{imn}^{kl} \eta^m \wedge \eta^n \otimes \mathcal{P}_k \wedge \mathcal{P}_l. \quad (5.44)$$

where σ_{mni}^{kl} is the tensor defined in Equation (5.36), and corresponds to the *higher order structure functions* encoding the algebra of constraints (see Remark 4.74). Writing $s = \delta + d + \Delta + \text{‘more’}$, we first observe that each term on the last line has a different relationship with ghost momenta. When their respective action is not zero, the map δ makes the number of ghost momenta decrease by one, d does not change the number of ghost momenta, while Δ makes their number increase by one. To make things more explicit, we would define a new grading, the *pure antighost number*²⁴, which just counts the number of ghost momenta in an element of $\mathcal{C}^\infty(\mathfrak{P})$. This grading is different than the former *ghost number* defined earlier because it only sees the ghost momenta. Knowing both the ghost number and the pure antighost number of an element $u \in \mathcal{C}^\infty(\mathfrak{P})$ allows to know exactly how many ghosts and ghost momenta it contains, because its ghost number equates the difference of ghosts minus ghost momenta. The maps δ, d and Δ – although they all have ghost number +1 – inherit a pure antighost number of $-1, 0$ and $+1$, respectively, based on their action on monomials of ghost momenta. So the map s has ghost number +1 but contains terms of various pure antighost numbers. In particular we expect that the additional terms in the ‘more’ part of the map, have pure antighost number higher than +1 and increasing. For example we can show the following

Lemma 5.10. *There exists a derivation $s_{(2)}$ of $\mathcal{C}^\infty(\mathfrak{P})$ of ghost number 1 and pure antighost number 2 such that:*

$$[d, \Delta] + [\delta, s_{(2)}] = 0 \quad (5.45)$$

²⁴This is not a standard denomination, as it is usually merely called *antighost number*, see e.g. [Henneaux and Teitelboim, 1992]. We preferred to add the adjective *pure* in order to emphasize that this degree should be considered on par with the pure ghost number, and not with the (total) ghost number. The denomination *antighost* will become clear later, when we introduce the corresponding antighosts.

Proof. The map $[d, \Delta]$ is an element of $\wedge^1(\mathcal{P}_1, \dots, \mathcal{P}_p)^* \otimes \wedge^2(\mathcal{P}_1, \dots, \mathcal{P}_p) \otimes \mathcal{C}^\infty(T^*Q) \otimes \wedge^3(\eta^1, \dots, \eta^2)$. It is δ -closed, because:

$$[\delta, [d, \Delta]] = \underbrace{[[\delta, d], \Delta]}_{=0} - \underbrace{[d, [\delta, \Delta]]}_{=-d^2}$$

The first term on the right vanishes by Equation (5.31), while the second term is $\frac{1}{2}[d, [d, d]] = 0$, which is always zero, whatever the grading of d . The properties of the Koszul complex – Lemmas 15.30.2 and 15.30.3 for such a statement in this [stack](#) – imply that the chain complex $(K_\bullet \otimes \wedge^1(\mathcal{P}_1, \dots, \mathcal{P}_p)^* \otimes \wedge^3(\eta^1, \dots, \eta^2), [\delta, \cdot])$ is exact. Then, we deduce that $[d, \Delta]$ is δ -exact, i.e. there exists a derivation $s_{(2)}$ of $\mathcal{C}^\infty(\mathfrak{P})$ (of ghost number 1 and pure antighost number 2) such that Equation (5.45) is satisfied. \square

Let us now add $s_{(2)}$ to the map s and, acknowledging that there are still other terms left unknown beyond $s_{(2)}$, we compute s^2 :

$$s^2 = \delta^2 + [\delta, d] + d^2 + [\delta, \Delta] + [d, \Delta] + [\delta, s_{(2)}] + \dots$$

One can organize the resulting terms in a table depending on pure antighost numbers and observe that they give no contribution to s^2 :

Pure antighost number	Corresponding equation
-2	$\delta^2 = 0$
-1	$[\delta, d] = 0$
0	$d^2 + [\delta, \Delta] = 0$
+1	$[d, \Delta] + [\delta, s_{(2)}] = 0$
+2	$\Delta^2 + [d, s_{(2)}] + \dots$

The idea of BRST formalism is to use the properties of the Koszul complex to build by induction – as was done in Proposition 5.6 and Lemma 5.10 a family of maps $(s_{(k)})_{-1 \leq k \leq p}$ of ghost number 1 and pure antighost number k which belongs to $\wedge^1(\mathcal{P}_1, \dots, \mathcal{P}_p)^* \otimes \wedge^k(\mathcal{P}_1, \dots, \mathcal{P}_p) \otimes \mathcal{C}^\infty(T^*Q) \otimes \wedge^k(\eta^1, \dots, \eta^2)$, and is such that, for every $-1 \leq j \leq k$:

$$\sum_{\substack{-1 \leq m, n \leq j \\ m+n=j-1}} s_{(m)} \circ s_{(n)} = 0$$

where it is understood that $s_{(-1)} = \delta$, $s_{(0)} = d$ and $s_{(1)} = \Delta$. It stops at $k = p$ because $\wedge^{p+1}(\mathcal{P}_1, \dots, \mathcal{P}_p) = 0$ – this is a particularity of the irreducible case, and would not happen in the reducible one. Then the map of ghost number +1 defined as:

$$s = \sum_{-1 \leq i \leq p} s_{(i)} \tag{5.46}$$

satisfies by construction $s^2 = 0$. For more details on the computation see Sections 8.4, 8.5 and 9.3 in [[Henneaux and Teitelboim, 1992](#)].

Definition 5.11. *The total differential s of ghost number +1 defined in Equation (5.46) is called the BRST differential and turns the bi-graded vector space $M^{\bullet, \bullet}$ into a chain complex:*

$$0 \xrightarrow{s} M^{-p} \xrightarrow{s} \dots \xrightarrow{s} M^{-1} \xrightarrow{s} M^0 \xrightarrow{s} M^1 \xrightarrow{s} \dots \xrightarrow{s} M^p \xrightarrow{s} 0$$

where $M^k = \bigoplus_{0 \leq m, n \leq p, m-n=k} M^{m, n}$.

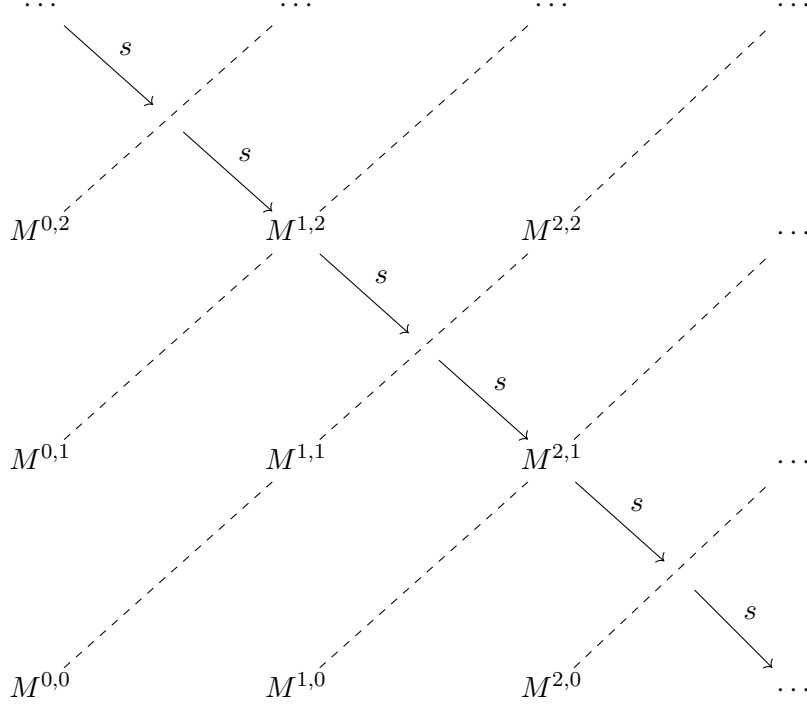
Be aware that although $\delta(\mathcal{P}_i) = \varphi_i$, since there are other components of various pure antighost numbers in s – in particular d , say, we do not necessarily have $s(\mathcal{P}_i) = \varphi_i$, because a minima d and Δ act as well. From Equations (5.42), (5.43) and (5.44), we can represent the action of these three differentials in $M^{\bullet,\bullet}$ by directional arrows:

$$\begin{array}{ccc}
 & & -\frac{1}{2}X_{C_{ij}^k}(f)\eta^i\wedge\eta^j\otimes\mathcal{P}_k \\
 & \nearrow \Delta & \\
 f & \xrightarrow{d} & X_i(d)\eta^i \\
 \downarrow \delta & & \\
 0 & &
 \end{array}$$

$$\begin{array}{ccc}
 & & -\frac{1}{6}\sigma_{imn}^{kl}\eta^i\wedge\eta^m\wedge\eta^n\otimes\mathcal{P}_l \\
 & \nearrow \Delta & \\
 \eta^k & \xrightarrow{d} & -\frac{1}{2}C_{ij}^k\eta^i\wedge\eta^j \\
 \downarrow \delta & & \\
 0 & &
 \end{array}$$

$$\begin{array}{ccc}
 & & \frac{1}{4}\sigma_{imn}^{kl}\eta^m\wedge\eta^n\otimes\mathcal{P}_k\wedge\mathcal{P}_l \\
 & \nearrow \Delta & \\
 \mathcal{P}^i & \xrightarrow{d} & -\frac{1}{2}C_{ij}^k\eta^i\otimes\mathcal{P}_k \\
 \downarrow \delta & & \\
 \varphi_i & &
 \end{array}$$

This provides a visual understanding of Equation (5.26) (to go further, see [Browning and McMullan, 1987]). We see also that the differential s sends an element of $M^{m,n}$ to a family of elements that sit on a diagonal of total degree $m - n + 1$. We can graphically represent the action of the BRST differential s on the bigraded vector space $M^{\bullet,\bullet}$, as it acts ‘diagonally’:



For example, starting from $M^{0,1}$, $\delta(M^{0,1}) \subset M^{0,0}$, $d(M^{0,1}) \subset M^{1,1}$, $\Delta(M^{0,1}) \subset M^{2,2}$, $s_{(2)}(M^{0,1}) \subset M^{3,3}$, etc. The k -th group of cohomology of s is defined as usual:

$$H^k(s) = \frac{\text{Ker}(s : M^k \rightarrow M^{k+1})}{\text{Im}(s : M^{k-1} \rightarrow M^k)}$$

The importance of the differential s relies on the following fact (Theorem 8.3 in [Henneaux and Teitelboim, 1992] and possibly early formulated in [Figuroa-O’Farrill and Kimura, 1991b]):

Theorem 5.12. *The cohomology of s is equal to the cohomology of d modulo δ , that is to say:*

$$H^k(s) = H^k(d)$$

for every k (understanding that for $-p \leq k \leq -1$, $H^k(d) = 0$). In particular, the zero-th group of cohomology of s coincides with that of d and we have:

$$H^0(s) \simeq \mathcal{C}^\infty(\Sigma_{ph})$$

Remark 5.13. In mathematical language, Theorem 5.12 can be advantageously reformulated in terms of *spectral sequences* [Dubois-Violette, 1987]. The equality of cohomologies is then understood as the fact that the spectral sequence degenerates after the second page (because δ is an exact differential), which is precisely the horizontal cohomology of the vertical cohomology $H^p(d, H^q(\delta))$.

Remark 5.14. It turns out that the graded Poisson bracket on $\mathcal{C}^\infty(\mathfrak{P})$ (see Section 5.3) descends to $H^0(s)$ and that the isomorphism $H^0(s) \simeq \mathcal{C}^\infty(\Sigma_{ph})$ is actually an isomorphism of graded Poisson algebras – where the Poisson bracket on the latter algebra is the canonical one obtained through Poisson reduction. Thus the BRST formalism is a way to perform Poisson reduction on coisotropic submanifolds – as stated in Proposition 3.97 – without passing to the quotient, but through cohomological techniques. This quite important and useful alternative approach to Poisson reduction has been coined by mathematicians *homological Poisson reduction* [Stasheff, 1997] or *homological symplectic reduction* [Figuroa-O’Farrill and Kimura, 1991b, Kimura, 1993].

Remark 5.15. The differential s is strictly equal to the sum $\delta + d$ only when the algebra of gauge transformations is a true Lie algebra, in the sense that the structure functions C_{ab}^c are structure *constant*. Then, we have that $X_{C_{ij}^k} = 0$, so the right-hand side of Equation (5.27) is zero, meaning that $\Delta(f) = 0$. Moreover, the algebra of gauge transformations is closed in the sense of Definition 4.69, so Lemma 4.73 implies that the term σ_{mni}^{kl} vanishes, so both $\Delta(\eta^k)$ and $\Delta(\mathcal{P}_i)$ vanish as well. This means that the derivation Δ is strictly zero, so we have $d^2 = 0$ on the bigraded space $M^{\bullet,\bullet}$ which is then a *bicomplex*. In that case the BRST differential is $s = \delta + d$, and no higher terms.

To summarize, we have extended the classical phase space T^*Q by adding ghosts and their conjugate ghost momenta, and we defined a differential s on the space of smooth functions $\mathcal{C}^\infty(\mathfrak{P})$, so that the zero-th cohomology group of this space is precisely the set of gauge-invariant functions on Σ , i.e. the classical observables. The main advantage of this approach is that it is purely algebraic because it only depends on the structure of the algebra of first-class constraints. The gauge conditions are not implemented so it provides a way of characterizing the complicated space $\mathcal{C}^\infty(\Sigma_{ph})$ in terms of the original canonical variables on T^*Q , together with ghosts, ghost momenta and a simple differential. Replacing gauge symmetries by a rigid symmetry controlled by the differential s is quite helpful because it allows to substitute the action depending on complicated variables (parametrizing the constraint surface Σ) with an action that is defined over the extended phase space \mathfrak{P} , hence preserving manifest covariance and locality²⁵. Integrating over these additional degrees of freedom in the new action yields infinities and divergences, which are luckily compensated by the presence of ghosts. Quantization of a classical theory with purely first-class systems turns out to be much more satisfying. For a wider discussion on these questions, see Section 8.5 and 9.5 in [Henneaux and Teitelboim, 1992].

Remark 5.16. The present section dealt with the BRST formalism of an *irreducible* first-class system. For a reducible system, the Koszul complex becomes more complicated and was improved by Tate, becoming a Koszul-Tate resolution. The treatment of such reducible systems of first-class constraints is dealt with in Section 5.4.

5.3 Graded geometry and its associates

Let us now provide some mathematical background to the graded geometric tools developed in Section 5.2. We need to start from the notion of *graded vector spaces* – see Definition A.1 – and elaborate from it. When E is a graded vector space, we would like to associate to it an ‘algebra of functions’. Due to the various degrees of elements of E , this is not straightforward so we should first give a meaning to the dual space E^* . Let $E = (E_i)_{i \in \mathbb{Z}}$ be a \mathbb{Z} -graded vector space, then we call *homogeneous element* a vector e who belongs to only one E_i , and we denote by $|e| = i$ its degree. Each space E_i admits a dual space $(E_i)^*$ where elements have by convention a degree $-i$. The dual space E^* is the sum of all the $(E_i)^*$ and as such, it is a graded vector space. The degree of the dual elements is chosen so that, for every $e \in E_i$ and $\alpha \in (E_i)^*$, the sum of the degrees of e and α vanishes:

$$|e| + |\alpha| = i - i = 0$$

so that the real number $\alpha(e)$ has indeed degree 0 – as would be expected from a real number.

A monomial on $(E_i)^*$ is an element denoted $\alpha_1 \wedge \alpha_2 \wedge \dots \wedge \alpha_p$ and its degree is $p \times (-i) = -ip$. The wedge product in this formula is *not* the wedge product from the exterior algebra of $(E_i)^*$,

²⁵See the introduction of [Fisch et al., 1989] for a glimpse at why gauge transformations and ghosts are necessary to ensure locality of physical theories. See these [lecture notes](#) to understand why fixing the Coulomb gauge in electromagnetism give a non-local theory.

as by convention we set it to satisfy the following identity:

$$\alpha_k \wedge \alpha_l = (-1)^{(-i)^2} \alpha_k \wedge \alpha_l$$

so it is not necessarily skew-symmetric if i is even. In particular on $(E_0)^*$, this wedge product does not coincide with the wedge product of the exterior algebra of $(E_0)^*$, because it is symmetric. This is a convention specific to graded geometry to which we will stick from now on. This wedge product extends canonically to the whole of E^* as a *graded commutative product*. Indeed, for any homogeneous elements $\alpha, \beta \in E^*$, we set:

$$\alpha \wedge \beta = (-1)^{|\alpha||\beta|} \beta \wedge \alpha$$

We see that it is graded because the degrees $|\alpha|$ and $|\beta|$ appear, and it is commutative because if ever one of them is even, the sign results in a $+1$. Building on Definition A.1, this discussion allows us to define the following notion:

Definition 5.17. A graded commutative algebra is a graded algebra $(A = \bigoplus_{i \in \mathbb{Z}} A_i, \wedge)$ such that the product satisfies the graded commutativity identity:

$$a \wedge b = (-1)^{|a||b|} b \wedge a \tag{5.47}$$

for any two homogeneous elements $a, b \in A$.

Example 5.18. A typical example of a graded commutative algebra is the *graded symmetric algebra* of a graded vector space E . The graded symmetric algebra $S(E)$ over a graded vector space E is the quotient of the tensor algebra of E by the ideal generated by the elements of the form: $a \otimes b - (-1)^{|a||b|} b \otimes a$. The wedge product (5.47) is then the graded symmetric product making this algebra graded commutative. The algebra of formal power series $\widehat{S}(E)$ built over $S(E)$ is also an example of graded commutative algebra.

Definition 5.19. Let $E = (E_i)_{i \in \mathbb{Z}}$ be a graded vector space. The graded commutative algebra $(\widehat{S}(E^*), \wedge)$ (resp. $S(E^*)$) of formal power series on E^* (resp. of polynomials on E^*) is called the algebra of functions on E (resp. the algebra of polynomial functions on E).

Remark 5.20. Notice that we used the symmetric algebra of E and not the exterior algebra, and at first sight it may seem contradictory as we used a wedge product. However, in fact the wedge product (5.47) is a *graded symmetric product*, and thus it is the adequate product on the graded symmetric algebra of E . Moreover, Proposition 5.21 explains that in fact the symmetric algebra is somehow a hidden exterior algebra.

In order to provide some examples of Definition 5.19, we need to introduce a new notion which, although coming from algebraic topology – is endogenous to the world of graded geometry. Let $E = (E_i)_{i \in \mathbb{Z}}$ be a graded vector space. The *suspension operator* s – also denoted $[-1]$ – on a graded vector space $E = \bigoplus_i E_i$ has the property of shifting the degree of homogeneous elements by $+1$. Namely,

$$(sE)_i = E_{i-1}$$

so that if x has degree $i - 1$, then sx has degree i . Obviously, this operator has an inverse, called the *desuspension operator*, which shifts the degree of homogeneous elements by -1 . The latter is also denoted $[1]$ in the literature. Then we define $E[1]$ to be the graded vector space $E[1] = \bigoplus_{i \in \mathbb{Z}} (E[1])_i$ such that:

$$E[1]_i = E_{i+1}$$

and we call it the *desuspension of E* . We obviously have $E[1][-1] = E[-1][1] = E$. By using successively n times the desuspension functor, we can shift a graded vector space by degree $-n$:

$$E[n]_i = E_{i+n}$$

By duality, the degree of the functions on $E[1]$ have their degree increased by $+1$, i.e. the dual space $(E[1]_i)^*$ has degree $-i$, although being isomorphic to $(E_{i+1})^*$, which has degree $-i - 1$. The notation $[1]$ precisely indicates that the degree of the dual spaces have increased by one. Because the degree of the space is odd, formal power series over E terminate and we have the nice following observation:

Proposition 5.21. *Let E be a vector space and let $E[1]$ be its desuspended version. The exterior algebra of E^* is then canonically isomorphic to the algebra of (polynomial) functions of E :*

$$\Lambda^\bullet(E^*) \simeq S(E[1])$$

and it is so that the polynomial degree of an element $u \in \Lambda^\bullet(E^*)$ is precisely its degree as an element of $S(E[1])$.

Example 5.22. The Chevalley-Eilenberg algebra $\Lambda^\bullet(\mathfrak{g}^*)$ of a Lie algebra \mathfrak{g} is the algebra of functions of the graded vector space $\mathfrak{g}[1]$, which is understood as the Lie algebra \mathfrak{g} concentrated in degree -1 .

Example 5.23. By convention, we denote by $T[1]M$ the tangent bundle $TM[1]$ where tangent vectors are considered to carry a degree -1 . The wedge product of the differential forms – Equation (1.21) – is the graded commutative product on the graded algebra $\Omega^\bullet(M)$, where the degree of a differential p -form is p . Example (5.28) shows that it can be understood as the algebra of functions of the *shifted tangent bundle* $T[1]M$.

The natural geometric generalization of graded vector spaces are *graded manifolds*. In differential geometry, an n -dimensional smooth manifold is defined as a topological space which is locally homeomorphic to \mathbb{R}^n , and such that two coordinate charts are smoothly compatible. In super and graded geometry, there is a similar idea (developed by Rogers [Rogers, 1980] and DeWitt [DeWitt, 1984]): a graded manifold is in some sense locally homeomorphic to some product $\mathbb{R}^n \times E$, where E is a graded vector space. However in the literature, graded manifolds are usually and were originally defined from the dual point of view: from the *sheaf of functions* rather than from a set of coordinate charts, even if the two definitions are equivalent. This dual conception goes back to Berezin and Leites for supermanifolds [Berezin and Leites, 1975, Leites, 1980], and has naturally been extended to graded manifolds by Kostant [Kostant, 1977]. The two notions happen to be equivalent, as was shown by Batchelor [Batchelor, 1980]. See [Cattaneo and Schätz, 2011] and [Fairon, 2017] for a comprehensive introduction.

The space of functions on a smooth manifold and its restrictions to open sets can be seen as a sheaf \mathcal{C}^∞ , that is an application from the topology of M taking values in the category of commutative algebras, and satisfying some compatibility conditions over open sets (see [Mac Lane and Moerdijk, 1994] for further details). A topological space X together with a sheaf of rings \mathcal{O} is called a *ringed space*, thus any smooth manifold M in the usual sense is a ringed space, with structure sheaf \mathcal{C}^∞ . Moreover, the smooth case has the additional property that the germs of smooth functions at a given point of M is a *local ring*. A ring is *local* if it admits a unique non-trivial maximal ideal, different from the ring itself. In the present case, the maximal ideal of the ring of germs of smooth functions at a point $x \in M$ is the one generated by the functions vanishing at this point. One wants to mimick that property of germs in graded geometry, so that one additionally requires that the sheaf of functions on a graded manifolds have the property that their stalks are local rings. Since the ring of polynomial over a finite set of variables is a not local, but the ring of formal power series over a finite set of variables is, one usually leans toward the latter to define graded manifolds [Fairon, 2017, Kotov and Salnikov, 2021]. See in particular Remark 2.19 in [Fairon, 2017] which explains that relying on polynomial functions is often too restrictive because it would be an obstacle to differentiability of functions.

Definition 5.24. Let E be a graded vector space and M be a smooth manifold. A graded manifold is a locally ringed space $\mathcal{M} = (M, \mathcal{O}_{\mathcal{M}})$ where the structure sheaf $\mathcal{O}_{\mathcal{M}}$ is locally of the form $\mathcal{C}^{\infty}(U) \otimes \widehat{S}(E^*)$ for some open set $U \subset M$. Then the smooth manifold M is called the base manifold, or body. We say that the graded manifold \mathcal{M} is positively graded (resp. negatively) if the grading of E is strictly negative (resp. positive).

As the space of functions on \mathcal{M} is locally isomorphic to the space of functions on some open set of M tensored with the functions on the graded vector space E , as defined in Definition 5.19, the sheaf $\mathcal{O}_{\mathcal{M}}$ is a sheaf of \mathbb{Z} -graded algebras. When $U = M$, we find that the structure sheaf $\mathcal{O}_{\mathcal{M}}(M)$ would be isomorphic to the space of smooth functions taking values in $\widehat{S}(E^*)$. It is thus tempting to identify E with the fiber of a graded vector bundle over M , that is to say: a vector bundle \mathcal{E} over M , whose fiber is the graded vector space E . That would be helpful because it would enable to work in local coordinates or to consider only graded vector bundles. And indeed this result is a famous theorem of M. Batchelor [Batchelor, 1979] which ensures that we can realize (non canonically) a positively graded manifold as a graded vector bundle over a smooth manifold:

Theorem 5.25. Batchelor (1979) Let $(M, \mathcal{O}_{\mathcal{M}})$ be a positively graded manifold, then there exists a graded vector bundle $\mathcal{E} \rightarrow M$ with fiber a negatively graded vector space E such that the structure sheaf $\mathcal{O}_{\mathcal{M}}$ is locally isomorphic to the sheaf of sections $\Gamma(S(\mathcal{E}^*))$, where by convention $S(\mathcal{E}^*)$ is the graded vector bundle with fiber $S(E^*)$.

Remark 5.26. Here, we do not explicitly need the use of formal power series because we are working only with positively graded manifold, and Remark 3.2 in [Fairon, 2017] explains that the problem raised by absence of locality of the stalks can be lifted. Another way of understanding this is that on a positively graded manifold, the polynomial degree of homogeneous functions has to be bounded (hence the power series is finite) [Kotov and Salnikov, 2021]. In any case, the statement of Theorem 5.25 can be extended to any kind of graded manifold whose grading is bounded below or above. Recent researches have even extended it to more general cases [Kotov and Salnikov, 2021].

This important result stated in Theorem 5.25 allows us to talk of (positively) graded manifolds in terms of (positively) graded vector bundles, which is simpler and more systematic. Using this identification, we will sometimes use the notation $\mathcal{E} \rightarrow M$ instead of \mathcal{M} , to refer to a graded manifold. Such graded manifold is often called *split* but we omit this distinction. Whenever the grading is positive (or negative), the structure sheaf $\mathcal{O}_{\mathcal{M}}$ is isomorphic to the sheaf of sections of the graded vector bundle which fiber is $S(E^*)$. The convention regarding positivity and negativity is chosen so that the name ‘positive graded manifold’ designates a graded manifold whose algebra of functions is positively graded. This emphasizes the prominent role of functions in graded geometry, since they are much easier to handle than vectors, as the following definition shows:

Definition 5.27. A morphism of graded manifolds from \mathcal{M} to \mathcal{N} (with respective base manifolds M and N) is the data of a smooth map $\phi : M \rightarrow N$ that we call the base map together with a morphism of sheaves $\Phi : \mathcal{O}_{\mathcal{N}} \rightarrow \mathcal{O}_{\mathcal{M}}$ over ϕ^* :

$$\Phi(fG) = \phi^*(f)\Phi(G)$$

for every $f \in \mathcal{C}^{\infty}(V)$ and $G \in \mathcal{O}_{\mathcal{N}}(V)$, for any open set $V \subset N$. We say that the morphism Φ covers the base map ϕ .

For convenience and clarity, we will write Φ for a morphism of graded manifold, without further mention of the base map.

Example 5.28. A famous example of graded manifold is the tangent space TM whose fiber degree has been shifted by one: it is denoted by $T[1]M$. Linear coordinate functions on $T[1]M$ are 1-forms on M and are considered to carry a degree $+1$. Since the degrees of the fibers have been shifted, any polynomial of such coordinate functions is a differential form, hence the sheaf of functions is $\mathcal{O}_{T[1]M} = \Omega^\bullet$. The formal power series terminate because the degree of the element generating the functions over $T[1]M$ is 1 (odd).

For a graded manifold \mathcal{M} whose grading is bounded below or above, and whose fiber – by Batchelor’s Theorem 5.25 – is the graded vector space $E = (E_i)_{i \in \mathbb{Z}}$, the dual space of E can be understood as the space of linear coordinate functions on the fiber E . That is to say, if E admits a decomposition into basis elements $(e_{i,p})_{1 \leq p \leq \dim(E_i)}$ of degree $|e_{i,p}| = i$, then we would understand $(e_i^q)_{1 \leq q \leq \dim(E_i)}$ as the dual basis: $e_i^q(e_{i,p}) = \delta_p^q$. Then the $(e_i^q)_{1 \leq q \leq \dim(E_i)}$ are playing the role of linear coordinates on the vector space E_i . Moreover, the body of \mathcal{M} , being a usual smooth manifold M , comes equipped with local coordinates x^j . As for usual differential geometry then, the graded manifold \mathcal{M} admits vector fields which can be understood as graded derivations of the algebra of functions $\mathcal{O}_{\mathcal{M}}$. In particular they can carry a degree.

Example 5.29. Given a graded manifold \mathcal{M} , we define its *Euler vector field* as the unique graded derivation \mathbf{E} of $\mathcal{O}_{\mathcal{M}}$ which, applied to a homogeneous function f , satisfies $\mathbf{E}(f) = |f|f$.

As in the classical smooth case, in coordinates a vector field would involve the dual elements $(e_i^q)_{1 \leq q \leq \dim(E_i)}$ and the local coordinates x^j . A typical example of a vector field on \mathcal{M} would be of the form:

$$X = \sum_{\substack{1 \leq j \leq n \\ 1 \leq m}} X_{a_1 \dots a_m}^j e_{i_1}^{a_1} \wedge \dots \wedge e_{i_m}^{a_m} \frac{\partial}{\partial x^j} + \sum_{i \in \mathbb{Z}} \sum_{\substack{q_i=1 \\ 1 \leq m_i}}^{\dim(E_i)} X_{a_1 \dots a_{m_i}}^{q_i} e_{i_1}^{a_1} \wedge \dots \wedge e_{i_{m_i}}^{a_{m_i}} \frac{\partial}{\partial e_i^{q_i}} \quad (5.48)$$

where $X_{a_1 \dots a_m}^j$ is a smooth function on M and there is implicit summation on contracted indices a_{i_k} . Moreover notice that the sum may not be finite. The degree of the term $X_{a_1 \dots a_m}^j e_{i_1}^{a_1} \wedge \dots \wedge e_{i_m}^{a_m} \frac{\partial}{\partial x^j}$ in the sum (5.48) is computed by counting the total degree of the monomial $e_{i_1}^{a_1} \wedge \dots \wedge e_{i_m}^{a_m}$. However, the degree of the term $X_{a_1 \dots a_{m_i}}^{q_i} e_{i_1}^{a_1} \wedge \dots \wedge e_{i_{m_i}}^{a_{m_i}} \frac{\partial}{\partial e_i^{q_i}}$ involves the degree of the derivative $\frac{\partial}{\partial e_i^{q_i}}$ which is by convention counted as negative (because $e_i^{q_i}$ is at the denominator):

$$|e_{i_1}^{a_1}| + \dots + |e_{i_{m_i}}^{a_{m_i}}| - |e_i^{q_i}| = -i_1 - \dots - i_{m_i} + i$$

There is a minus sign in front of the term $|e_i^{q_i}|$ because the coordinate function $e_i^{q_i}$ appears in the denominator $\frac{\partial}{\partial e_i^{q_i}}$.

A vector field X of homogeneous degree $|X| \in \mathbb{Z}$ acts on a function $f \in \mathcal{O}_{\mathcal{M}}$ in the usual way, through derivation. On a product of two functions $f \wedge g$, we have the graded identity:

$$X(f \wedge g) = X(f) \wedge g + (-1)^{|f||X|} f \wedge X(g)$$

Vector fields admit a *graded Lie bracket*, which is just the generalization of Equation (1.10). On two homogeneous vector fields X, Y it reads:

$$[X, Y] = X \circ Y - (-1)^{|X||Y|} Y \circ X \quad (5.49)$$

The degree of the Lie bracket $[X, Y]$ is by convention $|X| + |Y|$. This turns the space of graded derivations of $\mathcal{O}_{\mathcal{M}}$ – equivalently, the space of vector fields on \mathcal{M} – into a graded Lie algebra (see Definition 3.16).

Equation (5.49) has tremendous consequences: since a vector field can now have odd degree, it is now not necessarily true that the commutator of a vector field with itself vanishes. Indeed, if X has odd degree, then the sign $(-1)^{|X||X|}$ in Equation (5.49) is a minus sign and we have:

$$[X, X] = 2(X \circ X) \tag{5.50}$$

also denoted $2X^2$. There is a priori no reason that the right-hand side vanishes²⁶. For a vector field of even degree, Equation (5.49) is always 0. Since a vector field is a graded derivation of the algebra of functions \mathcal{O}_M , a homogeneous vector field of degree +1 which satisfies $[X, X] = 0$ is by Equation (5.50) a differential on \mathcal{O}_M . This is sufficiently important to be coined in a definition:

Definition 5.30. A differential graded manifold (or Q -manifold) is a graded manifold equipped with a degree +1 vector field Q which commutes with itself: $[Q, Q] = 0$. More generally, odd vector fields X such that $[X, X] = 0$ are called homological vector fields.

Remark 5.31. We speak of an NQ -manifold (or positively graded dg manifold) when the underlying graded manifold – seen as a graded vector bundle – involves only coordinate functions of positive degrees, i.e. when $E_i = 0$ for all $i \geq 0$. Coordinates on the base manifold have degree 0, whereas the fibers admit coordinate functions which are sections of their respective dual spaces, which are then supposed to be of degree greater than or equal to 1.

Example 5.32. The main example is the Lie algebroid (Definition 2.24), whose reformulation as a Q -manifold is due to Vaintrob [Vaintrob, 1997]. It eventually led Voronov to generalize this notion and define the possibly higher Lie algebroids to be Q -manifolds [Voronov, 2010], leading to the precise idea of Lie ∞ -algebroids. Lie algebroids can be defined with symmetric brackets on $\Gamma(A[1])$ instead of skew-symmetric ones on $\Gamma(A)$. On $A[1]$, sections have degree -1 whereas they have degree 0 when seen as sections of A . Thus, given a Lie algebroid A over M , we define on the sections of the suspended vector bundle $A[1]$ the following symmetric bracket:

$$\{x, y\} = [\tilde{x}, \tilde{y}]$$

for any sections x, y of $A[1]$, and where \tilde{x} is the representative of the section x in $\Gamma(A)$ (i.e. whose degree has been shifted by +1).

The space of functions on $A[1]$ is isomorphic to $\Gamma(S(A[1]^*))$ (the formal power series terminate because the degree of the fiber is odd), then it is sufficient to define the vector field Q on the smooth functions on M and on the sections of $A[1]^*$, and then extend to all of $\Gamma(S(A[1]^*))$ by derivation. However since Q is of degree one, then it sends smooth functions to sections of $A[1]^*$ and sections of $A[1]^*$ to sections of $S^2(A[1]^*)$. Hence we define:

$$\begin{aligned} \langle Q[f], x \rangle &= \rho(x)[f] \\ \langle Q[\alpha], x \odot y \rangle &= \rho(x)\langle \alpha, y \rangle - \rho(y)\langle \alpha, x \rangle - \langle \alpha, \{x, y\} \rangle \end{aligned}$$

for every $f \in C^\infty(M)$, $\alpha \in \Gamma(A[1]^*)$, and for any $x, y \in \Gamma(A[1])$, and where $\langle \cdot, \cdot \rangle$ denotes the pairing between A and A^* . We extend Q to all of $\Gamma(S(A[1]^*))$ by derivation, so that the cohomological property comes from the morphism property (2.3) and from the Jacobi identity of the Lie algebroid bracket:

$$\begin{aligned} \langle Q^2[f], x \odot y \rangle &= \left([\rho(x), \rho(y)] - \rho(\{x, y\}) \right) [f] \\ \langle Q^2[\alpha], x \odot y \odot z \rangle &= \left\langle \alpha, \{ \{x, y\}, z \} + \{ \{y, z\}, x \} + \{ \{z, x\}, y \} \right\rangle \end{aligned}$$

²⁶However, notice that the triple commutator $[X, [X, X]]$ always vanishes, whatever the degree of X .

for every $f \in \mathcal{C}^\infty(M)$, $x, y, z \in \Gamma(A[1])$ and $\alpha \in \Gamma(A[1]^*)$. This explains the one-to-one correspondence between Lie algebroids and Q -manifold structures of degree 1 on $A[1]$.

On a Lie algebroid A , the cohomological vector field Q actually turns out to be the Lie algebroid differential d_A defined on the Chevalley-Eilenberg complex $\wedge^\bullet A^*$. This sheds light on other situations:

1. Since TM is a Lie algebroid, the shifted tangent bundle $T[1]M$ is a NQ -manifold, and the cohomological vector field Q is the de Rham differential on the algebra of functions $\Omega^\bullet(M)$.
2. The same idea applies to Poisson manifolds: by Proposition 3.39, a Poisson manifold M induces a Lie algebroid structure on T^*M . Then the cohomological vector field on $T^*M[1]$ is the Poisson differential $d_\pi = [\pi, \cdot]_{SN}$, see Definition 3.18.

The notion of differential graded manifold applies to the extended phase space \mathfrak{P} of Section 5.2:

$$\mathfrak{P} = T^*Q \otimes \wedge^1(\mu_1, \dots, \mu_p) \otimes \wedge^1(u^1, \dots, u^p) \quad (5.51)$$

This is a graded manifold with body T^*Q and fibers of (ghost) degree $+1$ and -1 . The ghosts η^i are the dual coordinates to the μ_i , while the ghost momenta \mathcal{P}_j are dual coordinates to the u^j . Then, what we have defined as the algebra of functions on \mathfrak{P} , which is involving the wedge products of ghosts and ghost momenta:

$$\mathcal{C}^\infty(\mathfrak{P}) = \mathcal{C}^\infty(T^*Q) \otimes \wedge^\bullet(\eta^1, \dots, \eta^p) \otimes \wedge^\bullet(\mathcal{P}_1, \dots, \mathcal{P}_p)$$

is indeed the algebra of functions on \mathfrak{P} as defined in Definition 5.19. The whole goal of Section 5.2 was to define a degree $+1$ differential s on $\mathcal{C}^\infty(\mathfrak{P})$ satisfying Theorem 5.12. As a derivation of the algebra of functions $\mathcal{C}^\infty(\mathfrak{P})$, it can be understood as a vector field on \mathfrak{P} . By Equation (5.50), the cohomological property $s^2 = 0$ is then equivalent to the vanishing of the self commutator: $[s, s] = 0$. This proves that the extended phase space of Section 5.2 is a differential graded manifold, with cohomological vector field s .

The action of the linear maps δ, d and Δ on smooth functions, ghosts and ghost momenta – see Equations (5.42), (5.43) and (5.44) – gives us some informations to what does $s = \delta + d + \Delta + \dots$ look like in local coordinates²⁷:

$$s = \varphi_i \frac{\partial}{\partial \mathcal{P}_i} + \eta^i X_i - C_{ij}^k \eta^j \wedge \mathcal{P}_k \frac{\partial}{\partial \mathcal{P}_i} - \frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \frac{\partial}{\partial \eta^k} + \Delta + \dots \quad (5.52)$$

The first-term is the differential δ while the next three terms correspond to d . Inspired by Remark 5.8, we have the intuition that (at least part of) the differential s is a Hamiltonian vector field on \mathfrak{P} . By Equation (5.27) we know that Δ contains the term $-\frac{1}{2}\{C_{ij}^k, \cdot\} \eta^i \wedge \eta^j \wedge \mathcal{P}_k$, and this is precisely the contribution that is missing to the sum $-C_{ij}^k \eta^j \wedge \mathcal{P}_k \frac{\partial}{\partial \mathcal{P}_i} - \frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \frac{\partial}{\partial \eta^k}$ so that it can be written as $\{-\frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \wedge \mathcal{P}_k, \cdot\}$. Gathering these terms together, Equation (5.52) can be recasted as follows:

$$s = \left\{ \varphi_i \eta^i - \frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \wedge \mathcal{P}_k, \cdot \right\} + \Delta' + \dots \quad (5.53)$$

where the new term Δ' is a derivation of ghost number 1 and pure antighost number 1 defined as:

$$\Delta' = \frac{1}{4} \sigma_{imn}^{kl} \eta^m \wedge \eta^n \wedge \mathcal{P}_k \wedge \mathcal{P}_l \frac{\partial}{\partial \mathcal{P}_i} - \frac{1}{6} \sigma_{imn}^{kl} \eta^i \wedge \eta^m \wedge \eta^n \wedge \mathcal{P}_l \frac{\partial}{\partial \eta^k}$$

²⁷From now on, we will use the wedge product instead of the tensor product between ghosts and ghost momenta, because in the graded geometry setup we now see them as *graded coordinates* which can legitimately graded commute, following Rule (5.47).

where σ_{imn}^{kl} is the tensor defined in Equation (5.36). This new term can be explained as follows: Δ' is what remains of Δ after one had singled out the term $-\frac{1}{2}\{C_{ij}^k, \cdot\}\eta^i \wedge \eta^j \wedge \mathcal{P}_k$ from it. It turns out that this term is almost exact as:

$$\Delta' = \left\{ \frac{1}{12} \sigma_{imn}^{kl} \eta^i \wedge \eta^m \wedge \eta^n \wedge \mathcal{P}_k \wedge \mathcal{P}_l, \cdot \right\} - \frac{1}{12} \{ \sigma_{imn}^{kl}, \cdot \} \eta^i \wedge \eta^m \wedge \eta^n \wedge \mathcal{P}_k \wedge \mathcal{P}_l \quad (5.54)$$

where the last term certainly cancels with a term in $s_{(3)}$. This discussion hints towards the idea of recasting the BRST differential s as a Hamiltonian vector field on the graded manifold \mathfrak{P} . In order to explore this possibility, one needs to discuss the notion of Poisson and symplectic structures in graded geometry.

In order to generalize Poisson and symplectic structures to the graded context, one first needs to allow the notion of graded Lie bracket to carry a degree (the following notions are taken from [Cattaneo et al., 2006]). A *graded Lie algebra of degree n* is a graded vector space A endowed with a graded Lie bracket on $A[n]$, in the sense of Definition 3.16. Such a bracket can be seen as a degree $-n$ Lie bracket on A , i.e. as bilinear operation $[[\cdot, \cdot]] : A \otimes A \rightarrow A$ satisfying $[[A_i, A_j]] \subset A_{i+j-n}$, as well as the graded antisymmetry and graded Jacobi relations:

$$\begin{aligned} [[a, b]] &= -(-1)^{(|a|-n)(|b|-n)} [[b, a]] \\ [[a, [[b, c]]]] &= [[[a, b], c]] + (-1)^{(|a|-n)(|b|-n)} [[b, [[a, c]]]] \end{aligned}$$

Then Definition 3.1 immediately generalizes to the graded case:

Definition 5.33. A graded Poisson algebra of degree n or n -Poisson algebra is a graded vector space $A = \bigoplus_{i \in \mathbb{Z}} A_i$ equipped with two bilinear products \wedge and $\{\cdot, \cdot\}$, such that:

1. (A, \wedge) is a graded commutative algebra;
2. $(A, \{\cdot, \cdot\})$ is a n -graded Lie algebra;
3. the Lie bracket is a graded derivation of the associative product:

$$\{a, b \wedge c\} = \{a, b\} \wedge c + (-1)^{|b|(|a|-n)} b \wedge \{a, c\} \quad (5.55)$$

for any homogeneous elements $a, b, c \in A$.

As a particular case, when $n = 1$, A is said to be a Gerstenhaber algebra.

Remark 5.34. Using the graded commutativity of the graded Poisson bracket, together with Equation (5.55), one can show by induction the following formula:

$$\{a \wedge b, c\} = (-1)^{|b|(|c|-n)} \{a, c\} \wedge b + a \wedge \{b, c\} \quad (5.56)$$

Example 5.35. The algebra of function on the extended phase space \mathfrak{P} of Section 5.2 and denoted $\mathcal{C}^\infty(\mathfrak{P})$ (see Equation (5.21)) is a graded Poisson algebra (of degree 0) when equipped with the Poisson bracket $\{\cdot, \cdot\}$ which restricts to the canonical one on smooth functions on T^*Q , and to the one defined in Equation (5.25) when restricted to ghosts and ghost momenta. The Poisson bracket between a smooth function and a ghost or ghost momentum is defined to be zero:

$$\{f, \eta^i\} = 0 \quad \text{and} \quad \{f, \mathcal{P}_i\} = 0$$

Otherwise it should satisfy the usual derivation property (5.55). It will be given an explicit formula in Proposition 5.46. Notice that this graded Poisson bracket has total ghost number 0 (as said before) but has pure antighost number 1. So it will decrease the pure antighost number of a function $F \in \mathcal{C}^\infty(\mathfrak{P})$.

Example 5.36. The algebra of polyvector fields $\mathfrak{X}^\bullet(M)$ presented in Section 3.1 – not to be confused with the shifted algebra $\mathcal{V}^\bullet(M) = \mathfrak{X}^\bullet(M)[1]$ – is a Gerstenhaber algebra when equipped with the Schouten-Nijenhuis bracket, because it is of degree -1 on $\mathfrak{X}^\bullet(M)$.

From the observation of Example 5.35, mathematicians have managed to generalize the treatment of irreducible first-class constraints which is called *Dirac reduction* in order to apply it to abstract Poisson algebras. This reduction process has been explained in the discussion following Definition 4.82. It has been shown that this reduction is equivalent to Poisson reduction of a Poisson algebra by a coisotropic ideal – also called *Sniatycki-Weinstein reduction* and explained in the discussion preceding Proposition 3.97 (see for example Section 3.1 of [Cattaneo et al., 2006], [Figuroa-O’Farrill and Kimura, 1991b]). Mathematically speaking, Dirac reduction has been explained in Remark 3.98 and in much more details in Section 2 of [Kimura, 1993]. Reformulated in terms of constraints, if I is a multiplicative ideal of a Poisson algebra P generated by a set of functions (the constraints), then the first-class constraints are the elements of I which are also part of the normalizer $N(I)$ of I (with respect to the Poisson bracket):

$$\{\text{first-class constraints}\} = N(I) \cap I$$

We denote this set by I' and the second-class constraints are the remaining generators.

The algebraic analogues of the gauge invariant functions in the sense of Definition 4.82 are the elements f of P such that $\{f, I'\} \in I$. We denote this set of functions by $N(I, I')$. Then, the *Poisson reduction of P by I* – also called *Dirac reduction* in [Blacker et al., 2022] – then essentially consists of taking the quotient of $N(I, I')$ by the ideal I , i.e. taking the quotient of the gauge invariant functions by the ideal of constraints, as was explained in the discussion below Definition 4.82. Proposition 1.23 in Reference [Blacker et al., 2022] also explains under what conditions the Poisson algebra from Sniatycki-Weinstein reduction coincides with that of Dirac reduction.

Moreover, References [Figuroa-O’Farrill and Kimura, 1991b] and [Kimura, 1993] establish another kind of reduction, in the symplectic context, called *homological symplectic reduction*. It is based on the ideas controlling the BRST formalism, and leads to Proposition 5.57. Under adequate circumstances, Sniatycki-Weinstein reduction, Dirac reduction and homological reduction coincide. The paper [Stasheff, 1997] answers the question asked at the end of [Kimura, 1993] and generalizes the BRST-Homological reduction to every Poisson algebras. See Table 1 of [Blacker et al., 2022] for an overview of the various approaches in Poisson and Symplectic reduction.

As the notion of Poisson algebra has been defined, we can now straightforwardly extend most of the material introduced in Section 3 to the graded case. In particular, the notion of graded Poisson manifold or Hamiltonian vector field makes totally sense:

Definition 5.37. *A n -graded Poisson manifold is a graded manifold \mathcal{M} such that the algebra of function $\mathcal{O}_{\mathcal{M}}$ is a n -graded Poisson algebra.*

Example 5.38. Given the discussion in Example 5.35, the extended phase space \mathfrak{P} is a 0-graded Poisson manifold.

As there exist graded Poisson manifolds, there exist also graded symplectic manifold. The definition of the latter relies on the notion of closed non-degenerate two forms (see Definition 3.48). This requires to understand what is a differential form on a graded manifold $\mathcal{M} \approx M \times E$. On a classical smooth manifold, the differential p -forms are generated from the differential one-forms as their exterior algebra. In that context, a differential one-form is a linear coordinate on the tangent bundle, i.e. $dx^i \left(\frac{\partial}{\partial x^j} \right) = \delta_j^i$. We would use the same idea

in the graded case. By convention, coordinate functions on \mathcal{M} are made for one part of the coordinates x^j on the body M , and the basis elements e_i^a of the graded vector space $(E_i)^*$ for every i such that $E_i \neq 0$. Then, we define the differential one-forms, and denote them dx^j and de_i^a , to be linear coordinates on the tangent bundle $T\mathcal{M}$ (see [Cattaneo and Schätz, 2011] for further details and material). The notation d here symbolizes the (graded) de Rham derivative, and should not be confused with the differential modulo δ of Section 5.2.

By convention the elements dx^j and de_i^a have the same degrees – respectively 0 and $-i$ – as the elements x^j and e_i^a (recall that elements of $(E_i)^*$ have degree $-i$). This convention extends to the exterior algebra generated by dx^j and de_i^a . For example, if $\eta = dx^j \wedge de_i^a \wedge de_m^b$ then the degree of η is $|\eta| = -i - m$, while if $\eta = x^k e_r^c \wedge e_s^h dx^j \wedge de_i^a \wedge de_m^b$ then $|\eta| = -r - s - i - m$. In both cases, η is a 3-forms as the form degree (1-form, 2-form, 3-form, etc.) indicates how many arguments the differential form can absorb, while the degree of the differential form (induced by the degrees of basis elements of E^*) is a totally new feature from graded geometry. However, the form degree adds up to the degree to form what we call the *total degree*: in this convention, the elements dx^j and de_i^a have degree $+1$ and $-i + 1$, respectively. This convention is just defined to compute the commutativity of two differential one forms:

$$\begin{aligned} dx^j \wedge dx^k &= (-1)^{1 \times 1} dx^k \wedge dx^j \\ dx^j \wedge de_i^a &= (-1)^{1 \times (-i+1)} de_i^a \wedge dx^j \\ de_i^a \wedge de_j^b &= (-1)^{(-i+1) \times (-j+1)} de_j^b \wedge de_i^a \end{aligned}$$

In particular if i is even, we have that $de_i^a \wedge de_i^a \neq 0$. So much for the conventions on differential forms on a graded manifold. These observations allow us to set the following notion:

Definition 5.39. *A n -graded symplectic manifold is a graded manifold equipped with a non-degenerate closed differential two-form ω of degree n .*

Example 5.40. As in the classical geometry case, the cotangent bundle of a graded manifold is a symplectic manifold of degree 0. Indeed, let $E = E_{-1}$ be a graded vector space concentrated in degree -1 . Linear coordinates on E are abusively denoted q^i and have degree $+1$. Then dq^i have degree $+1$ as well. But linear coordinates on the fiber of the cotangent bundle T^*E , the conjugate momenta associated to q^i , are denoted p_i and have degree -1 so that $p_i(dq^j) = \delta_i^j$ has indeed degree 0 as should be expected. Then the canonical symplectic form $\omega = \sum_{i=1}^n dp_i \wedge dq^i$ has degree 0. This argument applies to any graded manifold.

Now that we have some machinery we can provide some generalization of basic notions from Poisson and symplectic geometry. For example, a n -graded Poisson manifold in which the Poisson bivector is non-degenerate is a symplectic manifold and vis-versa. The Hamiltonian vector fields also admit their counterparts, as well as Poisson and symplectic vector fields. In order to define them we extend to the graded context the notion of Lie derivative. The graded Lie derivative is the obvious, straightforward generalization of Definition 3.33 to graded geometry. In particular, Cartan’s magic formula $\mathcal{L}_X(\eta) = d\iota_X(\eta) + \iota_X d\eta$ becomes:

$$\mathcal{L}_X(\eta) = d\iota_X(\eta) + (-1)^{|X|} \iota_X d\eta \tag{5.57}$$

More generally we can develop Cartan’s calculus along the same lines as in the classical case. In classical differential geometry, on a smooth manifold M , the commutator $[\cdot, \cdot]$ of operators on $\Omega^\bullet(M)$ is a graded commutator, in the sense that d and ι_X are derivations of the algebra of differential forms, of degree $+1$ and -1 , respectively. The degree here denotes the form degree. As the Lie derivative can be written as their commutator: $\mathcal{L}_X = [d, \iota_X]$, we deduce that it is of degree 0 (it does not increase nor decrease the form degree).

Definition 5.41. Let \mathcal{M} be a n -graded Poisson (resp. symplectic) manifold. Denote by $\{.,.\}$ (resp. by ω) the graded Poisson bracket (resp. the symplectic two-form). A vector field X of degree k is said to be Hamiltonian if there exists a function $f \in \mathcal{O}_{\mathcal{M}}$ of degree $n+k$ such that

$$X = \{f, .\} \quad (\text{resp. } df = \iota_X \omega) \quad (5.58)$$

A symplectic vector field on (\mathcal{M}, ω) is a vector field X such that:

$$\mathcal{L}_X(\omega) = 0$$

Remark 5.42. Notice the slight notational difference between the right-hand side of Equation (5.58) and Equation (3.28) in Remark 3.52. In the graded context, the position of the vector field with respect to the symplectic form (to the left or to the right) has indeed consequences on the overall sign.

In the graded case, on a graded manifold \mathcal{M} , we still continue to think of d as a degree $+1$ operator on $\Omega^\bullet(\mathcal{M})$, in the sense that it increases the *total degree* of the differential form by 1 (recall that the total degree of a differential form is the sum of its degree and of its form degree). But ι_X is now understood to be an operator of degree $|X| - 1$. Then Equation (5.59) can make more explicit the definition of the graded commutator in the definition of the Lie derivative:

$$\mathcal{L}_X = [d, \iota_X] \stackrel{\text{def}}{=} d\iota_X - (-1)^{|X|-1}\iota_X d \quad (5.59)$$

From this we deduce that the Lie derivative \mathcal{L}_X has degree $|X|$. Then, the classical relations $[\mathcal{L}_X, d] = 0$ and $\iota_{\mathcal{L}_X(Y)} = [\mathcal{L}_X, \iota_Y]$ straightforwardly extend to the graded context as:

$$[\mathcal{L}_X, d] \stackrel{\text{def}}{=} \mathcal{L}_X d - (-1)^{|X|} d\mathcal{L}_X = 0 \quad (5.60)$$

$$[\mathcal{L}_X, \iota_Y] \stackrel{\text{def}}{=} \mathcal{L}_X \iota_Y - (-1)^{|X|(|Y|-1)} \iota_Y \mathcal{L}_X = \iota_{\mathcal{L}_X(Y)} \quad (5.61)$$

By Equation (5.60), the Euler vector field can now act on differential forms since it has degree 0 so it commutes with the de Rham derivative:

$$\mathcal{L}_{\mathbf{E}}(de_i^a) = d\mathcal{L}_{\mathbf{E}}(e_i^a) = -i de_i^a$$

as $|e_i^a| = -i$. Then we have these fascinating results which are valid only in graded geometry [Cattaneo and Schätz, 2011]:

Lemma 5.43. Let (\mathcal{M}, ω) be a n -graded symplectic manifold.

1. If $n \neq 0$, then ω is exact.
2. Let X be a symplectic vector field of degree k . If $n+k \neq 0$, then X is Hamiltonian.

Proof. The Euler vector field satisfies: $\mathcal{L}_{\mathbf{E}}\omega = n\omega$. Since ω is closed, we are left with $d\omega(\mathbf{E}, .) = n\omega$ where d is the graded de Rham derivative. This implies $\omega = \frac{d\omega(\mathbf{E}, .)}{n}$. In order to prove Item 2., notice that by definition we have:

$$\mathcal{L}_{\mathbf{E}}(X) = kX, \quad \mathcal{L}_{\mathbf{E}}(\omega) = n\omega \quad \text{and} \quad \mathcal{L}_X(\omega) = d\iota_X\omega = 0$$

Since the Euler vector field has degree 0, the graded commutator in Equation (5.61) is a commutator, and we have $\iota_{\mathcal{L}_{\mathbf{E}}(X)} = \mathcal{L}_{\mathbf{E}}\iota_X - \iota_X\mathcal{L}_{\mathbf{E}}$. By setting $H = \iota_{\mathbf{E}}\iota_X\omega$ we obtain:

$$dH = d\iota_{\mathbf{E}}\iota_X\omega = \mathcal{L}_{\mathbf{E}}\iota_X\omega - \iota_{\mathbf{E}}d\iota_X\omega = \iota_{\mathcal{L}_{\mathbf{E}}(X)}\omega + \iota_X\mathcal{L}_{\mathbf{E}}(\omega) = (n+k)\iota_X\omega$$

Hence $\iota_X\omega = d\left(\frac{H}{n+k}\right)$, meaning that X is Hamiltonian. \square

This discussion shows that the extended phase space \mathfrak{P} of Section 5.2 is actually a cotangent bundle of some graded manifold. Indeed let \mathfrak{Q} – this is a gothic capital Q – be the graded manifold defined as the direct product of the configuration space Q and a p -dimensional graded vector space $E = E_{-1}$ concentrated in degree -1 . Then $\mathfrak{Q} = Q \times E$ is a graded manifold of degree -1 whose body is Q : one can then see \mathfrak{Q} as the suspension of a rank p vector bundle over Q . We say that \mathfrak{Q} is the *extended configuration space*. We denote by μ_1, \dots, μ_p a basis of E , so that their dual elements have degree $+1$ and serve as linear coordinates on E : they are the ghosts η^1, \dots, η^p . The momenta conjugate to the ghosts – i.e. the linear coordinate functions on the fibers of $T^*\mathfrak{Q}$ – are precisely the ghost momenta $\mathcal{P}_1, \dots, \mathcal{P}_p$, and given the explanation of Example 5.40 they have degree -1 . Notice that the ghost number precisely encapsulates the grading coming from the graded geometry. Then the extended phase space (5.51) is exactly the cotangent bundle of the extended configuration space:

$$\mathfrak{P} = T^*\mathfrak{Q}$$

Thus, the extended phase space is a graded manifold concentrated in degrees ± 1 , and its body is the cotangent bundle T^*Q . The coordinates on \mathfrak{P} are the q^i, p_j, η^a and \mathcal{P}_b , so that its algebra of functions, in the sense of Definition 5.19 coincides with the original definition, Equation (5.21). Moreover, as a cotangent bundle of a graded manifold, \mathfrak{P} admits a canonical symplectic form :

$$\omega = \sum_{i=1}^n dp_i \wedge dq^i + \sum_{m=1}^p d\mathcal{P}_a \wedge d\eta^a \quad (5.62)$$

whose associated graded Poisson bracket is the one defined in Equation (5.22).

Notice that not only \mathfrak{P} is a graded symplectic manifold, it is also a differential graded manifold as we have proven that there exists a cohomological vector field of degree $+1$ on \mathfrak{P} : the infamous BRST differential s of Section 5.2. A question then would be whether the homological graded vector field s is a *symplectic vector field*, i.e. do we have $\mathcal{L}_s(\omega) = 0$? Graded manifolds for which there is a compatibility between a homological vector field and a graded symplectic form are interesting on their own:

Definition 5.44. *A graded symplectic manifold endowed with a symplectic cohomological vector field is called a differential graded symplectic manifold, or QP-manifold for short.*

If \mathfrak{P} is a QP-manifold, i.e. if s is a symplectic vector field, then since ω has degree 0 and s has degree (ghost number) 1, Item 2. of Lemma 5.43 would imply that s is a Hamiltonian vector field, meaning that there exists a function Ω of ghost number 1 such that $s = \{\Omega, \cdot\}$. From the discussion leading to Equation (5.53), we already know that the first terms of the graded vector field s correspond to a Hamiltonian vector field $X = \left\{ \varphi_i \eta^i - \frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \wedge \mathcal{P}_k, \cdot \right\}$. This represents a huge indication that the BRST differential is a potential Hamiltonian vector field on the extended phase space \mathfrak{P} , and is such that the associated ‘Hamiltonian function’ Ω would start with the following terms:

$$\Omega = \varphi_i \eta^i - \frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \wedge \mathcal{P}_k + \dots$$

Given its centrality in the treatment and the quantization of physical classical systems in the BRST formalism, the graded function Ω deserves its own name:

Definition 5.45. *Any function $\Omega \in \mathcal{C}^\infty(\mathfrak{P})$ of ghost number 1 satisfying the following identity:*

$$s = \{\Omega, \cdot\} \quad (5.63)$$

is called a BRST charge or BRST generator. The cohomological property of the BRST differential s translates at the level of the BRST charge as:

$$\{\Omega, \Omega\} = 0 \quad (5.64)$$

This identity is called the classical master equation.

Notice that the classical master equation (5.64) is *not* a trivial equation. The best way to view it is to find a concrete formula for the canonical Poisson bracket on \mathfrak{P} (see Example 5.35). Moreover, it would help us making sense of Equations (5.63) in the graded case, as we are interested in writing $s = \{\Omega, \cdot\}$ in terms of derivatives of ghosts and ghost momenta. From Equations (5.56) and (5.25) one deduces that:

$$\begin{aligned} \{\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}, G\} &= (-1)^{m-1+n} \frac{\partial}{\partial \eta^{i_k}} (\eta^{i_1} \wedge \dots \wedge \eta^{i_m}) \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n} \frac{\partial G}{\partial \mathcal{P}_{i_k}} \\ &+ (-1)^{n-1} \eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \frac{\partial}{\partial \mathcal{P}_{i_k}} (\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \frac{\partial G}{\partial \eta^{i_k}} \end{aligned} \quad (5.65)$$

The signs in this equation are a consequence of the odd ghost number ± 1 of the ghosts and the ghost momenta. One can understand the sign $(-1)^{m-1+n}$ of the first term on the right-hand side as coming from a ‘right derivative’, i.e. a derivative – denoted $\overleftarrow{\partial}$ – that differentiates a function from the right (so we write it on the very left):

$$(\eta^{i_1} \wedge \dots \wedge \eta^{i_m}) \overleftarrow{\frac{\partial}{\partial \eta^{i_k}}} \equiv m \eta^{[i_1} \wedge \dots \wedge \eta^{i_{m-1}} \delta_{i_k}^{i_m]} \quad (5.66)$$

where the brackets symbolizes the full antisymmetry on the upper indices. By this antisymmetry, the right-hand side of Equation (5.66) is equal to $(-1)^{m-1} m \eta^{[i_2} \wedge \dots \wedge \eta^{i_m} \delta_{i_k}^{i_1]}$, so that one has:

$$(\eta^{i_1} \wedge \dots \wedge \eta^{i_m}) \overleftarrow{\frac{\partial}{\partial \eta^{i_k}}} = (-1)^{m-1} \frac{\partial}{\partial \eta^{i_k}} (\eta^{i_1} \wedge \dots \wedge \eta^{i_m}) \quad (5.67)$$

In other words, the first term on the right-hand side of Equation (5.65) can be rewritten as:

$$(-1)^{m-1+n} \frac{\partial}{\partial \eta^{i_k}} (\eta^{i_1} \wedge \dots \wedge \eta^{i_m}) \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n} \frac{\partial G}{\partial \mathcal{P}_{i_k}} = (\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \overleftarrow{\frac{\partial}{\partial \eta^{i_k}}} \frac{\partial G}{\partial \mathcal{P}_{i_k}}$$

because one has to pass the right-derivative over all the ghost momenta (bringing a sign $(-1)^n$) and then apply Equation (5.67).

The second sign $(-1)^{n-1}$ on the right-hand side of Equation (5.65) can be interpreted along the same lines. With a right-derivative, the second term on the right-hand side of Equation (5.65) can be rewritten as:

$$(-1)^{n-1} \eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \frac{\partial}{\partial \mathcal{P}_{i_k}} (\mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \frac{\partial G}{\partial \eta^{i_k}} = (\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}) \overleftarrow{\frac{\partial}{\partial \mathcal{P}_{i_k}}} \frac{\partial G}{\partial \eta^{i_k}}$$

Since the right-derivative does have to pass over the ghosts, we do not have any dependency in their number m . The advantage of using right-derivatives is the absence of signs, that were otherwise present in Equation (5.65) and possibly intriguing. If one denotes the usual partial

derivative (from the left) with a right arrow on the top, the canonical graded Poisson bracket on the extended phase space then reads:

$$\{F, G\} = \sum_{i=1}^n \frac{\partial F}{\partial q^i} \frac{\partial G}{\partial p_i} - \frac{\partial F}{\partial p_i} \frac{\partial G}{\partial q^i} + \sum_{a=1}^p F \frac{\overleftarrow{\partial}}{\partial \eta^a} \frac{\overrightarrow{\partial}}{\partial \mathcal{P}_a} G + F \frac{\overleftarrow{\partial}}{\partial \mathcal{P}_a} \frac{\overrightarrow{\partial}}{\partial \eta^a} G \quad (5.68)$$

The plus sign between the two last terms can be interpreted as twice a minus sign: the first one would be the one usually appearing in a Poisson bracket (e.g. in the first sum), while the second one would be induced by the anticommutativity of the two derivatives which carry an odd ghost number, see e.g. Equation (5.23).

We can also provide a formula of the graded Poisson bracket which does not involve such derivatives from the right. Indeed, by making $\frac{\partial}{\partial \mathcal{P}_{i_k}}$ passing over $\eta^{i_1} \wedge \dots \wedge \eta^{i_m}$ from the right, Equation (5.65) can be rewritten as follows:

$$\begin{aligned} \{\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}, G\} &= (-1)^{m-1+n} \frac{\partial}{\partial \eta^{i_k}} \left(\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n} \right) \frac{\partial G}{\partial \mathcal{P}_{i_k}} \\ &+ (-1)^{n-1+m} \frac{\partial}{\partial \mathcal{P}_{i_k}} \left(\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n} \right) \frac{\partial G}{\partial \eta^{i_k}} \end{aligned} \quad (5.69)$$

Since $n + m$ corresponds to the ghost number of the function $\eta^{i_1} \wedge \dots \wedge \eta^{i_m} \wedge \mathcal{P}_{j_1} \wedge \dots \wedge \mathcal{P}_{j_n}$, we straightforwardly extend Equation (5.69) to more general functions on \mathfrak{P} :

Proposition 5.46. *The canonical graded Poisson bracket on $C^\infty(\mathfrak{P})$ associated to the canonical symplectic structure (5.62) on the extended phase space is defined by:*

$$\{F, G\} = \sum_{i=1}^n \frac{\partial F}{\partial q^i} \frac{\partial G}{\partial p_i} - \frac{\partial F}{\partial p_i} \frac{\partial G}{\partial q^i} - (-1)^{\text{gh}(F)} \sum_{a=1}^p \frac{\partial F}{\partial \eta^a} \frac{\partial G}{\partial \mathcal{P}_a} + \frac{\partial F}{\partial \mathcal{P}_a} \frac{\partial G}{\partial \eta^a} \quad (5.70)$$

where F is a homogeneous function of ghost number $\text{gh}(F)$.

Remark 5.47. This graded Poisson bracket is often called the *big bracket* [Kosmann-Schwarzbach and Rubtsov, 2010].

Equation (5.70) is intriguing because it involves the ghost number of F and a minus sign, although we understand their justification from the discussion leading to Equation (5.68). Both of these equations are important as they can be used in different contexts. For example, the advantage of Equation (5.68) is that there are no sign depending on the ghost numbers of F or G while its disadvantage is that it involves right-derivatives. On the other hand, Equation (5.70) has the disadvantage of involving the ghost number of the function F , but it has the advantage of allowing concrete computations, e.g. by making clear what is the Hamiltonian vector field associated to F :

$$\{F, \cdot\} = \sum_{i=1}^n \frac{\partial F}{\partial q^i} \frac{\partial}{\partial p_j} - \frac{\partial F}{\partial p_j} \frac{\partial}{\partial q^i} - (-1)^{\text{gh}(F)} \sum_{a=1}^p \frac{\partial F}{\partial \eta^a} \frac{\partial}{\partial \mathcal{P}_a} + \frac{\partial F}{\partial \mathcal{P}_a} \frac{\partial}{\partial \eta^a} \quad (5.71)$$

Notice that both Equations (5.68) and (5.70) are independent from the way we write F and G (in which order we organize the ghosts and the ghost momenta).

We can now apply these observations and Proposition 5.46 to the equations characterizing the BRST charge (see Definition 5.45). First, from Equation (5.70), we understand that the classical master equation (5.64) is *not* a trivial equation, since Ω has ghost number 1:

$$\{\Omega, \Omega\} = \frac{\partial \Omega}{\partial q^i} \frac{\partial \Omega}{\partial p_i} - \frac{\partial \Omega}{\partial p_i} \frac{\partial \Omega}{\partial q^i} + \frac{\partial \Omega}{\partial \eta^a} \frac{\partial \Omega}{\partial \mathcal{P}_a} + \frac{\partial \Omega}{\partial \mathcal{P}_a} \frac{\partial \Omega}{\partial \eta^a} = 2 \left(\frac{\partial \Omega}{\partial q^i} \frac{\partial \Omega}{\partial p_i} + \frac{\partial \Omega}{\partial \eta^a} \frac{\partial \Omega}{\partial \mathcal{P}_a} \right) \quad (5.72)$$

We used the fact that $\frac{\partial\Omega}{\partial p_i} \frac{\partial\Omega}{\partial q^i} = -\frac{\partial\Omega}{\partial q^i} \frac{\partial\Omega}{\partial p_i}$ because both terms have ghost number 1, while $\frac{\partial\Omega}{\partial \mathcal{P}_a} \frac{\partial\Omega}{\partial \eta^a} = \frac{\partial\Omega}{\partial \eta^a} \frac{\partial\Omega}{\partial \mathcal{P}_a}$ because here, they have even ghost numbers. Equation (5.72) is obviously not necessarily zero. Moreover, from Equation (5.71), we deduce that the BRST differential in Equation (5.63) can be understood as follows:

$$s = \frac{\partial\Omega}{\partial q^j} \frac{\partial}{\partial p_j} - \frac{\partial\Omega}{\partial p_j} \frac{\partial}{\partial q^j} + \frac{\partial\Omega}{\partial \eta^a} \frac{\partial}{\partial \mathcal{P}_a} + \frac{\partial\Omega}{\partial \mathcal{P}_a} \frac{\partial}{\partial \eta^a} \quad (5.73)$$

If we assume that the BRST differential is a Hamiltonian vector field, we read from Equations (5.53) and (5.54) the first terms in the BRST charge:

$$\Omega = \varphi_i \eta^i - \frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \wedge \mathcal{P}_k + \frac{1}{12} \sigma_{imn}^{kl} \eta^i \wedge \eta^m \wedge \eta^n \wedge \mathcal{P}_k \wedge \mathcal{P}_l + \dots \quad (5.74)$$

Notice that the first term has pure antighost number 0, the second term has pure antighost number 1, and we expect the higher terms to have higher pure antighost numbers. As the BRST differential s can be decomposed by pure antighost number – see Equation (5.46) – we can then decompose the BRST charge by pure antighost numbers:

$$\Omega = \sum_{i=0}^{p+1} \Omega_{(i)} \quad (5.75)$$

with $\Omega_{(0)} = \varphi_i \eta^i$, $\Omega_{(1)} = -\frac{1}{2} C_{ij}^k \eta^i \wedge \eta^j \wedge \mathcal{P}_k$ and $\Omega_{(2)} = \frac{1}{12} \sigma_{imn}^{kl} \eta^i \wedge \eta^m \wedge \eta^n \wedge \mathcal{P}_k \wedge \mathcal{P}_l$. The higher order terms encode how complicated the algebra of gauge symmetries is, e.g. see [Browning and McMullan, 1987] and subsection 9.4.3 in [Henneaux and Teitelboim, 1992]. This decomposition has an intrinsic importance because usually physicists and mathematicians show the existence of the BRST charge Ω by precisely building the sequence of terms $\Omega_{(i)}$ one after another:

Proposition 5.48. *The BRST differential s is a Hamiltonian vector field with respect to the canonical graded symplectic form (5.62) on the extended phase space \mathfrak{P} .*

Proof. See e.g. Section 9.3 in [Henneaux and Teitelboim, 1992] or Section 3.3 in [Schätz, 2008]. \square

The constructions of the BRST charge made in the proof of Proposition 5.48 are all similar and can be formalized under the theory of *homological perturbation* in pure mathematics (see [Tătar and Tătar, 1994] for a clear understanding of the statement, together with references therein for a complete overview of the homological perturbation theory related to BRST formalism in the mid-1990's). While in homological perturbation theory, mathematicians build the BRST differential directly, in general physicists historically built the BRST charge Ω step by step, using the exactness property of the Koszul complex. The definition of the component $\Omega_{(i)}$ is set only up to exact terms, so there is always some liberty in choosing $\Omega_{(i)}$. Then, a priori there exists an infinite number of BRST charges associated to the BRST differential. Moreover the choice of constraints is not unique, although the constraint surface is, adding possibly more liberty in the choice of even the first term $\Omega_{(0)}$ and thus all subsequent ones. However, it turns out that the BRST charge is essentially unique:

Proposition 5.49. *The BRST charge $\Omega \in C^\infty(\mathfrak{P})$ associated to the BRST differential s is unique up to canonical transformations of the extended phase space \mathfrak{P} .*

Proof. This statement is proven in subsection 9.3.3 in [Henneaux and Teitelboim, 1992]. \square

Remark 5.50. We conclude that the extended phase space \mathfrak{P} equipped with its canonical graded symplectic form (5.62) and the BRST differential s is a QP -manifold. Equivalently, since the symplectic form corresponds to the non-degenerate Poisson bracket defined on the algebra of functions $C^\infty(\mathfrak{P})$ in Example 5.35, we deduce that $C^\infty(\mathfrak{P})$ is a differential graded Poisson algebra.

Example 5.51. When the gauge transformations are generated by a true Lie algebra, the BRST differential is only made of two terms $s = \delta + d$ (see Remark 5.15). Then, the BRST charge only contains the first two terms appearing in Equation (5.74):

$$\Omega = \varphi_a \eta^a - \frac{1}{2} C_{ab}^c \eta^a \wedge \eta^b \wedge \mathcal{P}_c$$

and the C_{ab}^c are constant. For more examples related to (possibly abelian) Lie algebras, see Section 9.4 of [Henneaux and Teitelboim, 1992].

Example 5.52. When the gauge transformations form a closed algebra in the sense of Definition 4.69, we will see that the BRST charge of Example 5.51 is not modified, except that the C_{ab}^c become smooth functions over Q . Since the constraints are irreducible, Theorem 4.79 establishes that the almost Lie algebroid E associated to the constraints by Proposition 4.77 is a foliation Lie algebroid, i.e. the vector bundle E fits in the following short exact sequence:

$$0 \longrightarrow E \xrightarrow{\rho} TQ \longrightarrow TQ/\rho(E) \longrightarrow 0$$

By denoting by X_a the Hamiltonian vector field on T^*Q associated to the constraint φ_a (not to be confused with the vector fields $\rho(e_a)$, defined on Q), we have in local coordinates:

$$X_a = \rho_a^i \frac{\partial}{\partial x^i} - p_j \frac{\partial \rho_a^j}{\partial x^k} \frac{\partial}{\partial p_k} \quad (5.76)$$

Then, using Equation (4.89), one finds that the commutator of two such Hamiltonian vector fields does not close:

$$[X_a, X_b] = C_{ab}^c X_c + \varphi_c \frac{\partial C_{ab}^c}{\partial x^i} \frac{\partial}{\partial p_i} \quad (5.77)$$

As expected, the Hamiltonian vector fields X_a close on the constraint surface, when $\varphi_c = 0$.

The shifted Lie algebroid $E[1]$ can be identified with the extended configuration space \mathfrak{Q} and $T^*E[1]$ with the extended phase space \mathfrak{P} . Let us compute the components of the BRST differential in the case where the system of constraints corresponds to a closed algebra of gauge transformations. From Equations (5.42), (5.43) and (5.44) we deduce that $\Delta(f) = -\frac{1}{2} \frac{\partial C_{ab}^c}{\partial x^i} \frac{\partial f}{\partial p_i} \eta^a \wedge \eta^b \otimes \mathcal{P}_c$ while, by Equation (4.84), we deduce that $\Delta(\eta^k) = 0$ and $\Delta(\mathcal{P}_i) = 0$. From this, we obtain that

$$\Delta = -\frac{1}{2} \frac{\partial C_{ab}^c}{\partial x^i} \eta^a \wedge \eta^b \wedge \mathcal{P}_c \frac{\partial}{\partial p_i}$$

It does not contain any other term. The form of Δ for systems of constraints admitting a closed algebra of gauge transformations is thus of a very particular form.

From this discussion, one observes that under the present assumptions, Equation (5.53) only contains the terms which are in the Poisson bracket, that is to say, the BRST charge for irreducible systems of constraints associated to a closed gauge algebra is:

$$\Omega = \varphi_a \eta^a - \frac{1}{2} C_{ab}^c \eta^a \wedge \eta^b \wedge \mathcal{P}_c \quad (5.78)$$

This is a generalization of what we saw in Example 5.51, but now the C_{ab}^c are smooth functions on Q . The fact that the BRST charge has only two terms concentrated in pure antighost numbers

0 and +1 is typical of irreducible constraints forming a closed algebra of gauge transformations. Then, using Equation (5.52) or Equation (5.73), the BRST differential s reads:

$$s = \varphi_a \frac{\partial}{\partial \mathcal{P}_a} + \eta^a X_a - C_{ab}^c \eta^b \wedge \mathcal{P}_c \frac{\partial}{\partial \mathcal{P}_a} - \frac{1}{2} C_{ab}^c \eta^a \wedge \eta^b \frac{\partial}{\partial \eta^c} - \frac{1}{2} \frac{\partial C_{ab}^c}{\partial x^i} \eta^a \wedge \eta^b \wedge \mathcal{P}_c \frac{\partial}{\partial p_i} \quad (5.79)$$

The first term corresponds to the differential δ , while the next three terms correspond to d , the differential modular δ , and the last term indeed corresponds to Δ . Notice that with these data, on the constraint surface, we do not have $d^2(\mathcal{P}_a) \approx 0$.

Eventually, one sees that on the zero section of $\mathfrak{P} = T^*E[1]$, which is defined by the equations $\mathcal{P}_a = 0$ and $p_i = 0$ – and hence $\varphi_a = 0$ there as well – and which is also by definition isomorphic to $E[1]$, the BRST differential s given in Equation (5.79) reduces to the cohomological vector field Q of Example 5.32. The relationship between this cohomological vector field and the BRST charge (5.78) is straightforward, as one only needs to replace \mathcal{P}_c by $\frac{\partial}{\partial \eta^c}$:

$$Q = \eta^a \rho_a^i \frac{\partial}{\partial x^i} - \frac{1}{2} C_{ab}^c \eta^a \wedge \eta^b \frac{\partial}{\partial \eta^c}$$

Moreover, on the zero section of T^*Q , defined by the set of equations $p_i = 0$ and identified with Q , the Hamiltonian vector fields X_a defined in Equation (5.76) coincide with $\rho(e_a)$, and the last term of Equation (5.77) vanishes, so that we obtain the usual morphism identity (2.3). Thus, we can say that φ_a is the Hamiltonian lift of $\rho(e_a) \in \mathfrak{X}(Q)$ to T^*Q and that the BRST charge Ω is the Hamiltonian lift of the cohomological vector field $Q \in \mathfrak{X}(E[1])$ to $T^*E[1]$. All of this discussion shows that all the geometric data of the Lie algebroid E is in some sort contained inside the extended phase space $\mathfrak{P} = T^*E[1]$, as its zero section. This adds up to the already quite deep relationship existing between Lie algebroids and constrained systems, as was first established in Theorem 4.79.

Remark 5.53. The extended phase space in Example 5.52 is what is called a *double vector bundle* over Q , as it fits into the following commutative square:

$$\begin{array}{ccc} T^*E & \longrightarrow & E \\ \downarrow & & \downarrow \\ T^*Q & \longrightarrow & Q \end{array}$$

See [Mackenzie, 1992] for more informations on this particular kind of geometric structures.

5.4 BRST formalism for reducible constraints

In Section 5.2 we introduced the BRST formalism for an irreducible set of (first-class) constraints. This choice was made for pedagogical reasons since understanding the BRST formalism in the irreducible case is already not trivial. We will indeed see that reducibility of the constraints implies that the cost of preserving exactness of the Koszul complex is to introduce other spaces, making the complex much bigger (both in length and in width).

In Section 4, we defined irreducibility of the constraints as the property of being minimally functionally independent (see Definition 4.42). This can be alternatively stated as the fact that these constraints can serve as transverse coordinates to the constraint surface. By the regularity conditions on both the primary and secondary constraints surfaces – see Scholie 4.19 and 4.37 –

we deduce that locally, one can always find a set of independent, i.e. irreducible, constraints, on which the rest of the constraints would depend. However, this set of local constraints might not be sufficient to define the constraint surface globally, as they can be dependent elsewhere in the phase space so they may not be convenient to handle the physics on the constraint surface, e.g. because of manifest symmetry breaking or topological obstructions that cannot be probed by local constraints. Moreover, even if such a global splitting between independent constraints and dependent ones would be possible, it spoils in practice manifest Lorentz invariance or locality in physical space (see Footnote 25). For these reasons, it is sometimes better to keep working with a set of (globally defined) constraints which are functionally dependent, without assuming that any definite splitting has been performed. We then say that the constraints are *reducible* (see Definition 4.42). See subsection 1.1.8 in [Henneaux and Teitelboim, 1992] for additional reasons about the need to use a set of reducible constraints.

A set of constraints being reducible means that there are more constraints than the codimension of the constraint surface Σ , and that locally, in the neighborhood of every point of the constraint surface, we can pick up $\text{codim}(\Sigma)$ constraints in the set such that all the other constraints are consequences of the former. See subsection 1.1.2 in [Henneaux and Teitelboim, 1992] for an understanding of this perspective. An alternative but equivalent reformulation involves the null eigenvectors of the matrix controlling the Noether identities: see footnote 4 on page 21 in [Rothe and Rothe, 2010], or Section 2 in [Gomis et al., 1995] as well as subsection 3.1.9 of [Henneaux and Teitelboim, 1992]. Eventually, the splitting of constraints into first-class and second-class constraints also interacts with the reducibility property, as is explained in subsection 1.3.4 of [Henneaux and Teitelboim, 1992]. What is said in Section 4.5 still applies to reducible constraints. In particular, first-class reducible constraints induce a regular foliation on the constraint surface, with the difference that the hamiltonian vector fields X_{φ_a} are not independent anymore.

From now on, we assume in the present section that the constraints are globally defined on the phase space, and that they are reducible and first-class. Let us set some conventions: we use the index a_0 to label the irreducible constraints φ_{a_0} and we assume that there are $A_0 \geq \text{codim}(\Sigma)$ of them. Since these constraints are not functionally independent, there exists a set of smooth functions $Z_{a_1}^{(1)a_0}$ on T^*Q which do not vanish everywhere on Σ and such that we have, for every a_1 (this is a rewriting of Equation (5.80) with the current convention of indices):

$$Z_{a_1}^{(1)a_0} \varphi_{a_0} = 0 \quad (5.80)$$

The index a_1 runs from 1 to some integer A_1 , which may be later specified. The functions $Z_{a_1}^{(1)a_0}$ are called *first-stage reducibility functions*²⁸, not to be confused with the first-stage, second-stage, third-stage constraints that were introduced in the Bergmann-Dirac algorithm of Section 4.3. The main property of the functions $Z_{a_1}^{(1)a_0}$ is that they exhaust the dependence relations between the constraints (at least on the constraint surface), i.e. if there are other functions λ^{a_0} such that $\lambda^{a_0} \varphi_{a_0} = 0$, then the latter are functionally dependent on the former: $\lambda^{a_0} \approx f^{a_1} Z_{a_1}^{(1)a_0}$. By Corollary 4.25, this is equivalent to writing, as strong equalities:

$$\lambda^{a_0} \varphi_{a_0} = 0 \quad \text{implies that} \quad \lambda^{a_0} = f^{a_1} Z_{a_1}^{(1)a_0} + g^{a_0 b_0} \varphi_{b_0} \quad (5.81)$$

with $g^{a_0 b_0} = -g^{b_0 a_0}$ so that $g^{a_0 b_0} \varphi_{a_0} \varphi_{b_0}$ is automatically zero by symmetry (see Theorem 10.1 in [Henneaux and Teitelboim, 1992]). The strong equality (5.81) additionally illustrates that there is always an ambiguity in choosing the functions $Z_{a_1}^{(1)a_0}$ because one can always add a combination of constraints to $Z_{a_1}^{(1)a_0}$ without changing Equation (5.80). Then, the value of the

²⁸The denomination is taken from [Gomis et al., 1995]; one can also talk about *order of reducibility*, as in [Henneaux and Teitelboim, 1992].

functions $Z_{a_1}^{(1)a_0}$ outside the constraint surface is not important, but we use strong equations because it is more convenient to work globally. The argument that the relevant content of the reducibility functions is only defined on the constraint surface can be made more explicit by using the hamiltonian vector fields X_{a_0} instead of the constraints in Equation (5.80), as is done in Section 10.2 of [Henneaux and Teitelboim, 1992].

Now we have two situations: either the first-stage reducibility functions $Z_{a_1}^{(1)a_0}$ form an independent set of functions – in which case we say that the theory is *first-stage reducible*, or they are functionally dependent. In the first case, by Scholie 4.37 it means that among the set of A_0 reducible constraints, there are $\text{codim}(\Sigma)$ independent ones, at least locally. Then there are $A_0 - \text{codim}(\Sigma)$ redundant functions and that the number of first-stage (independent) reducibility functions $Z_{a_1}^{(1)a_0}$ is $A_1 = A_0 - \text{codim}(\Sigma)$. Assuming a similar regularity condition for reducibility functions as in Scholie 4.37, it is always possible locally to find a set of independent first-stage reducibility functions but, as for the constraints themselves, it is sometimes more convenient to preserve some dependence between the $Z_{a_1}^{(1)a_0}$. In such a case, it means that there exists a set of smooth functions $Z_{a_2}^{(2)a_1} \in \mathcal{C}^\infty(T^*Q)$ – called *second-stage reducibility functions* – which do not uniformly vanish on Σ and such that:

$$Z_{a_2}^{(2)a_1} Z_{a_1}^{(1)a_0} \approx 0 \quad (5.82)$$

for every integers a_0 and a_2 , where a_2 runs from 1 to some integer A_2 , which may be later specified. Moreover, if one has non trivial second-stage reducibility constraints, then $A_1 > A_0 - \text{codim}(\Sigma)$, where $A_0 - \text{codim}(\Sigma)$ denotes the number of functionally independent first-stage reducibility functions²⁹. The same properties introduced earlier about exhaustion of the dependence and non-uniqueness of the functions $Z_{a_2}^{(2)a_1}$ apply here as well, meaning that:

$$\lambda_1^a Z_{a_1}^{(1)a_0} \approx 0 \quad \text{implies that} \quad \lambda^{a_1} = f^{a_2} Z_{a_2}^{(2)a_1} + g^{a_1 a_0} \varphi_{a_0} \quad (5.83)$$

for some smooth functions $f^{a_2}, g^{a_1 a_0} \in \mathcal{C}^\infty(T^*Q)$.

Here again, we have another embranchement: either the functions $Z_{a_2}^{(2)a_1}$ are functionally independent or they are not. In the first case, we say that the theory is *second-stage reducible*, and we deduce that the number of second-stage reducibility functions is $A_2 = A_1 - A_0 + \text{codim}(\Sigma)$. Notice that we have an alternating sum, as is usual in this kind of situations where we are confronted with a hierarchy of dependent algebraic equations. If the functions $Z_{a_2}^{(2)a_1}$ are functionally dependent, then $A_2 > A_1 - A_0 + \text{codim}(\Sigma)$, where $A_1 - A_0 + \text{codim}(\Sigma)$ denotes the number of functionally independent second-stage irreducibility functions. We then need to introduce a further set of *third-stage reducibility functions* $Z_{a_3}^{(3)a_2} \in \mathcal{C}^\infty(T^*Q)$ satisfying an equation similar to Equation (5.82). If these functions are independent then there are $A_3 = A_2 - A_1 + A_0 - \text{codim}(\Sigma)$ of them, if not we need to introduce *4-th stage reducibility functions*, etc. Eventually, we end up with a (possibly infinite) hierarchy of k -th stage reducibility functions $Z_{a_k}^{(k)a_{k-1}} \in \mathcal{C}^\infty(T^*Q)$ satisfying:

$$Z_{a_{k+1}}^{(k+1)a_k} Z_{a_k}^{(k)a_{k-1}} \approx 0 \quad (5.84)$$

If the hierarchy of functions stops at level L – including $L = \infty$ – we say that the theory is *L -th stage reducible*. The number of independent k -th reducibility functions is $\sum_{l=-1}^{k-1} (-1)^{k-1-l} A_l$, where by convention we have set $A_{-1} = \text{codim}(\Sigma)$. More details on this reasoning can be found in Section 10.2 of [Henneaux and Teitelboim, 1992] or Section 2.2 of [Gomis et al., 1995]. The case of irreducible constraints is obtained for $L = 0$.

As is explained in subsection 10.3.1 of [Henneaux and Teitelboim, 1992], for gauge theories whose stage of reducibility is $L \geq 1$, the Koszul complex introduced in Section 5.2:

²⁹Here, ‘functional independence’ should be understood in the minimal sense, as in Definition 4.42.

$$0 \longrightarrow K_{-p} \xrightarrow{\delta} K_{-p+1} \xrightarrow{\delta} \dots \xrightarrow{\delta} K_{-1} \xrightarrow{\delta} \mathcal{C}^\infty(T^*Q) \longrightarrow 0$$

does not define a resolution of $\mathcal{C}^\infty(\Sigma)$ anymore. Indeed, by Equation (5.80) the element $Z_{a_1}^{(1)b_0} \mathcal{P}_{b_0}$ is δ -closed but it cannot be δ -exact, as it would imply that it comes from a term like $h^{a_0 b_0} \mathcal{P}_{a_0} \wedge \mathcal{P}_{b_0}$, which in turn implies that $Z_{a_1}^{(1)b_0} = 2h^{a_0 b_0} \varphi_{a_0}$, which is not impossible because the functions $Z_{a_1}^{(1)b_0} = 2h^{a_0 b_0} \varphi_{a_0}$ are not supposed to vanish on the constraint surface. Thus, one needs to introduce a new space for each stage of reducibility. Let $K^{(1)} = \mathbb{R}^{A_1}$ and let us denote by $\mathcal{P}_1^{(1)}, \dots, \mathcal{P}_{A_1}^{(1)}$ its standard basis. We call these variables *ghost of ghost momenta* and we assign two different grading to them which generalize that of ghost momenta in Section 5.2: a ghost number of -2 and a pure antighost number of 2 .

In order to provide a more symmetric treatment of every (ghost of) ghost momenta, from now on we will refer to the ghost momenta \mathcal{P}_{a_0} as $\mathcal{P}_{a_0}^{(0)}$, and the corresponding vector space \mathbb{R}^{A_0} that they span as $K^{(0)}$. We then define $T_{-1} = K_{-1} = \mathcal{C}^\infty(T^*Q) \otimes K^{(0)}$ and $T_{-2} = \mathcal{C}^\infty(T^*Q) \otimes (\wedge^2 K^{(0)} \oplus K^{(1)})$; the latter is a vector space of ghost number -2 , as the lower index indicates. Then, one extends the map δ to T_{-2} by defining its action on $\mathcal{C}^\infty(T^*Q) \otimes K^{(1)}$:

$$\delta(\mathcal{P}_{a_1}^{(1)}) = Z_{a_1}^{(1)a_0} \mathcal{P}_{a_0} \quad (5.85)$$

The map δ is still of ghost number $+1$ because it sends variables of ghost number -2 to variables of ghost numbers -1 . On the one hand, by definition of the first-stage reducibility functions and the action of δ on the ghost momenta, we have that $\delta^2(\mathcal{P}_{a_1}^{(1)}) = 0$. On the other hand, if $\delta(\lambda^{a_0} \mathcal{P}_{a_0}) = 0$, the exhaustion property of the first-stage reducibility functions (Equation (5.81)) together with Equation (5.85) imply that one can write $\lambda^{a_0} \mathcal{P}_{a_0} = \delta\left(f^{a_1} \mathcal{P}_{a_1}^{(1)} - \frac{1}{2} g^{a_0 b_0} \mathcal{P}_{a_0} \wedge \mathcal{P}_{b_0}\right)$ for some functions $f^{a_1}, g^{a_0 b_0}$ such that $g^{a_0 b_0} = -g^{b_0 a_0}$. This result means that if an element of T_{-1} is δ -closed, then it is δ -exact, i.e. it is the image through δ of an element of T_{-2} . The first step of the resolution for reducible constraints has been built.

We now assume that the theory is at least second-stage reducible, i.e. that $L \geq 2$. Using Equation (5.85), the map δ can be extended to $\mathcal{C}^\infty(T^*Q) \otimes (\wedge^3 K^{(0)} \oplus K^{(1)} \wedge K^{(0)})$ by derivation so that it lands in T_{-2} , but notice that if an element of T_{-2} is a cocycle – i.e. δ -closed – there is no reason whatsoever that it is a coboundary – i.e. δ -exact. For example, Equation (5.82) together with Corollary 4.25 imply that there exists a smooth function $C_{a_2}^{a_0 b_0}$ such that:

$$Z_{a_2}^{(2)a_1} Z_{a_1}^{(1)a_0} = C_{a_2}^{a_0 b_0} \varphi_{b_0} \quad (5.86)$$

Multiplying both sides with φ_{a_0} , Equation (5.80) implies that the left-hand-side strongly vanishes. Then so does the right-hand side, which means that we can suppose that $C_{a_2}^{a_0 b_0} = -C_{a_2}^{b_0 a_0}$ (the assumptions of Theorem 10.1 in [Henneaux and Teitelboim, 1992] are satisfied). Then by Equation (5.86) the element $Z_{a_2}^{(2)a_1} \mathcal{P}_{a_1}^{(1)} + \frac{1}{2} C_{a_2}^{a_0 b_0} \mathcal{P}_{a_0} \wedge \mathcal{P}_{b_0}$ of T_{-2} is δ -closed but not δ -exact. Indeed, if it were, the first term would necessarily be in the image $\delta(K^{(1)} \wedge K^{(0)}) \subset T_{-2}$ so the coefficient $Z_{a_2}^{(2)a_1}$ would involve a constraint, i.e. vanish on the constraint surface. This is possible only if the functions $Z_{a_1}^{(1)a_0}$ are independent, which is not the case as we assume that the theory is at least second-stage reducible.

Thus, we set $K^{(2)} = \mathbb{R}^{A_2}$ and let us denote by $\mathcal{P}_1^{(2)}, \dots, \mathcal{P}_{A_2}^{(2)}$ its standard basis. We call these variables *3-rd order ghost momenta* and we assign two different grading to them: a ghost number of -3 and a pure antighost number of 3 . Then, one extends the map δ to $\mathcal{C}^\infty(T^*Q) \otimes K^{(2)}$ by:

$$\delta(\mathcal{P}_{a_2}^{(2)}) = Z_{a_2}^{(2)a_1} \mathcal{P}_{a_1}^{(1)} + \frac{1}{2} C_{a_2}^{a_0 b_0} \mathcal{P}_{a_0} \wedge \mathcal{P}_{b_0} \quad (5.87)$$

and we set $T_{-3} = \mathcal{C}^\infty(T^*Q) \otimes (\wedge^3 K^{(0)} \oplus K^{(1)} \wedge K^{(0)} \oplus K^{(2)})$ so that one can straightforwardly extend the map δ to T_{-3} . Then Equation (5.87) is so that every element of T_{-2} that is δ -closed is δ -exact, meaning that it can be written as the image of an element in T_{-3} . More precisely, if $\lambda^{a_1} \mathcal{P}_{a_1}^{(1)} + \frac{1}{2} h^{a_0 b_0} \mathcal{P}_{a_0}^{(0)} \wedge \mathcal{P}_{b_0}^{(0)} \in T_{-2}$ is δ -closed, then Equation (5.85) implies that:

$$(\lambda^{a_1} Z_{a_1}^{(1)b_0} + h^{a_0 b_0} \varphi_{a_0}) \mathcal{P}_{a_0}^{(0)} = 0$$

Then Equation (5.83) implies that there exist two smooth functions $f^{a_2}, g^{a_1 a_0}$ such that:

$$0 = f^{a_2} Z_{a_2}^{(2)a_1} Z_{a_1}^{(1)b_0} + g^{a_1 a_0} Z_{a_1}^{(1)b_0} \varphi_{a_0} + h^{a_0 b_0} \varphi_{a_0} = (-f^{a_2} C_{a_2}^{a_0 b_0} + g^{a_1 a_0} Z_{a_1}^{(1)b_0} + h^{a_0 b_0}) \varphi_{a_0}$$

where we used Equation (5.86) and the antisymmetry property of the coefficient $C_{a_2}^{a_0 b_0}$ to obtain the last expression. With Equation (5.81), one can show³⁰ that there exists a smooth function $g^{a_0 b_0 c_0}$ fully antisymmetric on the three indices such that:

$$h^{a_0 b_0} = f^{a_2} C_{a_2}^{a_0 b_0} - 2g^{a_1 [a_0} Z_{a_1}^{(1)|b_0]} + g^{a_0 b_0 c_0} \varphi_{c_0} \quad (5.88)$$

where the bracket in the exponent of the second-term on the right-hand side should be read as: $g^{a_1 [a_0} Z_{a_1}^{(1)|b_0]} = \frac{1}{2} (g^{a_1 a_0} Z_{a_1}^{(1)b_0} - g^{a_1 b_0} Z_{a_1}^{(1)a_0})$. Then, one can straightforwardly check that Equations (5.87) and (5.88) imply that:

$$\lambda^{a_1} \mathcal{P}_{a_1}^{(1)} + \frac{1}{2} h^{a_0 b_0} \mathcal{P}_{a_0}^{(0)} \wedge \mathcal{P}_{b_0}^{(0)} = \delta \left(f^{a_2} \mathcal{P}_{a_2}^{(2)} + g^{a_1 a_0} \mathcal{P}_{a_1}^{(1)} \wedge \mathcal{P}_{a_0}^{(0)} + \frac{1}{6} g^{a_0 b_0 c_0} \mathcal{P}_{a_0}^{(0)} \wedge \mathcal{P}_{b_0}^{(0)} \wedge \mathcal{P}_{c_0}^{(0)} \right)$$

The term in the parenthesis on the right is indeed an element of T_{-3} , so this equation proves the exactness of δ at T_{-2} . Thus, the second step of the resolution (for at least second-stage reducible theories) has been built.

Following the same arguments one can construct a full resolution of $\mathcal{C}^\infty(\Sigma)$ from a given set of reducible constraints. At each level, we set $K^{(k)} = \mathbb{R}^{A_k}$ and we denote by $\mathcal{P}_1^{(k)}, \dots, \mathcal{P}_{A_k}^{(k)}$ its standard basis³¹; we call them *k-th order ghost momenta* and we attribute to them the ghost number $-(k+1)$ and pure antighost number $k+1$. One additionally sets:

$$T_{-(k+1)} = \mathcal{C}^\infty(T^*Q) \otimes \bigoplus_{i=0}^k \wedge^{i+1} (K^{(0)} \oplus \dots \oplus K^{(k)}) \Big|_{-(k+1)}$$

The previously introduced T_{-1}, T_{-2} and T_{-3} indeed correspond to this formula, and one has, for $k \geq 3$:

$$T_{-(k+1)} = \mathcal{C}^\infty(T^*Q) \otimes (K^{(k)} \oplus K^{(k-1)} \wedge K^{(0)} \oplus K^{(k-2)} \wedge K^{(1)} \oplus \dots \oplus K^{(k-2)} \wedge K^{(0)} \wedge K^{(0)} \oplus K^{(k-3)} \wedge K^{(1)} \wedge K^{(0)} \oplus \dots \oplus \wedge^{k+1} K^{(0)}) \quad (5.89)$$

where by convention if two terms appear twice in the sum we consider only them once, e.g. in T_{-4} , Equation (5.89) involves $K^{(1)} \wedge K^{(0)} \wedge K^{(0)}$ and $K^{(0)} \wedge K^{(1)} \wedge K^{(0)}$, so we consider that they only appear once, the choice of which does not matter. For a L -th stage reducible theory (L being finite), we consider by convention that $K^{(k)} = 0$ for every $k > L$ so that $T_{-(k+1)}$ is still defined for every $k > L$. Notice that the product \wedge is graded commutative so $\wedge^p K^{(1)} \simeq S^p \mathbb{R}^{A_1}$ as the ghost number of $\mathcal{P}_{a_1}^{(1)}$ is even, and more generally $\wedge^p K^{(2j+1)} \simeq S^p \mathbb{R}^{A_{2j+1}}$. Then the family of $T_{-(k+1)}$ is a priori infinite and non trivial as soon as we have a first-stage reducible theory –

³⁰It is still unclear how to prove it although Equation (5.88) is the correct answer to reach.

³¹Be aware that here, A_k symbolizes the number of k -th order ghost momenta and not the number of *functionally independent* k -th order ghost momenta as in Chapter 10 of [Henneaux and Teitelboim, 1992].

while for an irreducible theory $\wedge^p K^{(0)} \simeq \wedge^p \mathbb{R}^{A_1}$ so the chain complex is finite (it corresponds to that made of the K_{-i}).

Let us now explain how can we extend the map δ to $K^{(k)}$ by induction. Assume that the differential has been defined up to order $2 \leq k < L$ so that it defines an exact sequence on the chain complex $T_{-(k+1)} \rightarrow T_{-k} \rightarrow \dots \rightarrow T_{-1} \rightarrow \mathcal{C}^\infty(T^*Q)$. It means that, for every $1 \leq i \leq k$ we have:

$$\delta(\mathcal{P}_{a_i}^{(i)}) = Z_{a_i}^{(i)a_{i-1}} \mathcal{P}_{a_{i-1}}^{(i-1)} + M_{a_i}^{(i)} \quad (5.90)$$

for some functions $M_{a_i}^{(i)} \in T_{-i}$ which does not involve any i -th order ghost momenta $\mathcal{P}_{a_{i-1}}^{(i-1)}$. So for example $M_{a_1}^{(1)} = 0$ for every a_1 (Equation (5.85)) while $M_{a_2}^{(2)} = \frac{1}{2} C_{a_2}^{a_0 b_0} \mathcal{P}_{a_0} \wedge \mathcal{P}_{b_0}$ for every a_2 (Equation (5.87)). At level $2 \leq k < L$ the reducibility equation (5.84) implies that there exists a smooth functions $C_{a_{k+1}}^{a_{k-1} a_0}$ such that:

$$Z_{a_{k+1}}^{(k+1)a_k} Z_{a_k}^{(k)a_{k-1}} = C_{a_{k+1}}^{a_{k-1} a_0} \varphi_{a_0}$$

This equation generalizes Equation (5.86) at order 1.

By Equation (5.90), the element $\delta(Z_{a_{k+1}}^{(k+1)a_k} \mathcal{P}_{a_k}^{(k)} - C_{a_{k+1}}^{a_{k-1} a_0} \mathcal{P}_{a_0}^{(0)} \wedge \mathcal{P}_{a_{k-1}}^{(k-1)})$ is an element of T_{-k} and a δ -cocycle, so by Lemma 10.A.1 in [Henneaux and Teitelboim, 1992], there exists an element $\overline{M}_{a_{k+1}}^{(k+1)} \in T_{-(k+1)}$ which does not possess any term $\mathcal{P}_{a_k}^{(k)}$ such that:

$$\delta\left(Z_{a_{k+1}}^{(k+1)a_k} \mathcal{P}_{a_k}^{(k)} - C_{a_{k+1}}^{a_{k-1} a_0} \mathcal{P}_{a_0}^{(0)} \wedge \mathcal{P}_{a_{k-1}}^{(k-1)}\right) = -\delta(\overline{M}_{a_{k+1}}^{(k+1)}) \quad (5.91)$$

Then we extend the map δ at order $k+1$ by setting $M_{a_{k+1}}^{(k+1)} = \overline{M}_{a_{k+1}}^{(k+1)} - C_{a_{k+1}}^{a_{k-1} a_0} \mathcal{P}_{a_0}^{(0)} \wedge \mathcal{P}_{a_{k-1}}^{(k-1)}$ so that we have Equation (5.90) at order $k+1$:

$$\delta(\mathcal{P}_{a_{k+1}}^{(k+1)}) = Z_{a_{k+1}}^{(k+1)a_k} \mathcal{P}_{a_k}^{(k)} + M_{a_{k+1}}^{(k+1)} \quad (5.92)$$

This definition is set so that we have precisely, by Equation (5.91), $\delta^2(\mathcal{P}_{a_{k+1}}^{(k+1)}) = 0$. Moreover, one can show (see the discussion below Theorem 10.2 in [Henneaux and Teitelboim, 1992]) that this definition implies in turn that the chain map δ is exact at $T_{-(k+1)}$, meaning that if an element of $T_{-(k+1)}$ is δ -closed, then it is the image of an element of $T_{-(k+2)}$. The resolution has thus been extended at level $k+1 \leq L$ by induction. The induction process stops when we reach the L -th order of reducibility, i.e. when we define Equation (5.92) for $\mathcal{P}_{a_L}^{(L)}$, or never stops if the theory is infinitely reducible.

At the end of the induction, we are equipped with a chain complex of vector spaces which is infinite in length:

$$\dots \longrightarrow T_{-p} \xrightarrow{\delta} T_{-p+1} \xrightarrow{\delta} \dots \xrightarrow{\delta} T_{-1} \xrightarrow{\delta} \mathcal{C}^\infty(T^*Q) \longrightarrow 0$$

The reason for the infinite length comes from the fact that the space T_{-2j} of even ghost number $-2j$ possesses $\wedge^j K^{(1)} \simeq S^j \mathbb{R}^{A_1}$ which never vanishes. By construction, this chain complex is a resolution of $\mathcal{C}^\infty(T^*Q) / \mathcal{I}_\Sigma \simeq \mathcal{C}^\infty(\Sigma)$ (see Theorem 10.3 in [Henneaux and Teitelboim, 1992]). This resolution is called the *Koszul-Tate resolution*, as Tate generalized the Koszul complex to non-regular sequences of elements of commutative rings. The Koszul complex associated to irreducible constraints is a particular case of that resolution.

The idea behind the BRST formalism in the reducible case is the same as that in the irreducible case: we use the Koszul-Tate resolution in order to define a δ -exact differential D on a particular bi-graded vector space, so that it becomes a differential when restricted to Σ .

We will see that because of the reducibility identity (5.80), this differential is *not* the δ -exact differential d introduced in Section 5.2 for the irreducible case, although the former is built from the latter. The zero-th cohomology group of this differential D would then correspond to the algebra of gauge-invariant functions. The BRST differential would then be defined as $s = \delta + D + \text{'more'}$ and would correspond to a particular homological vector field on the algebra of function on the extended phase space, which is much bigger than the one of the irreducible case.

As the name *ghost of ghost momenta* indicates, they are the conjugate momenta of so-called *ghosts of ghosts*. The latter generalize the original ghosts η^i by carrying higher degrees: to each k -th order ghost momentum $\mathcal{P}_{a_k}^{(k)}$ of ghost degree $-(k+1)$ correspond a k -th order ghost of ghost $\eta^{(k)a_k}$ of ghost degree $k+1$, which the convention that the original ghosts are denoted $\eta^{(0)a_0}$ instead of η^i . The k -th order ghosts of ghosts are linear coordinates on a degree $-(k+1)$ vector space $E_{-(k+1)}$, spanned by A_k vectors $\mu_{a_k}^{(k)}$. We then define the extended configuration space of a L -th stage reducible theory as the following positively graded vector bundle over the configuration space Q , that is to say: $\Omega = Q \times \bigoplus_{0 \leq i \leq L} E_{-(i+1)}$. In particular, if $L = \infty$ then the direct sum on the right-hand side possesses an infinite number of terms. The extended phase space is then the cotangent bundle of the extended configuration space, thus extending Equation (5.51):

$$\mathfrak{P} = T^*\Omega$$

As in the irreducible case, the algebra functions on the phase space then involves the ghosts for ghosts and their conjugate momenta:

$$\mathcal{C}^\infty(\mathfrak{P}) = \mathcal{C}^\infty(T^*Q) \otimes \wedge^\bullet(\eta^{(0)a_0}, \eta^{(1)a_1}, \dots, \eta^{(L)a_L}) \otimes \wedge^\bullet(\mathcal{P}_{a_0}^{(0)}, \mathcal{P}_{a_1}^{(1)}, \dots, \mathcal{P}_{a_L}^{(L)})$$

If the theory is infinitely reducible, the number of ghosts of ghosts and their momenta is infinite. The differential δ can then be considered as a derivation on $\mathcal{C}^\infty(\mathfrak{P})$ squaring to zero or, equivalently, as a homological vector field on the graded manifold \mathfrak{P} , once we set:

$$\delta(\eta^{(k)a_k}) = 0 \quad \text{for every } k \tag{5.93}$$

In the following we will not talk about ghost momenta before we have defined the differential D and we will work exclusively on Σ so the \approx sign can be interpreted as an equality (on the constraint surface only). Recall that in the irreducible case, we had a differential d modulo δ on $\mathcal{C}^\infty(\Sigma) \otimes \wedge^\bullet(\eta^1, \dots, \eta^{A_0})$ – later denoted by $\mathcal{C}^\infty(\Sigma) \otimes \wedge^\bullet(\eta^{(0)a_0})$ – satisfying Equations (5.12) and (5.13):

$$df = X_k(f)\eta^k \quad \text{and} \quad d\eta^k = -\frac{1}{2}C_{ij}^k \eta^i \wedge \eta^j$$

In the reducible case, we still define C_{ij}^k as the structure functions characterizing the brackets of the Hamiltonian vector fields, but there always exists an ambiguity in their definition as the Hamiltonian vector fields need not be independent (see Remark 4.65). The irreducible case had the following advantage that one can use the Jacobi identity for the constraints to deduce that $d^2(\eta^k) \approx 0$ and $d^2(\mathcal{P}_i) \approx 0$. More precisely, if the constraints are irreducible, the right-hand side of Equation (4.83) does not possess a term of the form $\tau_{imn}^I Z_I^k$ which appears for reducible theories and has no reason to vanish on the constraint surface, and we are left with the term $\sigma_{imn}^{kl} \varphi_l$ which vanishes on Σ . A compact form of such a situation is found in Equation (5.16). Rather, when the theory is reducible, the exhaustion property (5.81) of the 1-st stage reducibility functions implies that (in the current notations):

$$X_{[i}(C_{mn]}^k) + C_{[mn}^j C_{i]j}^k \approx f_{imn}^{a_1} Z_{a_1}^{(1)k}$$

where $f_{imn}^{a_1} \in \mathcal{C}^\infty(T^*Q)$ is fully antisymmetric in i, m, n . So we a priori do not have $d^2(\eta^k) \approx 0$ (although we still have $d^2(f) \approx 0$ as Equation (5.18) is not impacted by the reducibility).

As is explained in subsection 10.4.2 in [Henneaux and Teitelboim, 1992], d turns out to be a differential on a particular subalgebra of $\mathcal{C}^\infty(\Sigma) \otimes \wedge^\bullet(\eta^{(0)a_0})$ consisting of elements of the form $\alpha_{i_1 \dots i_l} \eta^{i_1} \wedge \dots \wedge \eta^{i_l}$ such that:

$$Z_{a_1}^{(1)i_1} \alpha_{i_1 \dots i_l} \approx 0 \quad (5.94)$$

The contracted index has no importance because the lower indices of α are fully antisymmetric. The algebra formed of such elements is called the *longitudinal ghost algebra*, and will be further denoted \mathcal{L} . It is generated by a $\text{codim}(\Sigma)$ subspace of $\wedge^1(\eta^1, \dots, \eta^{A_0})$ which corresponds to a subset of irreducible constraints within the set of reducible ones. Thus it inherits the grading corresponding to the ghost degree of $\wedge^\bullet(\eta^1, \dots, \eta^{A_0})$:

$$\mathcal{L}|_k \subset \wedge^k(\eta^1, \dots, \eta^{A_0})$$

for every $1 \leq k \leq A_0$. If one adds $\mathcal{C}^\infty(\Sigma)$ to \mathcal{L} at ghost number 0 then the zero-th cohomology group of $d : \mathcal{L}|_\bullet \rightarrow \mathcal{L}|_{\bullet+1}$ corresponds to the gauge invariant functions. However we do not possess any explicit basis for this subalgebra, and we would rather like to work in $\mathcal{C}^\infty(\Sigma) \otimes \wedge^\bullet(\eta^{(0)a_0})$ because we would like to keep the redundant ghosts for the same reasons that were invoked for reducible constraints.

The strategy is then to mimick the BRST formalism to d and this subalgebra, by finding a resolution of \mathcal{L} in terms of a chain complex involving the ghosts of ghosts. However, contrary to the argument presented in Section 5.2, in the present case the differential defining the resolution of the chain complex on which d is a differential will go in the reverse direction. More precisely, let σ be the $\mathcal{C}^\infty(\Sigma)$ linear map defined on ghosts of ghosts as:

$$\sigma(\eta^{(k)a_k}) \approx Z_{a_{k+1}}^{(k+1)a_k} \eta^{(k+1)a_{k+1}} \quad (5.95)$$

We attribute to the k -th order ghost of ghost $\eta^{(k)a_k}$ a (pure) ghost number $k+1$, and we use this grading to extend σ to a graded derivation on the entire algebra $\mathcal{C}^\infty(\Sigma) \otimes \wedge^{\geq 1}(\eta^{(0)a_0}, \eta^{(1)a_1}, \dots, \eta^{(L)a_L})$. From Equation (5.95) we deduce that σ has ghost number $+1$.

Notice that the definition of σ induces another grading on this algebra, it indeed increases the order of the ghosts of ghosts. Following subsection 10.4.4 in [Henneaux and Teitelboim, 1992], we call this grading the *auxiliary grading* (in the afore mentioned reference, the map σ is called Δ). It is defined as follows:

$$\text{aux}(f) = 0 \quad \text{and} \quad \text{aux}(\eta^{(k)a_k}) = k$$

In particular, the map d has auxiliary grading zero. This grading allows to understand the algebra $\mathcal{C}^\infty(\Sigma) \otimes \wedge^\bullet(\eta^{(0)a_0}, \eta^{(1)a_1}, \dots, \eta^{(L)a_L})$ as a bigraded complex $V^{\bullet, \bullet}$ where the first slot is the polynomial degree (starting at 1) in the algebra while the second slot is the auxiliary grading:

$$V^{m, n} = \mathcal{C}^\infty(\Sigma) \otimes \wedge^m(\eta^{(0)a_0}, \eta^{(1)a_1}, \dots, \eta^{(L)a_L})|_{\text{aux}=n}$$

Graphically, the bi-graded vector space $V^{\bullet, \bullet}$ corresponds to the tensor product of the following diagram with the algebra of smooth functions $\mathcal{C}^\infty(\Sigma)$ (the left and right angles symbolize the vector space spanned by what is inside):

$$\begin{array}{ccccccc}
\cdots & & \cdots & & \cdots & & \\
\uparrow \sigma & & \uparrow \sigma & & \uparrow \sigma & & \\
\langle \eta^{(2)a_2} \rangle & \xrightarrow{d} & \left\langle \begin{array}{l} \eta^{(0)a_0} \wedge \eta^{(2)a_2} \\ \oplus \eta^{(1)a_1} \wedge \eta^{(1)a_1} \end{array} \right\rangle & \xrightarrow{d} & \left\langle \begin{array}{l} \eta^{(0)a_0} \wedge \eta^{(0)a_0} \wedge \eta^{(2)a_2} \\ \oplus \eta^{(0)a_0} \wedge \eta^{(1)a_1} \wedge \eta^{(1)a_1} \end{array} \right\rangle & \xrightarrow{d} & \cdots \\
\uparrow \sigma & & \uparrow \sigma & & \uparrow \sigma & & \\
\langle \eta^{(1)a_1} \rangle & \xrightarrow{d} & \langle \eta^{(0)a_0} \wedge \eta^{(1)a_1} \rangle & \xrightarrow{d} & \langle \eta^{(0)a_0} \wedge \eta^{(0)a_0} \wedge \eta^{(1)a_1} \rangle & \xrightarrow{d} & \cdots \\
\uparrow \sigma & & \uparrow \sigma & & \uparrow \sigma & & \\
\langle \eta^{(0)a_0} \rangle & \xrightarrow{d} & \langle \wedge^2 \eta^{(0)a_0} \rangle & \xrightarrow{d} & \langle \wedge^3 \eta^{(0)a_0} \rangle & \xrightarrow{d} & \cdots
\end{array}$$

Notice that the vertical arrows have the opposite direction compared to the map δ in the bi-graded vector space $M^{\bullet,\bullet}$ in Section 5.2. However we preserve the homological property of these arrows, since by Equation (5.84), the map σ is a differential on $V^{\bullet,\bullet}$ which increases the auxiliary degree by +1. The kernel of σ on $V^{\bullet,0} = \mathcal{C}^\infty(\Sigma) \otimes \wedge^\bullet(\eta^{(0)a_0})$ is precisely the algebra of longitudinal ghosts \mathcal{L} since, for any $\alpha \approx \alpha_{i_1 \dots i_l} \eta^{i_1} \wedge \dots \wedge \eta^{i_l}$ the identity $\sigma(\alpha) = 0$ can be rewritten as Equation (5.94). We actually have a stronger result, that is: the differential $\sigma : V^{\bullet,\bullet} \rightarrow V^{\bullet,\bullet+1}$ of auxiliary degree +1 defines a resolution of \mathcal{L} (Theorem 10.4 in [Henneaux and Teitelboim, 1992]), i.e:

$$H^0(\sigma) \simeq \mathcal{L} \quad \text{and} \quad H^k(\sigma) = 0 \text{ for every } k \leq 1$$


One observes that d turns out to be a differential modulo σ (see Remark 5.7). It means that d and σ anti-commute – i.e. $d \circ \sigma + \sigma \circ d = 0$ – and that there exists a graded derivation $D^{(-1)}$ of $V^{\bullet,\bullet}$ of ghost number 0 and of auxiliary grading -1 such that:

$$d^2 = -[\sigma, D^{(-1)}]$$

Pursuing the analogy with the homological perturbation theory presented in Section 5.2, one should be able to extend d to the ghosts of ghosts, and define a differential D on the bigraded vector space $V^{\bullet,\bullet}$ which is such that $D = \sigma + d + \text{'more'}$. Notice that the situation is however quite different than the one addressed in Section 5.2, not only because the direction of the vertical arrows differs, but also because we add $A_1 + A_2 + \dots$ ghosts of ghosts to the bottom line, and their number is possibly quite different than A_0 . The subsections 10.4.5 and 10.4.6 in [Henneaux and Teitelboim, 1992] explain that, although the afore-mentioned difference with the BRST formalism in Section 5.2, the differential D exists and can be decomposed with respect to the auxiliary grading:

$$D = \sigma + d + \sum_{k \leq -1} D^{(k)}$$

where $\text{aux}(D^{(k)}) = k < 0$. We can graphically represent the action of the differential D on the bigraded vector space $M^{\bullet,\bullet}$, as it acts ‘diagonally’, but in another direction compared to the BRST differential:



From the diagram, we observe that the top left/bottom right diagonals have constant total degree – understood as the sum of the polynomial degree m with the auxiliary grading n – which actually turns out to coincide with the (pure) ghost number. Then we can compactly coin these diagonals under a unique grading $V^k = \bigoplus_{1 \leq m, 0 \leq n, m+n=k} V^{m,n}$ involving the ghost number k . They are shifted to one another via the differential $D : V^k \rightarrow V^{k+1}$, which has thus ghost number $+1$, just like d . We additionally complete the family of graded spaces V^k by a vector space at degree zero $V^0 = \mathcal{C}^\infty(\Sigma)$, on which D acts like d , so that we have the following chain complex:

$$0 \xrightarrow{D} \mathcal{C}^\infty(\Sigma) \xrightarrow{D} V^1 \xrightarrow{D} V^2 \xrightarrow{D} V^3 \xrightarrow{D} \dots$$

The k -th group of cohomology of D is then defined as usual:

$$H^k(D) = \frac{\text{Ker}(D : V^k \rightarrow V^{k+1})}{\text{Im}(D : V^{k-1} \rightarrow V^k)}$$

And then, the differential D has a property similar to that of Theorem 5.12 (this is Theorem 10.5 in [Henneaux and Teitelboim, 1992]):

Theorem 5.54. *The cohomology of D is equal to the cohomology of d modulo σ , that is to say:*

$$H^k(D) = H^k(d)$$

for every $k \geq 0$ (understanding that for $k \leq -1$, $H^k(d) = 0$). In particular, the zero-th group of cohomology of D coincides with that of d on $\mathcal{C}^\infty(\Sigma) \oplus \mathcal{L}$ and we have:

$$H^0(D) \simeq \mathcal{C}^\infty(\Sigma_{ph})$$

Pushing further the analogy with the irreducible case, we are now in possession with a non-negatively graded chain complex (V^\bullet, D) playing the same role as (W^\bullet, d) in Section 5.2. The

former complex V^\bullet is the straightforward generalization of the latter to the reducible context. We can then proceed to apply the BRST formalism presented in Section 5.2, where we replace (W^k, d) by (V^\bullet, D) . By Equation (5.18), we know that $d^2 f \approx 0$ so the only component of D that acts on the functions on Σ is $D^{(0)} = d$ as it does not need any contribution from other components $D^{(k)}$ to satisfy the nilpotency condition. Then, the differential D extends to a derivation of ghost number +1 on the following tensor product:

$$S^\bullet = \mathcal{C}^\infty(T^*Q) \otimes_{\mathcal{C}^\infty(\Sigma)} V^\bullet$$

Indeed, only $D^{(0)} = d$ acts on $\mathcal{C}^\infty(T^*Q)$ via Equation (5.12), while the other components of D act on the ghosts and ghosts of ghosts, vanishing on the smooth functions. Be aware however that now $D^2 \neq 0$, as was the case for d in the Section 5.2. The BRST formalism for the irreducible case will however solve the problem, by showing that D is a differential modulo δ , and that there exists a total differential $s = \delta + D + \text{'more'}$ such that $H^0(s) = H^0(D)$.

More precisely, there are as many ghost of ghost momenta as there are ghosts of ghosts, so we are in a symmetric situation as in the irreducible case. To pursue the analogy, let $M^{\bullet,\bullet}$ be the following bi-graded vector space:

$$M^{m,n} = S^m \otimes_{\mathcal{C}^\infty(T^*Q)} T_{-n}$$

It should not to be confused with the one in Section 5.2, although they play the same role, as we have for the pairs of lowest integers:

$$\begin{aligned} M^{0,0} &= \mathcal{C}^\infty(T^*Q), & M^{1,0} &= \mathcal{C}^\infty(T^*Q) \otimes \wedge^1(\eta^{(0)a_0}), & M^{0,1} &= \mathcal{C}^\infty(T^*Q) \otimes \wedge^1(\mathcal{P}_{a_0}^{(0)}), \\ M^{2,0} &= \mathcal{C}^\infty(T^*Q) \otimes (\wedge^2(\eta^{(0)a_0}) \oplus \wedge^1(\eta^{(1)a_1})), & M^{0,2} &= \mathcal{C}^\infty(T^*Q) \otimes (\wedge^2(\mathcal{P}_{a_0}^{(0)}) \oplus \wedge^1(\mathcal{P}_{a_1}^{(1)})) \\ M^{1,1} &= \mathcal{C}^\infty(T^*Q) \otimes \wedge^1(\eta^{(0)a_0}) \otimes \wedge^1(\mathcal{P}_{a_0}^{(0)}) \end{aligned}$$

In analogy with the irreducible case, we extend δ to $M^{\bullet,\bullet}$ using Equation (5.93). On the other hand, it is much more difficult to find how to extend D to the ghost of ghost momenta. In the irreducible case, we have found an explicit expression (5.33), so that Proposition 5.6 held. The present situation, being much more complex as D contains more components than the mere longitudinal differential d , will not be solved explicitly.

$$0 \xrightarrow{s} M^{-p} \xrightarrow{s} \dots \xrightarrow{s} M^{-1} \xrightarrow{s} M^0 \xrightarrow{s} M^1 \xrightarrow{s} \dots \xrightarrow{s} M^p \xrightarrow{s} 0$$

As in Section 5.2, the bottom left/top right diagonals in the bi-graded vector space $M^{\bullet,\bullet}$ have constant ghost number and can then be encapsulated into a compact form $M^k = \bigoplus_{0 \leq m,n, m-n=k} M^{m,n}$. As we have seen that T_\bullet is not bounded below, and S_\bullet is not bounded above (for the same reason), the graded space M^k is neither bounded above nor below, contrary to the irreducible case. Then, Section 10.5 in [Henneaux and Teitelboim, 1992] proves the existence and unicity (up to canonical transformations in the extended phase space) of a BRST differential:

Theorem 5.55. *There exists a differential $s = \delta + D + \text{'more'}$ of ghost number +1, making M^\bullet a chain complex:*

$$\dots \xrightarrow{s} M^{-2} \xrightarrow{s} M^{-1} \xrightarrow{s} M^0 \xrightarrow{s} M^1 \xrightarrow{s} M^2 \xrightarrow{s} \dots$$

and such that the cohomology of s is equal to the cohomology of D modulo δ , that is to say:

$$H^k(s) = H^k(D)$$

for every $k \geq 0$ (understanding that for $k \leq -1$, $H^k(D) = 0$). In particular, the zero-th group of cohomology of s coincides with that of D and we have, by Theorem 5.54:

$$H^0(s) \simeq \mathcal{C}^\infty(\Sigma_{ph})$$

This theorem answers the question of the existence of a BRST differential in the reducible case. The proof of this result in [Henneaux and Teitelboim, 1992] is based on showing the existence and uniqueness of the BRST charge Ω by homological perturbation. In order to make sense of it, let us notice that the content of Section 5.3 also applies to the reducible case: we have introduced the extended configuration space $\mathfrak{Q} = Q \times \bigoplus_{0 \leq i \leq L} E_{-i+1}$ and the extended phase space is $\mathfrak{P} = T^*\mathfrak{Q}$, so that $\mathcal{C}^\infty(\mathfrak{P}) = M^{\bullet, \bullet}$. The canonical graded symplectic form on the extended phase space is then a (possibly infinite) sum of terms generalizing the canonical graded symplectic form (5.62) met in the irreducible case:

$$\omega = \sum_{i=1}^n dp_i \wedge dq^i + \sum_{k=0}^L \sum_{a_k=1}^{A_k} d\mathcal{P}_{a_k}^{(k)} \wedge d\eta^{(k)a_k}$$

The corresponding graded Poisson bracket generalizes that of Proposition 5.70 so that, when evaluated on the k -th order ghosts of ghosts and their associated momenta, we obtain a generalization of Equation (5.25) that takes into account the order k :

$$\{\eta^{(k)a_k}, \mathcal{P}_{b_k}^{(k)}\} = -(-1)^{k+1} \{\mathcal{P}_{b_k}^{(k)}, \eta^{(k)a_k}\} = \delta_{b_k}^{a_k}$$

Any other combination vanishes. Thus, the generalization of Equation (5.70) to the reducible case is:

$$\{F, G\} = \sum_{i=1}^n \frac{\partial F}{\partial q^i} \frac{\partial G}{\partial p_i} - \frac{\partial F}{\partial p_i} \frac{\partial G}{\partial q^i} + \sum_{k=0}^L (-1)^{(k+1)(1+\text{gh}(F))} \sum_{a_k=1}^{A_k} \frac{\partial F}{\partial \eta^{(k)a_k}} \frac{\partial G}{\partial \mathcal{P}_{a_k}^{(k)}} - (-1)^{k+1} \frac{\partial F}{\partial \mathcal{P}_{a_k}^{(k)}} \frac{\partial G}{\partial \eta^{(k)a_k}}$$

while the generalization of Equation (5.68) is:

$$\{F, G\} = \sum_{i=1}^n \frac{\partial F}{\partial q^i} \frac{\partial G}{\partial p_i} - \frac{\partial F}{\partial p_i} \frac{\partial G}{\partial q^i} + \sum_{k=0}^L \sum_{a_k=1}^{A_k} F \frac{\overleftarrow{\partial}}{\partial \eta^{(k)a_k}} \frac{\overrightarrow{\partial}}{\partial \mathcal{P}_{a_k}^{(k)}} G - (-1)^{k+1} F \frac{\overleftarrow{\partial}}{\partial \mathcal{P}_{a_k}^{(k)}} \frac{\overrightarrow{\partial}}{\partial \eta^{(k)a_k}} G$$

The sign $-(-1)^{k+1}$ corresponds to the product of two signs: the first is the usual minus sign appearing in the usual Poisson bracket on smooth manifolds (as in the first sum), while the sign $(-1)^{k+1}$ descends from a generalization of Equation (5.23) to the reducible setting and takes into account the ghost number of the ghosts of ghosts and their conjugate momenta:

$$\frac{\partial}{\partial \eta^{(k)a_k}} \wedge \frac{\partial}{\partial \mathcal{P}_{b_k}^{(k)}} = (-1)^{k+1} \frac{\partial}{\partial \mathcal{P}_{b_k}^{(k)}} \wedge \frac{\partial}{\partial \eta^{(k)a_k}}$$

Then, the BRST differential s is a homological vector field on the extended phase space \mathfrak{P} because $s^2 = 0$, while the proof of Theorem 5.55 in [Henneaux and Teitelboim, 1992] shows that it is actually a Hamiltonian vector field with respect to the canonical graded Poisson bracket, so that Proposition 5.48 and Equation (5.63) hold as well in the reducible case. As in the irreducible case, the BRST differential s and the BRST charge Ω can be decomposed by components labelled by pure antighost numbers. However, while the number of components was finite in the irreducible case (see Equations (5.46) and (5.75)), here the sums are unbounded above:

$$s = \sum_{-1 \leq i} s_{(i)} \quad \text{and} \quad \Omega = \sum_{0 \leq i} \Omega_{(i)}$$

They are such that $s_{(-1)} = \delta$ and $s_{(0)} = D$, from which one can deduce $\Omega_{(0)}$ and $\Omega_{(1)}$ as in Equation (5.74).

5.5 BRST quantization of Hamiltonian systems under constraints

To summarize what we have shown so far: given a system of first-class constraints, we construct an extended phase space \mathfrak{P} by adding ghosts and ghost momenta. There exists a differential on $\mathcal{C}^\infty(\mathfrak{P})$ of ghost number $+1$ – equivalently, a degree $+1$ vector field on \mathfrak{P} – whose degree 0 cohomology coincides with the classical observables (Theorem 5.12 in the irreducible case and Theorem 5.55 in the reducible one), equivalently described by $\mathcal{C}^\infty(\Sigma_{ph})$ or by the gauge invariant functions on Σ . In particular, what we know now is that there is a one-to-one correspondence between classical observables on Σ and s -closed – we say *BRST-closed* – degree 0 functions on \mathfrak{P} , up to s -exact – we say *BRST-exact* – functions:

$$\begin{array}{ccc} \text{gauge-invariant} & & \text{BRST-closed } \mathcal{O} \in \mathcal{C}^\infty(\mathfrak{P}) \text{ of degree 0,} \\ O \in \mathcal{C}^\infty(\Sigma) & \xleftarrow{1-1} \tilde{O} \in \mathcal{C}^\infty(\Sigma_{ph}) & \xleftarrow{1-1} \text{up to BRST-exact terms} \end{array}$$

We will call *BRST observables* such BRST-closed degree 0 elements of $\mathcal{C}^\infty(\mathfrak{P})$, with the understanding that two such functions should be identified if their difference is a BRST-exact term.

Recall that by abuse of denomination (Definition 4.82), we called *gauge-invariant function* any smooth function $f \in \mathcal{C}^\infty(T^*Q)$ whose bracket with the constraints – still and always first-class in our context – is weakly vanishing:

$$\{\varphi_{a_0}, f\} \approx 0 \quad (5.96)$$

A classical observable $\tilde{O} \in \mathcal{C}^\infty(\Sigma_{ph})$ – equivalently a smooth map $O \in \mathcal{C}^\infty(\Sigma)$ which is gauge invariant i.e. constant along the gauge orbits – gives rise to an infinite number of such gauge-invariant globally defined functions, which differ from another only outside Σ . In Definition 4.82 these functions are called *gauge-invariant extensions of O* . From now on we will mostly consider such gauge-invariant functions and extensions, as they are defined over the whole phase space, which is easier to manipulate than Σ . A gauge-invariant extension of the classical observable O will also be denoted O because, although we know that it is not defined only on Σ but the entire phase space T^*Q , it is often obvious and harmless to conflate the notations.

As for the BRST charge, any BRST observable can be decomposed with respect to the pure antighost number: $\mathcal{O} = \sum_{p \geq 0} \mathcal{O}_{(p)}$, where $\mathcal{O}_{(p)}$ has pure antighost number p and the sum being finite but possibly very long. Then the relationship between a classical observable O and a BRST observable \mathcal{O} is that the pure antighost number 0 component $\mathcal{O}_{(0)}$ of the BRST observable \mathcal{O} should be a gauge-invariant extension of the classical observable O :

$$\mathcal{O} = O + \sum_{p \geq 1} \mathcal{O}_{(p)}$$

where here O should be understood as a smooth function on T^*Q satisfying (5.96) extending the true classical observable \tilde{O} outside Σ . We say that \mathcal{O} is a *BRST-invariant extension of O* , in reference to the gauge-invariant extensions of classical observables as introduced in Definition 4.82. More generally, we set:

Definition 5.56. For any $f \in \mathcal{C}^\infty(T^*Q)$ we call BRST-invariant extension of f any BRST-closed function $F \in \mathcal{C}^\infty(\mathfrak{P})$ such that:

$$F = f + \sum_{p \geq 1} F_{(p)} \quad (5.97)$$

where $F_{(p)}$ is the component of F of pure antighost number p .

To mimick Hamiltonian formalism on classical observables (or more generally on gauge-invariant extensions), we need a Poisson bracket (in fact a symplectic structure) on our algebra of BRST observables. Hopefully, there is a Poisson structure on $H^0(s)$ induced from the canonical symplectic structure on $\mathcal{C}^\infty(\mathfrak{P})$ described in Example 5.35, and it has the following nice property:

Proposition 5.57. *The Poisson (in fact symplectic) structure induced on $H^0(s)$ is Poisson equivalent to the Poisson (in fact symplectic) structure obtained on $\mathcal{C}^\infty(\Sigma_{ph})$ by Poisson reduction.*

Proof. It is the pair of Theorems 11.1 and 11.2 in [Henneaux and Teitelboim, 1992]. A more mathematical formulations in [Figueroa-O’Farrill and Kimura, 1991b] and [Kimura, 1993], while [Stasheff, 1997] is a generalization of the latter. \square

Moreover, by Proposition 5.48 we know that $s = \{\Omega, \cdot\}$ so BRST-closedness of a function \mathcal{O} is equivalent to the identity $\{\Omega, \mathcal{O}\} = 0$. A BRST-exact function F is written as $F = \{\Omega, G\}$ and is then identified with the null observable when it has ghost number 0. Recall that the action of the BRST differential on the ghost momenta is of the form:

$$s(\mathcal{P}_{a_0}^{(0)}) = \varphi_{a_0} + \dots$$

where the dots symbolize terms of higher pure antighost number. Then, we deduce that this sum represents a BRST-invariant extension of the constraints φ_{a_0} . In other words, the BRST-exact function of $\mathcal{C}^\infty(\mathfrak{P})$ defined as:

$$\Phi_{a_0} = \{\Omega, \mathcal{P}_{a_0}^{(0)}\} \tag{5.98}$$

is a BRST-invariant extension of φ_{a_0} . As expected, we see that this extension, being exact and of ghost number 0, corresponds to the null observable, which corresponds to the fact that φ_{a_0} is a constraint.

Now let $f \in \mathcal{C}^\infty(T^*Q)$ and assume that it admits a BRST-invariant extension F as given in Equation (5.97). The term $F_{(1)}$ is of the form $F_{(1)} = F_{a_0}^{b_0} \eta^{(0)a_0} \wedge \mathcal{P}_{b_0}^{(0)}$, so that $s(f + F_{(1)}) = \eta^{(0)a_0} X_{a_0}(f) - F_{a_0}^{b_0} \eta^{(0)a_0} \varphi_{b_0} + \dots$. From this we deduce that on the constraint surface, $X_{a_0}(f) \approx 0$, i.e. that f is a gauge-invariant function in the sense of Definition 4.82. The converse is also true: if f is a gauge-invariant function in the above sense, by Theorem 5.12, $H^0(s)$ is isomorphic to the gauge-invariant functions on Σ , so there should exist a BRST-closed function F such that Equation (5.97) holds. We can even give an idea of the first term of this BRST-invariant extension. Since $X_{a_0}(f) \approx 0$, by Proposition 4.24 (extended globally) there exists smooth functions $F_{a_0}^{b_0}$ such that $X_{a_0}(f) = F_{a_0}^{b_0} \varphi_{b_0}$. Then, given the discussion leading to Equation (5.98), we have $F = f + F_{a_0}^{b_0} \eta^{(0)a_0} \wedge \mathcal{P}_{b_0}^{(0)} + \dots$. We thus have the following result:

Proposition 5.58. *Given a degree 0 BRST-closed function, the component of pure antighost number 0 is a gauge-invariant function. Conversely, any gauge-invariant function admits a BRST-invariant extension.*

Both the first-class Hamiltonian H' (see Definition 4.54), the total Hamiltonian H_T (see Definition 4.26) and the extended Hamiltonian H_E (see Definition 4.60) are first-class functions. Hence, they satisfy Equation (4.92) so they are gauge-invariant functions. By Proposition 5.58, there exists a BRST-invariant extension $\mathcal{H} \in \mathcal{C}^\infty(\mathfrak{P})$ of the first-class Hamiltonian H'^{32} . We call this function the *BRST-invariant Hamiltonian*. One could use this Hamiltonian to define

³²Since we only work with first-class constraints, the first-class Hamiltonian does not possess any constraints.

the time evolution of functions on the extended phase space. In particular, on the ghosts and ghost momenta, we have:

$$\dot{\eta}^{(0)a_0} = \{\eta^{(0)a_0}, \mathcal{H}\} = V_{b_0}^{(0)a_0} \eta^{(0)b_0} + \dots \quad (5.99)$$

$$\dot{\mathcal{P}}_{a_0}^{(0)} = \{\mathcal{P}_{a_0}^{(0)}, \mathcal{H}\} = -V_{a_0}^{(0)b_0} \mathcal{P}_{b_0}^{(0)} + \dots \quad (5.100)$$

where the dots stand for terms of higher pure antighost number (see Theorem 2 of Section 11.5 in [Rothe and Rothe, 2010] to justify the form of the first term). Although the ghost number is always preserved, the pure antighost number is not preserved by time evolution unless the dots identically vanish. From these equations, and from similar equations for higher order ghosts of ghosts and their conjugate momenta (possibly involving polynomial of lower order such fields on the right-hand side), one can unambiguously define the time evolution of any BRST-invariant extension F of a gauge-invariant function f :

$$\dot{F} = \{F, \mathcal{H}\} \quad (5.101)$$

Notice that The function \dot{F} is BRST-invariant because both F and \mathcal{H} are. Notice that as usual two functions \dot{F} and \dot{G} are identified if they differ by a BRST-exact term. Since $H^0(s)$ is Poisson isomorphic to the algebra of classical observables i.e. gauge-invariant functions on Σ , we deduce that Equation (5.101) is equivalent to the following one:

$$\dot{f} \approx \{f, H'\}$$

which is, by definition of classical observable, equivalent to $\{f, H_T\}$ and $\{f, H_E\}$.

Since the BRST-invariant extensions are BRST-closed, we deduce that the dynamics of such an extension F does not change if one picks up another representant of the BRST-invariant Hamiltonian: $\mathcal{H} \mapsto \mathcal{H} + \{K, \Omega\}$, where $K \in \mathcal{C}^\infty(\mathfrak{P})$ has ghost number -1 . Indeed, one has:

$$\{F, \mathcal{H} + \{K, \Omega\}\} = \{F, \mathcal{H}\} + \{\{F, K\}, \Omega\} + \underbrace{\{K, \{F, \Omega\}\}}_{=0} = \dot{F} + \text{a BRST-exact term} \quad (5.102)$$

Then we see that $\dot{F} = \{F, \mathcal{H}\}$ is cohomologically identified with $\{F, \mathcal{H} + \{K, \Omega\}\}$, proving that changing the Hamiltonian by a BRST-exact term (which is cohomologically trivial) does not change the dynamics of a BRST-invariant extension. This can be mathematically explained as follows: the original choice of BRST-Hamiltonian was totally arbitrary (in the limit imposed by the BRST-closedness condition) for components of pure antighost number higher than 1, so this arbitrariness should be reflected in the invariance of the dynamics of classical observables under a change of such components of pure antighost degrees higher than 1 by the addition of a BRST-exact term $\mathcal{H} + \{K, \Omega\}$. Notice that the transformation $\mathcal{H} \mapsto \mathcal{H} + \{K, \Omega\}$ also modifies the components of pure antighost degree 0, but at this level any such choice of function K is physically equivalent to fixing the values of Lagrange multipliers for the first-class constraints and add them to the first-class Hamiltonian:

$$\mathcal{H} \mapsto \mathcal{H} + \{K, \Omega\} \quad \xleftarrow{\text{physically equivalent to}} \quad H' \mapsto H' + \lambda^i \varphi_i$$

Then choosing a function K is physically equivalent in the canonical Hamiltonian formalism to passing from the first-class Hamiltonian H' to the extended Hamiltonian H_E with fixed values of the Lagrange multipliers, and we know that the classical observables "do not see" the difference between the first-class, total or extended Hamiltonian because they are insensitive to the presence of first-class constraints in these Hamiltonians. Thus, by Equation (5.102) and the innocuity of K on the dynamics of BRST-invariant extensions, we recover the classical

result that the presence of first-class constraints in the extended Hamiltonian does not have any incidence on the dynamics of classical observables. Choosing another K is then equivalent to changing the value of the Lagrange multipliers associated to the first-class constraints in the extended Hamiltonian H_E . Choosing a representant of the BRST-invariant Hamiltonian via a choice of a function K can then be understood as fixing a BRST-invariant extension of the extended Hamiltonian H_E for which the values of the Lagrange parameters have been chosen (i.e. fixing a gauge). As $\mathcal{H} + \{K, \Omega\}$ represents a BRST-invariant extension of such a gauge-fixed Hamiltonian, the function K is called the *gauge-fixing fermion*, and $\mathcal{H} + \{K, \Omega\}$ is called the *unitarizing Hamiltonian* in subsection 11.5.1 of [Rothe and Rothe, 2010], or by abuse of denomination the *gauge-fixed Hamiltonian* in Section 11.2 of [Henneaux and Teitelboim, 1992]. As a final remark on this topic, notice that the dynamics of the ghosts and ghost momenta – as determined in Equations (5.99) and (5.100) – are in general modified under a different choice of K , as they are not BRST-closed.

As is usual in Hamiltonian mechanics one can define an action principle from a given Hamiltonian. In our case, the extended Hamiltonian induces an *extended action*:

$$S_E[q, p, \lambda] = \int_{\mathbb{R}} p_i \dot{q}^i - H_E(q, p, \lambda) dt$$

where λ symbolizes the Lagrange multipliers associated to the first-class constraints. We know that this action generates much more solutions than the *total action* corresponding to the total Hamiltonian:

$$S_T[q, p, \lambda] = \int_{\mathbb{R}} p_i \dot{q}^i - H_T(q, p, \lambda) dt$$

This is because we have arbitrarily added all secondary first-class constraints to the total Hamiltonian to form the extended Hamiltonian. The solutions of the equations of motion induced by the action S_T correspond precisely to the solutions of the Euler-Lagrange equations (see the discussion leading to the system of equations (4.39)). See Section 3.3 in [Henneaux and Teitelboim, 1992] for a discussion on these differences. Now, since the gauge-fixed Hamiltonian corresponds to the extended Hamiltonian with fixed values for the Lagrange multipliers, we can define a *gauge-fixed BRST action* as:

$$S_{BRST}[q, p, \eta, \mathcal{P}] = \int_{\mathbb{R}} p_i \dot{q}^i + \sum_{k \geq 0} \mathcal{P}_{a_k}^{(k)} \dot{\eta}^{(k)a_k} - \mathcal{H} - \{K, \Omega\} dt \quad (5.103)$$

This action governs the dynamics in the extended phase space \mathfrak{P} and is BRST-invariant up to a boundary term, see Section 11.2 in [Henneaux and Teitelboim, 1992] for more details on this topic.

Having defined an action principle in the BRST formalism, one can now make sense of the Faddeev-Popov action in terms of the BRST-formalism. It turns out that the Faddeev-Popov action is obtained from the present formalism by promoting the Lagrange multipliers λ^k of the extended Hamiltonian to dynamical variables. This then requires the addition of their conjugate momenta b_k – which are supposed to identically vanish to single out the original phase space T^*Q as a constraint surface – and a pair of extra conjugate variables (ρ^k, \bar{c}_k) which do not change the cohomology of the BRST differential, because we impose:

$$s(\bar{c}_k) = b_k \quad \text{and} \quad s(b_k) = 0$$

We call the \bar{c}_k the *antighosts*, they have pure antighost number -1 (negative) and ghost number -1 , while the ρ^k are their associated *antighost momenta* and have ghost number $+1$. These variables $\lambda^k, b_k, \rho^k, \bar{c}_k$ extend the phase space \mathfrak{P} further with a corresponding Poisson (in fact symplectic) structure $\{\rho^k, \bar{c}_l\} = \delta_l^k$, and are said to be the *non-minimal sector* of this newly

defined total phase space, while the original ghosts of ghosts $\eta^{(k)a_k}$ and ghost of ghost momenta $\mathcal{P}_{a_k}^{(k)}$ form the *minimal sector* (although the distinction can be shown to be inessential). Then the Faddeev-Popov action can be understood as a gauged-fixed action in this total phase space. We refer to subsections 11.3.1–11.3.3 in [Henneaux and Teitelboim, 1992] and Section 6.2 in [Gomis et al., 1995] for a thorough explanation of this phenomena.

Remark 5.59. One can come back to the Lagrangian form of the action (5.103) by replacing the conjugate momenta by their on-shell values obtained from the equations of motion. This Lagrangian would then be BRST-invariant when all the momenta in the BRST charge are replaced by their on-shell values. These BRST-transformations are called *Lagrangian BRST transformations* and, as the Noether theorem states that to any invariance of the Lagrangian corresponds a conserved charge, it can be proven that the Noether charge corresponding to the Lagrangian BRST transformations is precisely the BRST charge Ω . See subsection 11.3.4 in [Henneaux and Teitelboim, 1992] and [Nirov and Razumov, 1993] for more details.

Now that we have established the correspondence between the Hamiltonian treatment of physical observable and the BRST formalism, we can proceed to quantize the latter. In Section 5.1 we have shown that Dirac canonical quantization relied on the idea that one can associate a quantum operator to a classical observable. Then, the Hilbert space \mathcal{H} of quantum states would be found "by hand" as an irreducible representation of the Lie algebra of such quantized operators. The Lie bracket would be given by the commutator of the operator, itself defined from the Poisson bracket defined on the classical observables (see e.g. Equation (5.1) in the unconstrained case, or (5.4) for pure second-class systems, allowing the reduction from a mixed system to a pure first-class system). Quantizing the BRST-formalism would first require to allow the quantization of (ghosts of) ghosts and their conjugate momenta: $\eta^{(k)a_k} \mapsto \hat{\eta}^{(k)a_k}$ and $\mathcal{P}_{a_k}^{(k)} \mapsto \hat{\mathcal{P}}_{a_k}^{(k)}$; this would allow to quantize any function on \mathfrak{P} (up to ordering problems). Second, since the extended phase space is a graded symplectic manifold (without any constraints at play), the Hilbert space \mathcal{H} of quantum states would be a representation of the commutator of operator:

$$[\hat{F}, \hat{G}] = i\hbar\{F, G\} \quad (5.104)$$

where on the right-hand side we have the graded Poisson bracket defined on the extended phase space \mathfrak{P} (see Example 5.35). Notice that the operators defined on this Hilbert space carry a degree – as they were carrying a degree (the ghost number) as function on the graded symplectic manifold \mathfrak{P} – and the commutator on the left-hand side of Equation (5.104) is a *graded commutator*, that is to say:

$$[\hat{F}, \hat{G}] = \hat{F} \circ \hat{G} - (-1)^{\text{gh}(F)\text{gh}(G)} \hat{G} \circ \hat{F} \quad (5.105)$$

This rule is consistent with the graded antisymmetry of the Poisson bracket on the right-hand side of Equation (5.104).

Under all these conventions, the BRST charge Ω becomes a linear operator on the Hilbert space \mathcal{H} of quantum states, that we will still denote Ω for convenience. When $F = G = \Omega$, the classical master equation (5.64) implies that the right-hand side of Equation (5.104) vanishes. Using Equation (5.105), the left-hand side then becomes:

$$[\Omega, \Omega] = 2\Omega^2 = 0 \quad (5.106)$$

The linear operator Ω is then a nilpotent operator. As the BRST charge was a real function in the classical theory, we require that the corresponding linear operator Ω is self-adjoint:

$$\Omega^* = \Omega \quad (5.107)$$

Henneaux and Teitelboim state in subsection 14.1.1 of [Henneaux and Teitelboim, 1992] that because of Equations (5.106) and (5.107), the Hilbert space \mathcal{H} does possess negative norm states (see also Section 10.6 in [Gomis et al., 1995] for a discussion about this norm and unitarity). As a physically relevant quantity in the classical BRST formalism was BRST-invariant extensions, which are BRST-closed functions of $\mathcal{C}^\infty(\mathfrak{P})$, we define the *BRST observables* on the Hilbert space as the linear operators A which graded commute with the BRST linear operator:

$$[A, \Omega] = 0 \quad (5.108)$$

Here again, if the ghost number of A is odd, then the commutator is an anticommutator because the ghost number of Ω is 1. Quantum states in the Hilbert space also carry a ghost number, which can be either an integer or a half-integer, but both situations cannot coexist as the property of having integer or half-integer ghost numbers is eventually tied to the number of constraints (see subsection 14.1.2 of [Henneaux and Teitelboim, 1992]).

As the BRST-formalism has been developed for systems of first-class constraints, one would expect that the space of ‘true’ physical quantum states is a particular subspace of the Hilbert space \mathcal{H} . In the classical Dirac quantization scheme, this subspace had been obtained by imposing Equation (5.6). By Proposition 5.58, since the BRST-invariant extensions correspond to the gauge-invariant functions, the BRST quantization scheme straightforwardly implies that admissible quantum state $|\psi\rangle$ should satisfy the following identity:

$$\Omega |\psi\rangle = 0 \quad (5.109)$$

up to BRST exact terms. By this, we mean that two admissible states $|\psi_1\rangle$ and $|\psi_2\rangle$ should be identified if their difference lies in the image of Ω :

$$|\psi_1\rangle \sim |\psi_2\rangle \quad \text{if and only if} \quad |\psi_1\rangle = |\psi_2\rangle + \Omega |\chi\rangle \quad (5.110)$$

So, the admissible states are BRST-closed, up to BRST-exact terms, where closedness and exactness are defined with respect to the degree 1 nilpotent linear operator Ω . Equation (5.109) and (5.110) define a subspace $\mathcal{S} \subset \mathcal{H}$ of the Hilbert space of quantum states, which is then isomorphic to the cohomology space of the linear operator $\Omega : \mathcal{H} \rightarrow \mathcal{H}$, called the *BRST-state cohomology* and denoted $H_{st}(\Omega)$ for ‘state’. In other words, the physical admissible states correspond to the cohomology classes of the BRST operator Ω ³³.

By Equation (5.108), the BRST observables leave \mathcal{S} invariant. This implies that a BRST-exact linear operator i.e. of the form $A = [K, \Omega]$ vanishes on the subspace \mathcal{S} . This is expected as such operators cohomologically correspond to the zero observables in classical physics. In particular the quantization of the BRST-extensions of the first-class constraints, being BRST-exact (see Equation (5.98)), vanish on \mathcal{S} , which is reminiscent of Equation (5.6). By identity (5.106), the adjoint operator $\text{ad}_\Omega = [\Omega, \cdot]$ defines a differential on the space of commutators. The ad_Ω -closed elements are the BRST observables while the ad_Ω -exact operators are the trivial observables vanishing on the subspace \mathcal{S} . Logically, two BRST observables A_1 and A_2 should be identified if their difference is a trivial observable, in the following sense:

$$A_1 \sim A_2 \quad \text{if and only if} \quad A_1 = A_2 + [K, \Omega]$$

Thus, the BRST observables are parametrized by the cohomology space of the differential $\text{ad}_\Omega : \text{End}(\mathcal{S}) \rightarrow \text{End}(\mathcal{S})$, called the *BRST-operator cohomology* and denoted $H_{op}(\Omega)$ for ‘operator’, in order to distinguish it from the cohomology of $H_{st}(\Omega)$ of admissible quantum states. See subsection 14.1.4 in [Henneaux and Teitelboim, 1992] for more details, and Section 14.2 for general theorems on these cohomologies. In particular Theorem 14.2 which states the following:

³³This result can also be seen at the classical level, where the physical phase space can be recovered from a similar homological condition governed by the classical BRST charge [McMullan, 1992].

Proposition 5.60. *The BRST-operator cohomology $H_{op}(\Omega)$ is isomorphic with the algebra of linear operators acting on the BRST-state cohomology $H_{st}(\Omega)$.*

Remark 5.61. In general one only consider as relevant the operators of ghost number 0 because they correspond to the quantization of BRST-invariant extensions in the classical BRST formalism. However, regarding the degree of the states, one may not impose such a constraint because it happens that the states of the Hilbert space \mathcal{H} have half-integer degrees. See subsection 14.2.5 in [Henneaux and Teitelboim, 1992] for more furnished explanations.

This concludes the overview of BRST quantization of constrained system, and how one avoids several problems posed by first-class constraints when using BRST formalism. Notice that this overview is not at all exhaustive, as we did not address many problems, issues and topics such as states with negative norms, unitarity, Fock spaces etc. These topics have less intrinsic geometric values and we leave the reader interested to know more to study the physical aspects of BRST quantization directly in Chapter 14 of [Henneaux and Teitelboim, 1992] or Chapter 11 of [Rothe and Rothe, 2010].

As a final word, we have so far studied gauging procedures in the canonical Hamiltonian formalism only. This formalism is adapted to non-relativistic classical physical theories. However, when turning to field theories, it is sometimes more efficient to stick to the Lagrangian formalism as Lorentz invariance is manifest in the latter but not in the former. This raises two problems: first, the ‘phase space’ of field theories is infinite dimensional, as the indices labelling the variables (the fields) are not finite anymore but continuous (see Section B.4). Then, we cannot straightforwardly apply the tools of differential and Poisson/symplectic geometry, and the material presented in Chapter 4 has to be adapted. The phase space would be replaced by the space of *field histories* I (see subsection 17.1.2 in [Henneaux and Teitelboim, 1992]), and the constraints are replaced by the Euler-Lagrange equations of motion $\frac{\delta S_0}{\delta \varphi^i}$, whose zero level set define a (infinite dimensional) submanifold in I called the *stationary surface*.

One particularity of the Lagrangian formalism in field theories is that such equations of motions are not independent: the so-called *Noether identities* are dependence relations between them (see Chapter 3 in [Henneaux and Teitelboim, 1992], Chapter 2 in [Rothe and Rothe, 2010] or subsection 2.2 in [Barnich and Del Monte, 2018]):

$$R_\alpha^i \frac{\delta S_0}{\delta \varphi^i} = 0$$

The main point is that to each Noether identity corresponds a gauge transformation (this is the content of Noether’s second theorem). Since there always exist trivial Noether identities, for which $R_\alpha^i = M_\alpha^{ij} \frac{\delta S_0}{\delta \varphi^j}$ where M_α^{ij} is skew-symmetric in i, j , field theories necessarily admit gauge transformations (be they trivial). To preserve manifest Lorentz invariance (and possibly other convenient properties such as locality), these theories then require a treatment of gauge transformations in the Lagrangian formalism.

However, we do not have a proper phase space in the Lagrangian picture, as there is no canonical symplectic structure on the space of histories (or TQ even). Although there is a canonical Poisson bracket on the on-shell gauge invariant functions, called the Peierls bracket, it is not a good idea to restrict oneself to these physical observables because we can lose manifest physical symmetries of properties of the Lagrangian. Coincidentally, there exists a way of extending the space of field histories in a way similar to what happens in the BRST formalism by adding extra anti-commuting variables called the *antifields*. To each field φ^i corresponds an antifield φ_i^* to which we assign a pure antighost number $+1$ and ghost number -1 . Then, to each Noether identity – equivalently, gauge generator – we associate a ghost C^α of pure antighost number 0 and pure ghost number $+1$. To these ghosts we associate an antifield C_α^* of pure

antighost number 2 and ghost number -2 . These antifields of ghosts – which are *not* what we call (Faddeev-Popov) *antighosts*, see Section 6.2 in [Gomis et al., 1995] – will play the same role as ghost momenta play in BRST formalism (meaning: they define a resolution of the algebra of longitudinal forms). If the Noether identities are reducible – i.e. if the functions R_α^i are not functionally independent – then there exist additional ghosts of ghosts of higher ghost number, together with their associated antifields. Notice that the grading of the ghost and their ghost momenta is symmetric with respect to $1/2$, and not 0 as in the Hamiltonian BRST formalism. The resulting extended space is a graded (infinite dimensional) manifold.

On this extended space, one can define two kinds of differentials, similar to δ (not the same δ as in $\frac{\delta S_0}{\delta \varphi^i}$) and d introduced in the BRST formalism in Section 5.2:

$$\begin{aligned}\delta \varphi_i^* &= \frac{\delta S_0}{\delta \varphi^i}, & \delta C_\alpha^* &= R_\alpha^i \varphi_i^*, & \delta \phi_i &= \delta C^\alpha = 0 \\ d\varphi^i &= R_\alpha^i, & dC^\alpha &= -\frac{1}{2} f^\alpha{}_{\beta\gamma} C^\beta C^\gamma\end{aligned}$$

Then, as in the classical BRST formalism, one can find a differential $s = \delta + d + \dots$ on the functions of the extended phase space (hence depending on the fields φ^i , the ghosts C_α and the antifields $\varphi_i^*, C^{\alpha*}$, and possibly more ghosts and ghosts and their associated antifields), such that a theorem similar to Theorem 5.12 is satisfied. Notice that this differential is not a derivation. It turns out that the extended phase space can be equipped with a -1 -graded Poisson algebra structure³⁴, called the *antibracket* and usually denoted (\cdot, \cdot) . It has particular symmetries corresponding to this shifted Poisson structure and is compatible with the differential s .

Then it turns out that the differential s is antibracket-exact, i.e. there exist a function S of ghost number 0 on the extended phase space such that:

$$s = (S, \cdot)$$

This function contains the classical action S_0 at the zero-th order in pure antighost number, i.e. $S = S_0 +$ terms of higher pure antighost number. The cohomological condition $s^2 = 0$ becomes:

$$(S, S) = 0$$

which is called the *classical master equation* as in Equation (5.64) for the BRST charge. A solution of this master equation 5.5 is called a *proper solution* and the first two terms are:

$$S = S_0 + \varphi_i^* R_\alpha^i C^\alpha + \dots$$

The higher order terms encode how complicated the algebra of gauge symmetries is.

The antibracket on the functions of the extended phase space, together with this proper solution, opens the possibility to treat the gauge symmetries of a field theory in the Lagrangian formalism following similar lines as in the Hamiltonian BRST formalism. This procedure is called the *antifield formalism*, or the *Batalin-Vilkovisky formalism*, and is useful to provide a well-behaved path-integral to certain physical models. This formalism is quite interesting and has profound ramifications in field theories (Schwinger-Dyson equations, Slavnov-Taylor identities and Zinn-Justin equation) as well as in mathematical physics (quantum master equation, BV-algebras) but is beyond the scope of the present lecture. We refer to Chapters 17 and

³⁴Although the antibracket has ghost number 1, it should be understood as the bracket of a Gerstenhaber algebra, i.e. as a degree -1 Lie bracket. This grading may be obtained by reversing the ghost number of the ghosts and the antifields, although it would be a bit contradictory. Moreover, the algebraic structure of the space of functions on the extended phase space is what is called a *BV-algebra* [Cattaneo et al., 2006]. The letters B and V come from theoretical physicists Batalin and Vilkovisky.

18 of [Henneaux and Teitelboim, 1992] and Chapter 12 of [Rothe and Rothe, 2010] to have a thorough introduction to the topic from a physical perspective, or [Henneaux, 1990] for an earlier, alternative pedagogical presentation. On the other hand, the lecture notes [Mnev, 2017] and [Barnich and Del Monte, 2018] form quite a good introduction to the BV formalism from a mathematical perspective. Eventually, the review [Gomis et al., 1995] offers a complete overview of the topic and gauge theories in general, with many examples. These **lecture notes** provide a shorter and alternative presentation of the BRST-BV formalism, while reference [Gómez, 2016] presents an insightful interpretation from the perspective of perturbation theory. Eventually Chapter 19 of [Henneaux and Teitelboim, 1992] for a complete quantization of the electromagnetic field.

A Mathematical background in linear algebra

Let E be a (real) vector space of dimension n . Then it is isomorphic to \mathbb{R}^n . A basis of E is a set of n vectors – say e_1, \dots, e_n – that are linearly independent and that generate the whole vector space.

A.1 The tensor algebra, the symmetric algebra and the exterior algebra

The *tensor algebra* of E – denoted $T(E)$ – is an infinite family of vector spaces $T^0(E), T^1(E), T^2(E), \dots$ defined recursively as:

$$T^0(E) = \mathbb{R} \quad \text{and, for all } m \geq 0 \quad T^{m+1}(E) = E \otimes T^m E$$

with the convention that $\mathbb{R} \otimes E = E \otimes \mathbb{R} = E$. The symbol \otimes symbolizes a sort of multiplication, not between scalars but between vectors – or more generally tensors, hence the name. More precisely, this tensor product possesses the associativity and distributivity properties of the multiplication operator:

$$\begin{aligned} x_1 \otimes (x_2 \otimes x_3) &= (x_1 \otimes x_2) \otimes x_3 = x_1 \otimes x_2 \otimes x_3 \\ x_1 \otimes \dots \otimes (x_i + y) \otimes \dots \otimes x_m &= (x_1 \otimes \dots \otimes x_i \otimes \dots \otimes x_m) + (x_1 \otimes \dots \otimes y \otimes \dots \otimes x_m) \end{aligned}$$

for every $1 \leq i \leq m$ and every vectors $x_1, \dots, x_m, y \in E$. Notice however that the tensor product \otimes is *not* commutative, contrary to the usual multiplication on scalars. Since we are working on vector spaces, we assume that it is linear in every variable, that is, given any scalar $\lambda \in \mathbb{R}$:

$$\lambda(x_1 \otimes \dots \otimes x_m) = (\lambda x_1) \otimes \dots \otimes x_m = x_1 \otimes \dots \otimes (\lambda x_i) \otimes \dots \otimes x_m = x_1 \otimes \dots \otimes (\lambda x_m)$$

Thus, elements of the m -th tensor power of E – denoted $T^m(E)$ or sometimes $E^{\otimes m}$ – are literally products of vectors of E . This has to be contrasted (and not to be confused) with the cartesian product $E \times \dots \times E$ where multiplication by a scalar satisfies:

$$\lambda(x_1, \dots, x_m) = (\lambda x_1, \lambda x_2, \dots, \lambda x_m)$$

and where distributivity over addition is not satisfied:

$$\forall 1 \leq i \leq m \quad (x_1, \dots, x_i + y, \dots, x_m) = (x_1, \dots, x_i, \dots, x_m) + (0, \dots, y, \dots, 0)$$

This comes from the fact that the cartesian product $E \times \dots \times E$ actually corresponds to the direct sum $E^{\oplus m} = E \oplus \dots \oplus E$ (m -times). This discussion shows that $T^m(E)$ is of dimension n^m , whereas $E^{\oplus m}$ is of dimension $n \times m$. A basis of $T^m(E)$ is explicitly given by the following tensor products:

$$\left\{ e_{i_1} \otimes \dots \otimes e_{i_m} \mid 1 \leq i_1, \dots, i_m \leq n \right\} \quad (\text{A.1})$$

Additionally, associativity of the tensor product implies that:

$$T^k(E) \otimes T^l(E) \subset T^{k+l}(E) \quad (\text{A.2})$$

An algebra that is a graded vector space and whose product satisfies a similar condition as Equation (A.2) is called a *graded algebra*:

Definition A.1. A graded vector space is a family of vector spaces $E = (E_i)_{i \in \mathbb{Z}}$, indexed over \mathbb{Z} (not all E_i need be non-zero). The indices are integers and called degrees, and are denoted $|x| = i$ for any homogeneous element $x \in E_i$. We say that E is non-negatively graded (resp. non-positively graded) if $E = (E_i)_{i \geq 0}$ (resp. $E = (E_i)_{i \leq 0}$).

A graded algebra is a graded vector space $A = \bigoplus_{i \in \mathbb{Z}} A_i$ equipped with an associative \mathbb{R} -bilinear operation $\cdot : A \times A \rightarrow A$ which satisfies:

$$A_i \cdot A_j \subset A_{i+j}$$

Example A.2. A vector space E is a graded vector space where every $E_i = 0$ for $i \neq 0$ but $E_0 = E$.

Example A.3. The tensor algebra is a graded algebra, in which the grading corresponds to the length of the basis elements. This graded algebra is non-negatively graded.

Remark A.4. Remark 2.2.1 in [Mehta, 2006] explains that the correct class of graded vector spaces that we should consider are those whose grading is bounded below and above, and who are finite dimensional at every degree – these spaces are said to be of *finite dimension* [Kotov and Salnikov, 2021]. These conditions indeed ensure that their class is stable under the tensor product.

The tensor algebra $T(E)$ contains two particular subspaces³⁵: the one formed by linear combinations of fully symmetrized basis elements of $T(E)$ – it is the *symmetric algebra* $S(E)$, and the one formed by linear combinations of fully anti-symmetrized basis elements of $T(E)$ – it is the *exterior algebra* $\bigwedge^\bullet(E)$. Both will be graded algebra, with respect to their product.

Remark A.5. When we write a bullet \bullet as an index or an exponent we want to emphasize that the space is graded, e.g. $\bigwedge^\bullet(E) = \bigwedge^0(E) \oplus \bigwedge^1(E) \oplus \dots \oplus \bigwedge^n(E)$.

Both the symmetric algebra and the exterior algebra are actually graded sub-vector spaces of $T(E)$, that is to say: they both decompose as a family of vector spaces $S(E) = \bigoplus_{m=0}^{\infty} S^m(E)$ and $\bigwedge^\bullet(E) = \bigoplus_{m=0}^n \bigwedge^m(E)$, which are such that $S^m(E), \bigwedge^m(E) \subset T^m(E)$, for every $m \geq 0$. The graded space $S(E)$ is the subspace of $T(E)$ that is invariant under the action of any permutation σ on the labels of the basis vectors. More precisely, for every $m \geq 1$, the space $S^m(E)$ is generated (as a vector subspace of $T^m(E)$) by the following elements:

$$e_{i_1} \odot e_{i_2} \odot \dots \odot e_{i_m} = \frac{1}{m!} \sum_{\sigma \in S_m} e_{i_{\sigma(1)}} \otimes e_{i_{\sigma(2)}} \otimes \dots \otimes e_{i_{\sigma(m)}} \quad (\text{A.3})$$

The symmetrized product \odot symbolizes that the tensor $e_{i_1} \odot e_{i_2} \odot \dots \odot e_{i_m}$ is invariant under the action of any permutation of m elements $\sigma \in S_m$. In particular, invariance under the permutation (1 2) reads:

$$e_{i_1} \odot e_{i_2} = e_{i_2} \odot e_{i_1}$$

Hence the symmetric product is *commutative*. Any other combination of permutations leaves the product unchanged. The graded space $S(E)$ equipped with the product \odot is a (commutative) graded algebra because it satisfies a similar condition as Equation (A.2):

$$S^k(E) \odot S^l(E) \subset S^{k+l}(E)$$

Counting the number of ways one can choose m elements (with possible repetitions) among n basis vectors in order to construct the basis elements defined in Equation (A.3), one can check that one obtains all the basis elements by restricting oneself to $1 \leq i_1 \leq i_2 \leq \dots \leq i_m \leq n$:

$$\left\{ e_{i_1} \odot \dots \odot e_{i_m} \mid 1 \leq i_1 \leq i_2 \leq \dots \leq i_m \leq n \right\} \quad (\text{A.4})$$

Then, the dimension of the space $S^m(E)$ is $\binom{n+m-1}{m}$, thus one can see that it increases with m . The symmetric algebra is thus infinite dimensional, as is the tensor algebra.

³⁵Actually the symmetric algebra and the exterior algebra are quotient of the tensor algebra, but there exists a canonical isomorphisms between those and the subspaces of E that we describe.

The exterior algebra, on the other hand, is generated (as a vector space) by elements of $T(E)$ invariant under *signed* permutations. Let us explain what it means. For every $m \geq 1$, the space $\bigwedge^m(E)$, whose elements are called *m-vectors* or *multivectors*, is generated (as a vector subspace of $T^m(E)$) by the following elements:

$$e_{i_1} \wedge e_{i_2} \wedge \dots \wedge e_{i_m} = \sum_{\sigma \in S_m} (-1)^\sigma e_{i_{\sigma(1)}} \otimes e_{i_{\sigma(2)}} \otimes \dots \otimes e_{i_{\sigma(m)}} \quad (\text{A.5})$$

where $(-1)^\sigma$ is the signature of the permutation σ . Using the *Levi-Civita symbol* $\epsilon_{\sigma(1)\dots\sigma(m)} = (-1)^\sigma \epsilon_{1\dots m}$, set with the convention that $\epsilon_{1\dots m} = 1$, one obtains the alternative, more physicists oriented, formula:

$$e_{i_1} \wedge e_{i_2} \wedge \dots \wedge e_{i_m} = \sum_{\sigma \in S_m} \epsilon_{\sigma(1)\dots\sigma(m)} e_{i_{\sigma(1)}} \otimes e_{i_{\sigma(2)}} \otimes \dots \otimes e_{i_{\sigma(m)}}$$

In particular the first few elements are:

$$\begin{aligned} \text{for } m = 0 & \quad \bigwedge^0(E) \simeq \mathbb{R} \\ \text{for } m = 1 & \quad \bigwedge^1(E) \simeq E \\ \text{for } m = 2 & \quad e_i \wedge e_j = e_i \otimes e_j - e_j \otimes e_i \\ \text{for } m = 3 & \quad e_i \wedge e_j \wedge e_k = e_i \otimes e_j \otimes e_k + e_j \otimes e_k \otimes e_i + e_k \otimes e_i \otimes e_j \\ & \quad - e_i \otimes e_k \otimes e_j - e_k \otimes e_j \otimes e_i - e_j \otimes e_i \otimes e_k \end{aligned}$$

There exists another convention, which is such that $x \wedge y = \frac{1}{2}(x \otimes y - y \otimes x)$ but this is not convenient for geometrical purposes, but which is the natural product when the exterior algebra is obtained through a quotient of the tensor algebra. These subtleties are discussed at large in Chapter 12 of [Lee, 2003] (Chapter 14 in the 2012 edition).

The *wedge product* \wedge is defined so that the tensor $e_{i_1} \wedge e_{i_2} \wedge \dots \wedge e_{i_m}$ is invariant under any *signed* permutation $(-1)^\sigma \sigma$ of m elements. For any permutation $\sigma \in S_m$, the general formula is the following:

$$e_{i_1} \wedge \dots \wedge e_{i_m} = (-1)^\sigma e_{i_{\sigma(1)}} \wedge \dots \wedge e_{i_{\sigma(m)}}$$

or, using the Levi-Civita symbol:

$$e_{i_1} \wedge \dots \wedge e_{i_m} = \epsilon_{\sigma(1)\dots\sigma(m)} e_{i_{\sigma(1)}} \wedge \dots \wedge e_{i_{\sigma(m)}}$$

where here, the Einstein summation convention is *not* used! To illustrate these rather abstract formulas, let us pick up the transposition (1 2) (of signature -1). Then, invariance of the bivector $e_{i_1} \wedge e_{i_2}$ under the action of the signed permutation $-(1\ 2)$ reads:

$$e_{i_1} \wedge e_{i_2} = -e_{i_2} \wedge e_{i_1} \quad (\text{A.6})$$

The minus sign on the right hand side is the signature of the transposition (1 2). Another example σ is the circular permutation (1 2 3) (of signature $+1$), which is such that e_{i_1} becomes e_{i_2} , e_{i_2} becomes e_{i_3} and e_{i_3} becomes e_{i_1} . This (signed) permutation leaving the trivector $e_{i_1} \wedge e_{i_2} \wedge e_{i_3}$ invariant means that:

$$e_{i_1} \wedge e_{i_2} \wedge e_{i_3} = e_{i_2} \wedge e_{i_3} \wedge e_{i_1} \quad (\text{A.7})$$

More generally, the rule of calculus in the exterior algebra is that, when permuting two elements, a sign appears only when the signature of the chosen transposition is -1 . In particular,

since it is often difficult to know the signature of a permutation, and since any permutation can be obtained from a sequence of transpositions (permutation of two elements), permuting elements two by two while multiplying by -1 until reaching the image of the desired (signed) permutation is a good technique to obtain the correct sign. Let us illustrate by looking up at the permutation $(1\ 3\ 2) = (1\ 3)(3\ 2)$ (its parity is even so its signature is $+1$). By first using the transposition $(3\ 2)$, and then $(1\ 3)$ one obtains:

$$e_{i_1} \wedge e_{i_2} \wedge e_{i_3} = -e_{i_1} \wedge e_{i_3} \wedge e_{i_2} = e_{i_3} \wedge e_{i_1} \wedge e_{i_2}$$

We can now study the (signed) action of the cycle $(1\ 2\ 3\ 4) = (1\ 2)(2\ 3)(3\ 4)$ (of signature -1) on $e_{i_1} \wedge e_{i_2} \wedge e_{i_3} \wedge e_{i_4}$, which can be obtained through three transpositions:

$$e_{i_1} \wedge e_{i_2} \wedge e_{i_3} \wedge e_{i_4} = -e_{i_1} \wedge e_{i_2} \wedge e_{i_4} \wedge e_{i_3} = e_{i_1} \wedge e_{i_3} \wedge e_{i_4} \wedge e_{i_2} = -e_{i_2} \wedge e_{i_3} \wedge e_{i_4} \wedge e_{i_1} \quad (\text{A.8})$$

Since the permutation $(1\ 2\ 3\ 4)$ is such that e_{i_1} becomes e_{i_2} , e_{i_2} becomes e_{i_3} , e_{i_3} becomes e_{i_4} and e_{i_4} becomes e_{i_1} , one observe that the sign in the right hand side of Equation (A.8) tells us that the parity of $(1\ 2\ 3\ 4)$ is odd. Additionally, we see that the action of the signed permutation $-(1\ 2\ 3\ 4)$ leaves $e_{i_1} \wedge e_{i_2} \wedge e_{i_3} \wedge e_{i_4}$ invariant.

Another efficient way of managing cyclic permutations – instead of decomposing them – is to take the leftmost element, and make it go right through all the terms, so that at each transposition with its neighbor, one adds a minus sign. At each step, we use Equation (A.6) so that we ensure that all expressions are equal. For example the signed action of $(-1)^{k-1}(1\ 2\ \dots\ k-1\ k)$ leaves the multivector $e_{i_1} \wedge \dots \wedge e_{i_m}$ invariant, and that can be shown by making e_{i_1} goes right through the $k-1$ vectors on its right:

$$\begin{aligned} e_{i_1} \wedge e_{i_2} \wedge \dots \wedge e_{i_k} \wedge \dots \wedge e_{i_m} &= -e_{i_2} \wedge e_{i_1} \wedge e_{i_3} \wedge \dots \wedge e_{i_k} \wedge \dots \wedge e_{i_m} \\ &= e_{i_2} \wedge e_{i_3} \wedge e_{i_1} \wedge \dots \wedge e_{i_k} \wedge \dots \wedge e_{i_m} \\ &= (-1)^{k-2} e_{i_2} \wedge \dots \wedge e_{i_{k-1}} \wedge e_{i_1} \wedge e_{i_k} \wedge \dots \wedge e_{i_m} \\ &= (-1)^{k-1} e_{i_2} \wedge \dots \wedge e_{i_k} \wedge e_{i_1} \wedge e_{i_{k+1}} \wedge \dots \wedge e_{i_m} \end{aligned}$$

The general rule is that for cyclic permutations of the form $(1\ 2\ \dots\ k-1\ k)$ the parity is the same as the parity of the integer $k-1$. It is as if the vector e_{i_1} had jumped over $k-1$ elements to get in the right place. This strategy could have been used in Equations (A.6), (A.7) and (A.8), where we obtain that the parity of a transposition is odd, the parity of a circular permutation of three elements is even, whereas the parity of a circular permutation of four elements is odd.

The properties of the wedge product implies in particular that for every $x \in E$, the bivector $x \wedge x$ is zero (in the vector space $\wedge^2(E)$). Thus, as soon as the same element of E appears twice in a multivector, then it is automatically zero. For example, let x_1, \dots, x_m be m linearly independent vectors of E (so in particular $1 \leq m \leq n$), then:

$$x_1 \wedge x_2 \wedge \dots \wedge x_m \neq 0$$

In the case that one of the x_i is a linear combination of the others, say $x_i = \sum_{j \neq i} \alpha_j x_j$, then the m -vector is zero, since:

$$x_1 \wedge \dots \wedge x_m = \sum_{j \neq i} \alpha_j \underbrace{x_1 \wedge \dots \wedge x_{i-1} \wedge x_j \wedge x_{i+1} \wedge \dots \wedge x_m}_{= 0}$$

This has a tremendous consequence: contrary to the symmetric algebra, the exterior algebra is bounded above. A multivector cannot be composed by more than m vectors, for otherwise it vanishes. Hence, contrary to the symmetric algebra, the exterior algebra is of finite dimension.

Due to the fact that the wedge product of two identical elements vanish, one can check that all the basis elements of $\wedge^m(E)$ are obtained by restricting oneself to $1 \leq i_1 < i_2 < \dots < i_m \leq n$, that is to say a basis is formed by the following multivectors:

$$\left\{ e_{i_1} \wedge \dots \wedge e_{i_m} \mid 1 \leq i_1 < i_2 < \dots < i_m \leq n \right\} \quad (\text{A.9})$$

Then one deduces that the dimension of the vector space $\wedge^m(E)$ is $\binom{n}{m}$. One can check that such a dimension is minimal and equal to 1 for $m = 0$ (i.e. when $\wedge^0(E) = \mathbb{R}$) and for $m = n$ (i.e. when $\wedge^n(E)$ is the one-dimensional vector subspace of $T(E)$ generated by the element $e_1 \wedge e_2 \wedge \dots \wedge e_n$). The direct sum $\wedge^\bullet(E) = \bigoplus_{m=0}^n \wedge^m(E)$ is then finite dimensional of total dimension 2^n .

Additionally, the definition of the wedge product has been made so that we have the following property:

$$(e_{i_1} \wedge \dots \wedge e_{i_p}) \wedge (e_{i_{p+1}} \wedge \dots \wedge e_{i_m}) = e_{i_1} \wedge \dots \wedge e_{i_m} \quad (\text{A.10})$$

In particular, the product is associative. This allows us to compute the wedge product of a k -multivectors and l -multi vectors. Notice that the wedge product satisfies Equation (A.10) precisely because of the absence of any scaling factor on the right hand side of Equation (A.5). The wedge product then defines a graded algebra structure on the exterior algebra (hence justifying the name), that is:

$$\wedge^k(E) \wedge \wedge^l(E) \subset \wedge^{k+l}(E)$$

More precisely, for any $\alpha \in \wedge^k(E)$ and any $\beta \in \wedge^l(E)$, then one has $\alpha \wedge \beta \in \wedge^{k+l}(E)$, and it satisfies the following identity:

$$\alpha \wedge \beta = (-1)^{kl} \beta \wedge \alpha$$

We say that the product is *graded commutative*, and the exterior algebra is thus a (*graded commutative graded algebra*).

Exercise A.6. The proof is left as an exercise.

Recall that the *dual* of the vector space E is the space denoted E^* of all *linear forms on E* , i.e. all the linear maps $\varphi : E \rightarrow \mathbb{R}$. While elements of E are called *vectors*, elements of E^* are called *covectors*. Given a basis e_1, \dots, e_n of E there is a privileged choice of a basis on E^* : the set of linear maps $e^1, \dots, e^n : E \rightarrow \mathbb{R}$, that are such that:

$$e^i(e_j) = \delta_j^i \quad (\text{A.11})$$

where here δ_j^i denotes the Kronecker symbol³⁶. Such a choice of basis on E^* can always be made. Notice the localization of the labels i, j : as indices on vectors, as exponents on covectors. This has some importance, and is related to Einstein summation convention: for example, imagine you have a vector $v = v^i e_i$ and a covector $\varphi = \varphi_j e^j$. In particular $\varphi \in E^*$ and can be understood as a linear form $\varphi : E \rightarrow \mathbb{R}$ which can act on v and define a real number (we assume Einstein summation convention throughout):

$$\varphi(v) = \varphi_j e^j(v^i e_i) = \varphi_j v^i e^j(e_i) = \varphi_j v^i \delta_i^j = \varphi_i v^i$$

We passed from the second term to the third by linearity of the dual basis. In the last implicit sum on the right-hand side, we say that the upper and lower indices have been *contracted*. The

³⁶Some author use the notation g_j^i instead. Moreover, some authors consider that the real Kronecker symbol is the one with one index up and one index down. In that case, when they write δ_{ij} they mean g_{ij} . We will try to use this convention in the present paper.

result should be an object which does not carry any index, which is precisely the case of the real number $\varphi(v)$.

One can define the tensor algebra of the dual E^* and since we are in finite dimension, $T(E^*) \simeq (T(E))^*$. The dual basis of this dual vector space can be obtained from the dual basis e^1, \dots, e^n and the definition of the tensor product. The action of the dual element $e^i \otimes e^j$ on $e_k \otimes e_l$ is given by the following:

$$e^i \otimes e^j(e_k \otimes e_l) = \delta_k^i \delta_l^j$$

This is equal to +1 if and only if $k = i$ and $l = j$. From this we deduce that the dual basis to the basis of $T(E)$ (see Equation (A.1)) is made of the following tensor products:

$$\left\{ e^{i_1} \otimes \dots \otimes e^{i_m} \mid 1 \leq i_1, \dots, i_m \leq n \right\} \quad (\text{A.12})$$

That is to say:

$$e^{i_1} \otimes \dots \otimes e^{i_m}(e_{j_1} \otimes \dots \otimes e_{j_m}) = \delta_{j_1}^{i_1} \dots \delta_{j_m}^{i_m} \quad (\text{A.13})$$

One can also define a symmetric algebra and an exterior algebra associated to the dual E^* , and we have the following isomorphisms because E is finite dimensional: $S(E^*) \simeq (S(E))^*$ and $\Lambda^\bullet(E^*) \simeq (\Lambda^\bullet(E))^*$. Notice however that the most obvious basis of $S(E^*)$ and $\Lambda^\bullet(E^*)$ are not the dual basis of (A.4) and (A.9). Indeed, using the definition of the symmetric product (see Equation (A.3)) on the dual basis e^1, \dots, e^n of E^* , one obtains a basis of $S(E^*)$, denoted by vectors of the form $e^{i_1} \odot \dots \odot e^{i_m}$ for $1 \leq i_1 \leq \dots \leq i_m \leq n$. However this basis is not dual to the basis of $S(E)$ given in Equation (A.4), for:

$$e^i \odot e^j(e_k \odot e_l) = \frac{e^i \otimes e^j + e^j \otimes e^i}{2} \frac{e_k \otimes e_l + e_l \otimes e_k}{2} = \frac{(\delta_k^i \delta_l^j + \delta_l^i \delta_k^j)}{2}$$

which is equal to $\frac{1}{2}$ when $k = i$ and $l = j$. Thus the element of $S^2(E^*)$ that would be considered to be dual to $e_i \odot e_j$ is $2e^i \odot e^j$. More generally a dual basis to the basis (A.4) is given by:

$$\left\{ m!(e^{i_1} \odot \dots \odot e^{i_m}) \mid 1 \leq i_1 \leq i_2 \leq \dots \leq i_m \leq n \right\}$$

This also forms a basis of $S(E^*)$, but one has to remember the factor. The same phenomenon occurs for $\Lambda^\bullet(E^*)$. Using the definition of the symmetric product (see Equation (A.5)) on the dual basis e^1, \dots, e^n of E^* , one obtains a basis of $\Lambda^\bullet(E^*)$, denoted by vectors of the form $e^{i_1} \wedge \dots \wedge e^{i_m}$ for $1 < i_1 < \dots < i_m < n$. However this basis is not dual to the basis of $\Lambda^\bullet(E)$ given in Equation (A.9), for:

$$e^i \wedge e^j(e_k \wedge e_l) = (e^i \otimes e^j - e^j \otimes e^i)(e_k \otimes e_l - e_l \otimes e_k) = 2(\delta_k^i \delta_l^j - \delta_l^i \delta_k^j) \quad (\text{A.14})$$

which is equal to 2 when $k = i$ and $l = j$. Thus the element of $\Lambda^2(E^*)$ that would be considered to be dual to $e_i \wedge e_j$ is $\frac{1}{2}e^i \wedge e^j$. More generally a dual basis to the basis (A.4) is given by:

$$\left\{ \frac{1}{m!}(e^{i_1} \wedge \dots \wedge e^{i_m}) \mid 1 \leq i_1 < i_2 < \dots < i_m \leq n \right\}$$

Let us now conclude this subsection by discussing the role of $T(E^*)$, $S(E^*)$ and $\Lambda^\bullet(E^*)$. The main point is that the tensor algebra of the dual, denoted $T(E^*)$, can be considered to be the vector space of *multi-linear forms on E* . Linear forms on E form the dual space $T^1(E^*) = E^*$. Bilinear forms on E are those functions $B : E \times E \rightarrow \mathbb{R}$ such that on the one hand it is linear in the first variable $B(\lambda x + \mu y, z) = \lambda B(x, z) + \mu B(y, z)$ (for every $\lambda, \mu \in \mathbb{R}$ and $x, y, z \in E$), and on the other hand a similar identity holds for the second variable. Bilinear forms on E

are precisely the elements of $T^2(E^*) = E^* \otimes E^* \simeq (E \otimes E)^*$. A priori bilinear forms may not be symmetric nor antisymmetric. More generally, a m multilinear form on E is an element $\Theta : E \times E \times \dots \times E \longrightarrow \mathbb{R}$ which is linear in each of its variables:

$$\Theta(x_1, x_2, \dots, \lambda x_k + \mu y, \dots, x_m) = \lambda \Theta(x_1, x_2, \dots, x_k, \dots, x_m) + \mu \Theta(x_1, x_2, \dots, y, \dots, x_m)$$

Although the following result is computational, it is important, and the proof is useful to understand how vectors and covectors interact.

Proposition A.7. *There is a canonical isomorphism between m multilinear forms on E and the elements of $T^m(E^*)$.*

Proof. Let Θ be a m multilinear form. Then evaluating it on a set of basis vectors e_{i_1}, \dots, e_{i_m} gives a real number $\Theta(e_{i_1}, \dots, e_{i_m})$ that we denote by $\Theta_{i_1 \dots i_m}$. Repeating the process for every combination of m basis vectors of E , one obtains a family of real numbers. Then, the object $\Theta_{i_1 \dots i_m} e^{i_1} \otimes \dots \otimes e^{i_m}$ (Einstein summation convention implied) is an element of $T^m(E^*)$.

Conversely, let Θ be an element of $T^m(E^*)$, and let us decompose it on the basis (A.12):

$$\Theta = \Theta_{i_1 \dots i_m} e^{i_1} \otimes \dots \otimes e^{i_m}$$

where $\Theta_{i_1 \dots i_m} \in \mathbb{R}$ and where the Einstein summation convention has been used. Then, picking up m vectors $x_1 = x_1^{j_1} e_{j_1}, x_2 = x_2^{j_2} e_{j_2}, \dots, x_m = x_m^{j_m} e_{j_m} \in E$ (Einstein summation convention implied on repeated indices) one can write, using Equation (A.13):

$$\begin{aligned} \Theta_{i_1 \dots i_m} e^{i_1} \otimes \dots \otimes e^{i_m} (x_1 \otimes \dots \otimes x_m) &= \Theta_{i_1 \dots i_m} x_1^{j_1} \dots x_m^{j_m} e^{i_1} \otimes \dots \otimes e^{i_m} (e_{j_1} \otimes \dots \otimes e_{j_m}) \\ &= \Theta_{i_1 \dots i_m} x_1^{j_1} \dots x_m^{j_m} \delta_{j_1}^{i_1} \dots \delta_{j_m}^{i_m} \\ &= \Theta_{i_1 \dots i_m} x_1^{i_1} \dots x_m^{i_m} \end{aligned}$$

We define this real number to be $\Theta(x_1, \dots, x_m)$. One can check that the assignment $(x_1, \dots, x_m) \longmapsto \Theta(x_1, \dots, x_m)$ is linear in each of its variable. Thus, Θ can be seen as a m multilinear form on E . \square

Now recall that – although it is not mathematically totally rigorous – we consider the symmetric algebra and the exterior algebra as subspaces of the tensor algebra. How do they fit in the picture? It turns out that $S^m(E^*)$ is the space of m multilinear forms that are fully symmetric, that is to say, those $\Xi \in T^m(E^*)$ such that, for any choice of permutation $\sigma \in S^m$:

$$\Xi(x_{i_1}, x_{i_2}, \dots, x_{i_m}) = \Xi(x_{\sigma(i_1)}, x_{\sigma(i_2)}, \dots, x_{\sigma(i_m)})$$

By Proposition A.7, the action of a symmetric m multilinear form on a set of m vectors x_1, \dots, x_m can be written as follows:

$$\Xi(x_1, \dots, x_m) = \Xi(x_1 \otimes \dots \otimes x_m) \tag{A.15}$$

when, on the right hand side, we understand that Ξ has been developed on the basis (A.12) of $T^m(E^*)$ and Equation (A.13) is used. On the other hand, the space $\wedge^m(E^*)$ is the space of m multilinear forms that are fully anti-symmetric, that is to say, those $\Omega \in T^m(E^*)$ such that, for any choice of permutation $\sigma \in S^m$:

$$\Omega(x_{i_1}, x_{i_2}, \dots, x_{i_m}) = (-1)^\sigma \Omega(x_{\sigma(i_1)}, x_{\sigma(i_2)}, \dots, x_{\sigma(i_m)})$$

where $(-1)^\sigma$ is the signature of σ . For example, given a bilinear form $B : E \times E \longrightarrow \mathbb{R}$, the bilinear form A defined by:

$$A(x, y) = B(x, y) - B(y, x)$$

is fully antisymmetric because $A(x, y) = -A(y, x)$. By Proposition A.7, the action of a fully antisymmetric m multilinear form $\Omega \in \wedge^m(E^*)$ on a set of m vectors x_1, \dots, x_m can be written as follows:

$$\Omega(x_1, \dots, x_m) = \Omega(x_1 \otimes \dots \otimes x_m) \quad (\text{A.16})$$

when, on the right hand side, we understand that Ξ has been developed on the basis (A.12) of $T^m(E^*)$ and Equation (A.13) is used. We often call the fully anti-symmetric multilinear forms on E *alternating tensors*.

Exercise A.8. Given a trilinear form $T \in T^3(E^*)$, check that the trilinear form R defined by:

$$R(x, y, z) = T(x, y, z) + T(y, z, x) + T(z, x, y) - T(x, z, y) - T(z, y, x) - T(y, x, z)$$

is fully antisymmetric.

Last but not least, let us give a formula to compute the value of an alternating tensor, when fed with a bunch of vectors. For every $1 \leq m \leq n$, it is only defined on *decomposable* elements of $\wedge^m(E^*)$, i.e. those of the form $\varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_m$, for some $\varphi_i \in E^*$. Here, in particular, the index is not a coordinate index. Pick up such a decomposable element, then one has:

$$\varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_m(x_1, \dots, x_m) = \det \begin{pmatrix} \varphi_1(x_1) & \varphi_1(x_2) & \dots & \varphi_1(x_{m-1}) & \varphi_1(x_m) \\ \varphi_2(x_1) & & & & \varphi_2(x_m) \\ \dots & & \dots & & \dots \\ \varphi_{m-1}(x_1) & & & & \varphi_{m-1}(x_m) \\ \varphi_m(x_1) & \varphi_m(x_2) & \dots & \varphi_m(x_{m-1}) & \varphi_m(x_m) \end{pmatrix} \quad (\text{A.17})$$

for every $x_1, \dots, x_m \in E$. This formula coincides with Equation (A.16) when $\Omega = \varphi_1 \wedge \dots \wedge \varphi_m$. Applying this formula to a decomposable alternating 2-tensor $\varphi_1 \wedge \varphi_2$, one has:

$$\varphi_1 \wedge \varphi_2(x_1, x_2) = \varphi_1(x_1)\varphi_2(x_2) - \varphi_1(x_2)\varphi_2(x_1) \quad (\text{A.18})$$

This equation, when $\varphi_1 \wedge \varphi_2 = e^i \wedge e^j$, and when $x_1 = e_k$ and $x_2 = e_l$, gives:

$$e^i \wedge e^j(e_k, e_l) = \delta_k^i \delta_l^j - \delta_l^i \delta_k^j \quad (\text{A.19})$$

which is precisely one half of Equation (A.14). Thus it gives +1 when $k = i$ and $l = j$. This shows that looking at elements of $\wedge^\bullet(E^*)$ as alternating multilinear forms on E is a well-defined and even legitimate thing to do. See Chapter 15 of the book [Bamberg and Sternberg, 1988a] for a good presentation on this topic, and Chapter 12 in [Lee, 2003] (Chapter 14 in the 2012 edition) for a detailed, mathematically oriented one (beware of the notations that are different than here!).

Exercise A.9. Show that for every alternating tensors $\Omega \in \wedge^m(E)$, then its evaluation on identical vectors is zero:

$$\Omega(x_1, \dots, x, \dots, x, \dots, x_m) = 0$$

From this result, deduce that for any choice of vectors x_1, \dots, x_m , if one of the x_i is a linear combination of the others, then $\Omega(x_1, \dots, x_m) = 0$.

A.2 Scalar product and Hodge star operator

Now suppose that E is additionally equipped with a *pseudo-Riemannian metric*, that is to say:

Definition A.10. A *pseudo-Riemannian metric* on a vector space E is a map g from $E \times E$ to \mathbb{R} which is:

1. *bilinear*, e.g. for the first term $g(\lambda x + \mu y, z) = \lambda g(x, z) + \mu g(y, z)$ for every $\lambda, \mu \in \mathbb{R}$ and $x, y, z \in E$ (and the same occurs for the second term)
2. *symmetric*, i.e. $g(x, y) = g(y, x)$ for every $x, y \in E$
3. *non-degenerate*, i.e. if $g(x, y) = 0$ for every $y \in E$ then $x = 0$

All three items are independent of the chosen basis of E . Given the definition of the symmetric algebra, one can see the metric as a bilinear map $g : E \odot E \rightarrow \mathbb{R}$. One can always define a metric on a vector space since one can check that, given a basis e_1, \dots, e_n , it is sufficient to define g from its action on the basis vectors e_i by:

$$g(e_i, e_i) = 1 \quad \text{and} \quad g(e_i, e_j) = 0 \quad \text{when } i \neq j, \quad (\text{A.20})$$

and then to formally extend it to all of E by assuming it is bilinear. Notice however that there exist alternative choices of scalar product that do not satisfy Equation (A.20), and more generally one writes³⁷:

$$g_{ij} = g(e_i, e_j)$$

The metric can then be represented, in a given basis e_1, \dots, e_n , as an $n \times n$ symmetric matrix G , whose components we write g_{ij} . Since the metric is symmetric, i.e. $g_{ij} = g_{ji}$, then there are only $\frac{n(n+1)}{2}$ independent coefficients in the matrix (the diagonal and the upper triangular part, or the diagonal and the lower triangular part). Being symmetric, the matrix can be diagonalized: the number p of positive eigenvalues determines what is called the *signature* of the metric, which is denoted by (p, q) , the number q being the number of negative eigenvalues. Notice that another convention uses the reverse order (q, p) . Sylvester's law of inertia ensures that the signature of the metric tensor g is invariant under any change of basis. There are no null eigenvalue because the metric is non-degenerate. In particular, the matrix G is invertible. Using these data, one can write the metric explicitly, as a bilinear symmetric form on $E \times E$:

$$g = g_{ij} e^i \odot e^j \in S^2(E^*)$$

where we have used the Einstein summation convention. Using the characterization of symmetric multilinear forms as elements of $T(E^*)$ (see Equation (A.15)), an explicit computation then shows that:

$$g(e_k, e_l) = g_{ij} e^i \odot e^j(e_k, e_l) = \frac{1}{2} g_{ij} (e^i \otimes e^j + e^j \otimes e^i)(e_k \otimes e_l) = \frac{1}{2} g_{ij} (\delta_k^i \delta_l^j + \delta_l^i \delta_k^j) = g_{kl}$$

because $g_{kl} = g_{lk}$. The result is precisely what we should expect.

Remark A.11. From now on, given a metric g , when we say *orthonormal basis*, we think of a basis e_1, \dots, e_n satisfying:

$$\begin{aligned} g(e_i, e_i) &= 1 && \text{for every } 1 \leq i \leq p \\ g(e_i, e_i) &= -1 && \text{for every } p+1 \leq i \leq n \\ g(e_j, e_k) &= 0 && \text{for every } j \neq k \end{aligned}$$

where p is the number of negative eigenvalues of g . In other words, we put the eigenvectors with positive eigenvalues (normed to 1) first in order. This is somewhat consistent with some conventions in Minkowski space whose metric's signature we set to $(3, 1) = (-+++)$: we often consider the time like coordinate to be either the fourth coordinate or the zeroth one. In any case, the first, second and third coordinates are space-like, and correspond to the positive eigenvalue $+1$. Obviously for a pseudo-Riemannian metric there is not negative eigenvalue, as for an Euclidean metric.

³⁷This notation is used in general relativity: space-time is a four dimensional manifold and the metric $g_{\mu\nu}$ is a notation for $g(\partial_\mu, \partial_\nu)$.

The metric induces a morphism of vector space $\tilde{g} : E \longrightarrow E^*$:

$$\begin{aligned} \tilde{g} : E &\longrightarrow E^* \\ x &\longmapsto g(x, \cdot) : y \longmapsto g(x, y) \end{aligned}$$

where we have used the Einstein summation convention. In particular, this definition tells us that acting on a basis vector with \tilde{g} reads:

$$\tilde{g}(e_i) = g_{ij} e^j$$


Moreover, the nondegeneracy of the metric is equivalent to the injectivity of this map, and hence, of its bijectivity (because E is finite dimensional).

Exercise A.12. The proof is left as an exercise.

The inverse map is denoted $\tilde{g}^{-1} : E^* \longrightarrow E$ and it can be used to define a metric g^{-1} on E^* induced from g , which is such that the maps \tilde{g}^{-1} and \tilde{g} are isometries³⁸. Following the same lines of argument as above, the scalar product g^{-1} gives rise to an isometry $\widetilde{g^{-1}} : E^* \longrightarrow E$, which actually is such that $\widetilde{g^{-1}} = \tilde{g}^{-1}$. That is to say, we shall have:

$$g^{-1}(\tilde{g}(x), \tilde{g}(y)) = g(x, y) \quad \text{and} \quad g(\tilde{g}^{-1}(\varphi), \tilde{g}^{-1}(\chi)) = g^{-1}(\varphi, \chi)$$

for every $x, y \in E$ and $\varphi, \chi \in E^*$. This is equivalent to the commutativity of the following diagram:



There is an $n \times n$ matrix H associated to the metric g^{-1} and the basis e^1, \dots, e^n . We adopt the convention that its components are written g^{ij} , with exponents, so that:

$$g^{-1}(e^i, e^j) = g^{ij}$$

The metric g^{-1} then corresponds to a bilinear symmetric operator on E^* expressed as:

$$g^{-1} = g^{ij} e_i \odot e_j \in S^2(E)$$

One can show that the matrix H is the inverse of the matrix G . This implies that the signature of the metric g^{-1} is the same as the one of g .

Exercise A.13. Using the symmetry of G and the fact that it is invertible, prove that $H = G^{-1}$.

³⁸This is a particular form of the Riesz representation theorem in mathematics, which actually applies to infinite dimensional Hilbert spaces.

The maps \tilde{g} and \tilde{g}^{-1} satisfy:

$$\tilde{g}(e_i) = g_{ij} e^j \in E^* \quad \text{and} \quad \tilde{g}^{-1}(e^k) = g^{kl} e_l \in E$$

where the Einstein summation convention has been used. These equations explain what people mean by saying that the metric raises and lowers the indices. For example, take a vector $x = x^i e_i \in E$, where the $x^i \in \mathbb{R}$ are the coordinates of x with respect to the basis $(e_i)_{1 \leq i \leq n}$. In physics in general one is only interested in the coordinates x^i , then lowering the indices i amounts to applying \tilde{g} to x :

$$\tilde{g}(x) = x^i g_{ij} e^j = x_j e^j$$

where we have defined $x_j := g_{ij} x^i$, which in the present context are thus the coordinates of $\tilde{g}(x)$ with respect to the basis $(e^i)_{1 \leq i \leq n}$ of E^* . Sometimes, the maps \tilde{g} and \tilde{g}^{-1} are called musical isomorphisms, and are denoted \flat (*flat*) and \sharp (*sharp*), respectively:

$$\flat : E \longrightarrow E^* \quad \text{and} \quad \sharp : E^* \longrightarrow E$$

They are inverse to one another. This notation is useful because it lightens the notation, by writing x^\flat instead of $\tilde{g}(x)$, and φ^\sharp instead of $\tilde{g}^{-1}(\varphi)$. Then, while $x = x^i e_i$, we have $x^\flat = x_j e^j$, with $x_j = g_{jk} x^k$. That is why we say that \flat lowers the indices (of the coordinates!) while the map \sharp raises them.

In particular, given a tensor $A_{i_1 \dots i_k}{}^{j_1 \dots j_l}$ one can raise and lower the indices using the musical isomorphisms, for example:

$$A_{i_1 \dots i_{r-1}}{}^{j_0}{}_{i_{r+1} \dots i_k}{}^{j_1 \dots j_l} \equiv g^{j_0 i_r} A_{i_1 \dots i_k}{}^{j_1 \dots j_l} \quad (\text{A.21})$$

where the Einstein summation convention has been used. This has the following mathematical meaning: $A = A_{i_1 \dots i_k}{}^{j_1 \dots j_l} e^{i_1} \otimes \dots \otimes e^{i_k} \otimes e_{j_1} \otimes \dots \otimes e_{j_l}$ is a mixed tensor belonging to $T^k(E^*) \otimes T^l(E)$. The left-hand side of Equation (A.21) is precisely the tensor obtained when one has applied $\sharp = \tilde{g}^{-1}$ on the r -th covariant leg of A . In other words (with Einstein summation convention):

$$\begin{aligned} A_{i_1 \dots i_{r-1}}{}^{j_0}{}_{i_{r+1} \dots i_k}{}^{j_1 \dots j_l} e^{i_1} \otimes \dots \otimes e^{i_{r-1}} \otimes e_{j_0} \otimes e^{i_{r+1}} \otimes \dots \otimes e^{i_k} \otimes e_{j_1} \otimes \dots \otimes e_{j_l} \\ = \left(\underbrace{\text{id}_{E^*} \otimes \dots \otimes \text{id}_{E^*}}_{r-1 \text{ terms}} \otimes \sharp \otimes \underbrace{\text{id}_{E^*} \otimes \dots \otimes \text{id}_{E^*}}_{k-r \text{ terms}} \otimes \underbrace{\text{id}_E \otimes \dots \otimes \text{id}_E}_l \right) (A) \end{aligned}$$

is an element of $T^{r-1}(E^*) \otimes E \otimes T^{k-r}(E^*) \otimes T^l(E)$.

The metric g can be extended to the exterior algebra $\bigwedge^\bullet(E)$ by using the Gram determinant. For every $1 \leq m \leq n$, we will define it first on *decomposable* multivectors, i.e. those elements of $\bigwedge^m(E)$ that are of the form $x_1 \wedge \dots \wedge x_m$ for $x_1, \dots, x_m \in E$, and then extend it to all of $\bigwedge^m(E)$ by linearity in each argument. More precisely, let $\alpha, \beta \in \bigwedge^m(E)$ be two decomposable multivectors, so that they can be written as $\alpha = x_1 \wedge \dots \wedge x_m$ and $\beta = y_1 \wedge \dots \wedge y_m$. Then we define the scalar product of α and β as:

$$\langle \alpha, \beta \rangle = \det \begin{pmatrix} g(x_1, y_1) & g(x_1, y_2) & \dots & g(x_1, y_{m-1}) & g(x_1, y_m) \\ g(x_2, y_1) & & & & g(x_2, y_m) \\ \dots & & & & \dots \\ g(x_{m-1}, y_1) & & & & g(x_{m-1}, y_m) \\ g(x_m, y_1) & g(x_m, y_2) & \dots & g(x_m, y_{m-1}) & g(x_m, y_m) \end{pmatrix}$$

The determinant on the right hand side resembles what is called the Gram determinant.

Exercise A.14. Prove that the so-called scalar product $\langle \alpha, \beta \rangle$ is symmetric and bilinear.

Exercise A.15. Prove that if e_1, \dots, e_n is an orthonormal basis of E (see Remark A.11 for a definition), the scalar product satisfies:

$$\langle e_{i_1} \wedge \dots \wedge e_{i_m}, e_{j_1} \wedge \dots \wedge e_{j_m} \rangle = g_{i_1 j_1} g_{i_2 j_2} \dots g_{i_m j_m} \quad (\text{A.22})$$

Then, since any m -multivector can be written as the linear sum of decomposable m -multivectors – such as the basis $(e_{i_1} \wedge \dots \wedge e_{i_m})_{1 \leq i_1 < \dots < i_m \leq n}$ of $\Lambda^m(E)$ – one can extend the inner product to the whole of $\Lambda^m(E)$ by enforcing linearity on each argument. For example, let $\alpha = \sum_i \alpha_i$ and $\beta_j = \sum_j \beta_j$ be two m -multivectors written as linear combinations of the decomposable multivectors α_i and β_j ; then we set:

$$\langle \alpha, \beta \rangle = \sum_{i,j} \langle \alpha_i, \beta_j \rangle$$

We apply the same idea at every level $1 \leq m \leq n-1$ (for $m=0$ and $m=n$ the space $\Lambda^m(E)$ is one-dimensional) so that the scalar product is defined on the entirety of the exterior algebra $\Lambda^\bullet(E)$. It can be shown that the left hand side does not depend on the decomposition of α and β in terms of decomposable multivectors (see a proof in the Appendix of Chapter 18 of [Bamberg and Sternberg, 1988a]). This definition also work on $\Lambda^\bullet(E^*)$ as well, when one takes g^{-1} instead of g .

Proposition A.16. *The so-called scalar product $\langle \alpha, \beta \rangle$ is non-degenerate so it indeed bears well its name.*

Proof. One can suppose that $1 \leq m \leq n-1$ and that the basis of E is orthonormal with respect to the metric g . Let α be an m -multivector such that:

$$\langle \alpha, \beta \rangle = 0 \quad \text{for every } \beta \in \bigwedge^m(E) \quad (\text{A.23})$$

The element α admits the following decomposition on the basis of $\Lambda^m(E)$:

$$\alpha = \alpha^{i_1 \dots i_m} e_{i_1} \wedge \dots \wedge e_{i_m}$$

where the Einstein summation convention has been used, and where we assumed $i_1 < \dots < i_m$. Apply Equation (A.23) to $e_{i_1} \wedge \dots \wedge e_{i_m}$, for some chosen $i_1 < \dots < i_m$. Then by Equation (A.22), one obtains $0 = \langle \alpha, e_{i_1} \wedge \dots \wedge e_{i_m} \rangle = \alpha^{i_1 \dots i_m}$. Doing this for every basis vector of $\Lambda^m(E)$, one proves that $\alpha = 0$. \square

Then, since $\dim(\Lambda^m(E)) = \dim(\Lambda^{n-m}(E)) = \binom{n}{m}$, one can use the inner product on the exterior algebra to identify $\Lambda^m(E)$ and $\Lambda^{n-m}(E)$, via what is called the *Hodge star operator*. Let $(e_i)_{1 \leq i \leq n}$ be a basis of E and denote by

$$\omega = \frac{1}{\sqrt{|\det(G)|}} e_1 \wedge \dots \wedge e_n \quad (\text{A.24})$$

the standard volume element of E , which also generates the one-dimensional space $\Lambda^n(E)$. The normalization is such that $\langle \omega, \omega \rangle = (-1)^q$, which depends on the number q of negative eigenvalues of the metric.

Exercise A.17. The proof is left to the reader.

Then, the Hodge star operator is a linear map³⁹ $\star : \Lambda^\bullet(E) \longrightarrow \Lambda^{n-\bullet}(E)$ which is defined on $\Lambda^m(E)$ (for every $1 \leq m \leq n$) by the following identity:

$$\alpha \wedge (\star \beta) = \langle \alpha, \beta \rangle \omega \quad (\text{A.25})$$

for every two m -multivectors $\alpha, \beta \in \Lambda^m(E)$ ⁴⁰. Notice that a choice of orientation for E has a direct consequence on the sign of the volume form, and thus on the definition of the Hodge star operator. This indirect definition can be made more explicit by looking at the effect of \star on a basis of $\Lambda^m(E)$. For any ordered subset $J = \{j_1, \dots, j_m\}$ of $\{1, \dots, n\}$ (i.e. such that $1 \leq j_1 < \dots < j_m \leq n$), let us call $J^c = \{1, \dots, n\} - J$ the ordered set corresponding to the remaining integer that do not belong to J . We denote the $n - m$ elements of J^c as j_{m+1}, \dots, j_n ; they are such that $j_k < j_l$ for $m < k < l$. Then, we denote:

$$e_J = e_{j_1} \wedge \dots \wedge e_{j_m} \quad \text{and} \quad e_{J^c} = e_{j_{m+1}} \wedge \dots \wedge e_{j_n}$$

Moreover, let σ_J be the permutation of n elements that sends the ordered set $\{1, \dots, n\}$ to $\{j_1, \dots, j_m, j_{m+1}, \dots, j_n\}$, i.e. it is such that $\sigma_J(k) = j_k$. Thus, under the action of σ_J , the n -form $e_1 \wedge \dots \wedge e_n$ becomes $e_J \wedge e_{J^c}$. Then, by fixing an ordered set $I = \{i_1, \dots, i_m\}$ of m elements $1 \leq i_1 < \dots < i_m \leq n$ and by computing every term of the form $e_J \wedge \star(e_I)$ for every ordered set $J = \{j_1, \dots, j_m\} \subset \{1, \dots, n\}$, one obtains the following formula:

$$\star(e_I) = \sum_{\substack{J \subset \{1, \dots, n\} \\ J \text{ ordered}}} (-1)^{\sigma_J} \frac{\langle e_J, e_I \rangle}{\sqrt{|\det(G)|}} e_{J^c} \quad (\text{A.26})$$

where $\langle e_J, e_I \rangle$ is a notation for $\langle e_{j_1} \wedge \dots \wedge e_{j_m}, e_{i_1} \wedge \dots \wedge e_{i_m} \rangle$, which is a minor of the Gram matrix.

Exercise A.18. Using Equation (A.25), show that in \mathbb{R}^3 with standard basis e_1, e_2, e_3 and with a metric $g = g_{ij} e^i \odot e^j$:

$$\begin{aligned} \star e_i &= \frac{1}{\sqrt{|\det(G)|}} \left(g_{1i} e_2 \wedge e_3 - g_{2i} e_1 \wedge e_3 + g_{3i} e_1 \wedge e_2 \right) \\ \star(e_i \wedge e_j) &= \frac{1}{\sqrt{|\det(G)|}} \left(\det \begin{pmatrix} g_{2i} & g_{2j} \\ g_{3i} & g_{3j} \end{pmatrix} e_1 - \det \begin{pmatrix} g_{1i} & g_{1j} \\ g_{3i} & g_{3j} \end{pmatrix} e_2 + \det \begin{pmatrix} g_{1i} & g_{1j} \\ g_{2i} & g_{2j} \end{pmatrix} e_3 \right) \end{aligned}$$

and check that these formulas are indeed those corresponding to Equation (A.26).

Finding an orthonormal basis with respect to the metric (see Remark A.11 for the definition) is equivalent to diagonalizing the associated matrix G , and rescale the diagonal values so that they become either 1 or -1 . The determinant is then the product of the diagonal values g_{ii} , and its absolute value is 1. For any ordered m -index $I = \{i_1, \dots, i_m\} \subset \{1, \dots, n\}$, one sets $\lambda_I = g_{i_1 i_1} g_{i_2 i_2} \dots g_{i_m i_m}$ the product of the m diagonal values. Then, assuming the basis e_1, \dots, e_n is orthonormal, one has:

$$\star(e_I) = (-1)^{\sigma_I} \lambda_I e_{I^c} = \pm e_{I^c} \quad (\text{A.27})$$

For an explicit formulation using coordinates, see the [nCatLab](#).

Let us illustrate Equation (A.26) on several examples. First, the following two identities:

$$\star(1_{\mathbb{R}}) = \omega \quad \text{and} \quad \star \omega = (-1)^q 1_{\mathbb{R}}$$

³⁹This notation means that for every $1 \leq m \leq n$ the Hodge star operator sends $\Lambda^m(E)$ to $\Lambda^{n-m}(E)$.

⁴⁰Notice that in general the Hodge star is usually defined on covectors. In that case the one shall use e^i instead of e_i , and the normalization factor $\sqrt{|\det(G)|}$ instead of $\sqrt{|\det(G^{-1})|}$.

are valid in every case, where $1_{\mathbb{R}}$ is the generator of $\bigwedge^0(E) = \mathbb{R}$, i.e. $1_{\mathbb{R}} = 1$. When $E = \mathbb{R}^2$ with the standard Euclidean metric and standard orientation, $\bigwedge^1(E)$ is two-dimensional as well and can be identified with E , so that the Hodge star operator can be seen as an endomorphism of E and coincides with a rotation by $\frac{\pi}{2}$. This can be checked on any orthonormal basis $\{e_1, e_2\}$ of $E = \bigwedge^1(E)$, since we have:

$$\star e_1 = e_2 \quad \text{and} \quad \star e_2 = -e_1$$

When $E = \mathbb{R}^3$ with the standard Euclidean metric and standard orientation, the two spaces $\bigwedge^1(E) \simeq E$ and $\bigwedge^2(E)$ are both three-dimensional and the Hodge star operator draws a relationship between the wedge product and the cross product:

$$\star(x \times y) = x \wedge y \quad \text{and} \quad \star(x \wedge y) = x \times y$$

Exercise A.19. Can you compute the effect of the Hodge star operator on Minkowski space? The Minkowski metric has signature $(3, 1)$.

Last but not least: the Hodge star operator is not an involution of the exterior algebra, but almost:

$$\star \star \alpha = (-1)^{m(n-m)+q} \alpha \tag{A.28}$$

for any m -multivector α , and where q is the number of negative eigenvalues in the signature (p, q) of the metric. This implies that, the inverse to the hodge star operator $\star : \bigwedge^m(E) \longrightarrow \bigwedge^{n-m}(E)$, is the operator $\star^{-1} : \bigwedge^{n-m}(E) \longrightarrow \bigwedge^m(E)$ defined by:

$$\star^{-1} = (-1)^{m(n-m)+q} \star \tag{A.29}$$

The final identity worth noticing is:

$$\langle \star \alpha, \star \beta \rangle = (-1)^q \langle \alpha, \beta \rangle \tag{A.30}$$

This equation proves that the Hodge star operator is almost an isometry of the exterior algebra, up to a sign.

Exercise A.20. Using Exercise A.18, prove that in \mathbb{R}^3 with the given metric g ,

$$\star \star e_i = (-1)^q e_i \quad \text{and} \quad \star \star (e_i \wedge e_j) = (-1)^q e_i \wedge e_j$$

Exercise A.21. Using Equation (A.27) and the fact that $(-1)^{\sigma_i^c} = (-1)^{\sigma_I+m(n-m)}$, prove Equation (A.28), that in turns implies Equation (A.30).

Finally notice that usually the Hodge star operator is defined on the exterior algebra of covectors $\bigwedge^\bullet(E^*)$. Pay heed to the differences that this implies: in particular one should use G^{-1} instead of G , and use exponents (resp. indices) in place of indices (resp. exponents).

B Supplementary and facultative material

B.1 Pseudo-Riemannian manifolds and Laplace-de Rham operator

In Section 2.5 we used the fact that smooth manifolds admit a tangent bundle, that associates a tangent space to every point, to define an orientation on manifolds. We can use the same strategy to define pseudo-Riemannian metrics on smooth manifolds. First we define a pointwise metric, and we require it to vary smoothly over the manifold.

Definition B.1. *Let M be an n -dimensional smooth manifold and let $x \in M$. A metric tensor on M is a smooth section g of the vector bundle $S^2(T^*M)$ that restricts at every point $x \in M$ to a pseudo-Riemannian metric $g_x : T_x M \times T_x M \rightarrow \mathbb{R}$. We call a smooth manifold equipped with a metric tensor a pseudo-Riemannian manifold; it is said Riemannian when the metric tensor is positive definite at every point.*

We know that the tangent space at a point is the best linear approximation of the manifold at that point. The metric tensor at this point is then fed by tangent vectors. However, since g is a smooth tensor, it can be fed by vector fields, and the result defines a smooth function (that is actually a way of characterizing smoothness of g):

$$g(X, Y) = g(Y, X) \in \mathcal{C}^\infty(M) \quad \text{for every } X, Y \in \mathfrak{X}(M)$$

The smoothness of the metric tensor g is characterized by the fact that the map $x \mapsto g_x(X_x, Y_x)$ is a smooth map, for every two smooth vector fields X and Y . The metric tensor g can be locally decomposed in a coordinate cotangent basis dx^1, \dots, dx^n over a coordinate chart U as:

$$g = g_{ij} dx^i \odot dx^j$$

where $g_{ij} \in \mathcal{C}^\infty(U)$ are smooth functions. So, in particular, with respect to the coordinate tangent frame:

$$g_{ij} = g(\partial_i, \partial_j)$$

Obviously since the metric tensor varies smoothly, its pointwise signature is constant over U (and more generally, over M), and is determined by the eigenvalues of the matrix-valued smooth function $G \in \mathcal{C}^\infty(U, \mathcal{M}_n(\mathbb{R}))$.

Remark B.2. Let D be a regular smooth distribution on M . Then, assume that we have a smoothly varying metric g_x defined on each subspace D_x and that is smoothly varying, in the sense that for every two smooth sections $X, Y \in \Gamma(D)$, the map $x \mapsto g_x(X_x, Y_x)$ is a smooth map over M . It is as if the metric was defined in the directions defined by D . We call this ‘metric tensor’ a *sub-Riemannian metric*, and the smooth manifold M a sub-Riemannian manifold. This metric defines a distance function on the manifold by integrating it over any horizontal path joining the two points:

$$d_\gamma(x, y) = \int_0^1 \sqrt{g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt$$

The *Carnot-Carathéodory distance* is then the infimum of all such distance, over all the horizontal paths:

$$d_{CC}(x, y) = \inf_{\text{horizontal } \gamma} \{d_\gamma(x, y)\}$$

This distance is very useful in sub-Riemannian geometry⁴¹. For example, Chow-Rashevskii theorem 2.72 can be restated as the following: “the topology induced by the Carnot-Carathéodory metric is equivalent to the intrinsic (locally Euclidean) topology of the manifold”.

⁴¹See these [introductory notes](#) to sub-Riemannian geometry.

A pseudo-Riemannian manifold that is additionally an oriented manifold has a distinguished volume form, that we now present. By Corollary 2.86, the fact that M is oriented means that there exist a nowhere vanishing globally defined volume form ω that is positively orientated at every point. The following argument explain that we can chose ω in a certain, adapted form. Since a metric tensor g induces a pointwise metric g^{-1} on the cotangent bundle that varies smoothly, then the following map $\sqrt{|\det(G^{-1})|} = \frac{1}{\sqrt{|\det(G)|}}$ is a smooth function which is well defined and nowhere vanishing. We also sometimes write $\sqrt{|g|}$ instead of $\sqrt{|\det(G)|}$. Since the metric associated to the coordinate cotangent frame dx^1, \dots, dx^n is g^{-1} , at the price of multiplying ω by a nowhere vanishing positive smooth function, we can always have:

$$\omega = \sqrt{|g|} dx^1 \wedge \dots \wedge dx^n \quad (\text{B.1})$$

This formula is the counterpart of Equation (A.24) in the context of smooth manifolds, where the exterior algebra $\Lambda^\bullet(E)$ is the vector bundle $\Lambda^\bullet T^*M$. Equation (B.1) is the local form of the standard volume element on a pseudo-Riemannian oriented manifold (M, g) .

Following the discussion in Section A.2, the pseudo-Riemannian metric g on M induces a pairing on the fibers of the exterior algebra of the cotangent bundle:

$$\begin{aligned} \langle \cdot, \cdot \rangle : \Omega^m(M) \times \Omega^m(M) &\longrightarrow \mathcal{C}^\infty(\mathbb{R}) \\ (\eta, \mu) &\longmapsto \langle \eta, \mu \rangle : x \longrightarrow \langle \eta_x, \mu_x \rangle \end{aligned}$$

for every $0 \leq m \leq n$. It is fiberwisely non-degenerate, but one needs to integrate the function on the right-hand side in order to define an inner product (\cdot, \cdot) on differential forms, via the following formula:

$$(\eta, \mu) = \int_M \langle \eta, \mu \rangle \omega \quad (\text{B.2})$$

for every $\eta, \mu \in \Omega^m(M)$, and every $0 \leq m \leq n$, and where ω is the distinguished volume form defined in Equation (B.1). This product may be divergent if the support of one of the arguments does not have compact support. It defines a L^2 norm on those differential forms η that are such that $(\eta, \eta) < +\infty$ (in particular compactly supported differential forms satisfy this condition).

The volume form defined in Equation (B.1) and the fiberwise inner product $\langle \cdot, \cdot \rangle$ also allow to define a Hodge star operator $\star : \Omega^m(M) \longrightarrow \Omega^{n-m}(M)$, as in Equation (A.25). One can then define a $\mathcal{C}^\infty(M)$ -linear operator $\delta : \Omega^\bullet(M) \longrightarrow \Omega^{\bullet-1}(M)$ that is ‘dual’ in some sense to the de Rham differential. On m -forms, it is defined as:

$$\delta = -(-1)^{n(m-1)+q} \star d \star$$

and sends m -forms to $m-1$ -forms. By construction, on 0-forms it is zero. Using the definition of the inverse star operator (see Equation (A.29)) $\star^{-1} : \Lambda^{n-m-1}(M) \longrightarrow \Lambda^{m-1}(M)$, one deduces that $\delta : \Omega^m(M) \longrightarrow \Omega^{m-1}(M)$ can also be written as:

$$\delta = (-1)^m \star^{-1} d \star \quad (\text{B.3})$$

Then, for every differential m -form η , and any differential $m+1$ -form μ , one has the following identity:

$$\star(\langle d\eta, \mu \rangle - \langle \eta, \delta\mu \rangle) = d(\eta \wedge \star\mu) \quad (\text{B.4})$$

The right-hand side is a n -form, that is why we used a star operator on the scalar in parenthesis on the left-hand side so that it becomes a n -form as well. Thus Equation (B.4) implies that δ is the adjoint of the de Rham differential, with respect to the inner product on differential forms:

$$(\eta, \delta\mu) = (d\eta, \mu)$$

for every $\eta \in \Omega^m(M)$ and $\mu \in \Omega^{m+1}(M)$, where $0 \leq m \leq n-1$.

Exercise B.3. Prove Equation (B.3).

Exercise B.4. Prove that the identity $d^2 = 0$ implies that $\delta^2 = 0$.

Definition B.5. We call the $C^\infty(M)$ -linear operator $\delta : \Omega^\bullet(M) \longrightarrow \Omega^{\bullet-1}(M)$ the codifferential. We define the Laplace-de Rham operator as the $C^\infty(M)$ -linear operator $\Delta_{dR} : \Omega^\bullet(M) \longrightarrow \Omega^\bullet(M)$ such that:

$$\Delta_{dR} = d \circ \delta + \delta \circ d$$

The first term of the Laplace-de Rham operator vanishes on smooth functions, i.e. 0-forms, so that we obtain minus the *Laplace-Beltrami operator*⁴²:

$$\Delta_{dR}(f) = -(-1)^q \star d \star df = -\Delta(f)$$

The difference in sign is a convention and descends from the additional sign in δ . The Laplace-de Rham operator is defined to be positive definite, whereas the Laplace-Beltrami operator is usually taken to be negative definite. Since the Laplace-Beltrami operator reads, in coordinates:

$$\Delta(f) = \frac{1}{\sqrt{|g|}} \partial_i \left(\sqrt{|g|} g^{ij} \partial_j(f) \right) \quad (\text{B.5})$$

we deduce that, in Minkowski space-time with signature $(3, 1)$ or, in physics notation, $(-, +, +, +)$, the Laplace-de Rham operator is the d'Alembertian:

$$\Delta_{dR} = \square = \frac{\partial^2}{\partial t^2} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2} \quad (\text{B.6})$$

Be aware however that, under the convention that the signature of the Minkowski metric is $(1, 3) = (+, -, -, -)$, the right-hand side can be written $\partial^\mu \partial_\mu$. On the contrary, with our convention of signature and using Equation (B.5), the right-hand side of Equation (B.6) reads $-\partial^\mu \partial_\mu$.

The codifferential δ and the Laplace-de Rham operator Δ_{dR} allow to characterize more precisely differential forms and de Rham cohomology. A differential form η that is such that $\delta\eta = 0$ is called *co-closed*, while if there is another differential form μ such that $\eta = \delta\mu$, we say that η is *co-exact*. Differential forms that lie in the kernel of the Laplacian, i.e. those η such that $\Delta_{dR}(\eta) = 0$, are called *harmonic*. We denote by $\mathcal{H}^m(M)$ the space of harmonic differential m -forms, for $0 \leq m \leq n$. Exact, co-exact and harmonic differential forms provide a nice decomposition of the space of differential forms:

Theorem B.6. Hodge decomposition *Let M be a compact Riemannian manifold, then for every $0 \leq m \leq n$, we have the following decomposition:*

$$\Omega^m(M) = d(\Omega^{m-1}(M)) \oplus \delta(\Omega^{m+1}(M)) \oplus \mathcal{H}^m(M)$$

This direct sum is orthogonal with respect to the inner product defined in Equation (B.2).

This decomposition is very useful to find a distinguished representative of de Rham cohomology classes, because the following corollary proves that each cohomology class has a unique harmonic representative:

Corollary B.7. *Let M be a compact Riemannian manifold, then for every $0 \leq m \leq n$ we have an isomorphism:*

$$H_{dR}^m(M) \simeq \mathcal{H}^m(M)$$

⁴²We cannot justify yet that the Laplace-Beltrami operator $\text{div} \circ \overrightarrow{\text{grad}}$ corresponds $(-1)^q \star d \star d$ but it will be shown later in the course.

Proof. Both proofs of the theorem and of the corollary can be found in Chapter 6 of [Warner, 1983]. \square

We conclude this section by the following beautiful remark: integrals and the Hodge star operator allow to write actions in a rather nice way. For example, integrating the Lagrangian density of Maxwell's electromagnetism $F_{\mu\nu}F^{\mu\nu}$ over a pseudo-Riemannian n -dimensional manifold M can be synthesized as (physical notation is on the left):

$$\mathcal{S}_M = \frac{1}{4} \int_M F_{\mu\nu}F^{\mu\nu} \sqrt{|g|} d^n x = \frac{1}{2} \int_M F \wedge \star F$$

One can also write Einstein-Hilbert action (without cosmological constant) as:

$$\mathcal{S}_{EH} = \int_M R \sqrt{|g|} d^n x = \int_M \star R$$

where R is the Ricci scalar. Then, more generally, integrating a Lagrangian density \mathcal{L} over an oriented pseudo-Riemannian smooth manifold M provides the following action:

$$\mathcal{S} = \int_M \star \mathcal{L}$$

Obviously in both cases there is a possible problem of convergence of the integral but we may either work only locally (physical quantities in classical physics do not have non-local properties) so that we can assume that the Lagrangian densities are compactly supported, or we can accept that the integral is not properly defined although while we admit only compactly supported variations of the fields (e.g. δA would be the compactly supported 'variation' of a connection 1-form A), then the induced variation $\delta \mathcal{S}$ would be well-defined (see Section II.4 of [Baez and Muniain, 1994]). This opens the possibility to work on physical theories from a geometric point of views. Gauge theories are precisely theories which benefit from such an approach.

B.2 The Poisson-sigma model

In Physics, a *sigma model* is a way of encoding an action functional from a smooth map sometimes denoted $\sigma : \Sigma \rightarrow M$, where Σ and M are smooth manifolds called respectively the *source* and *target* manifolds. Their dimension and the possibly additional structures (such as a pseudo-Riemannian metric or a Poisson structure on M) that these manifolds possess characterize the so-called sigma model. Sigma models are useful for the following reason: the dynamical fields of the physical theory correspond to the composite functions $\sigma^i = x^i \circ \sigma$ on the *target space*. For example the relativistic particle can be seen as a sigma model $X : \mathbb{R} \rightarrow \mathbb{M}^4$ (where \mathbb{M}^4 is Minkowski space), given by the action:

$$S \propto \int_{\mathbb{R}} \eta_{ij}(X(\tau)) \dot{X}^i(\tau) \dot{X}^j(\tau) d\tau$$

The trajectory of the particle in space time is parametrized by the proper time τ and is called the *world-line* of the particle. Notice that integration is made over the manifold Σ and not over M , because the independent variables are the coordinates over Σ .

Another example is the Nambu-Goto action for the bosonic relativistic open string is obtained from a sigma model $X : \Sigma \rightarrow M$, where Σ is a 2-dimensional smooth manifold (with boundaries) called a *world-sheet*, parametrized by a timelike coordinate τ and a spacelike coordinate σ , and M is a pseudo-Riemannian manifold representing spacetime. Then the Nambu-Goto action is:

$$S_{NG} \propto \int_{\Sigma} \sqrt{(g_{\mu\nu}(X) \dot{X}^\mu X'^\nu)^2 - \dot{X}^\mu \dot{X}_\mu X'^\nu X'_\nu} d\tau d\sigma$$

where $\dot{X} = \frac{\partial X}{\partial \tau}$ and $X' = \frac{\partial X}{\partial \sigma}$, and where $g_{\mu\nu}$ is the metric on M .

A gauge theory on a pseudo-Riemannian oriented manifold M may be seen as a particular kind of sigma model: it is characterized by a set of *gauge fields* corresponding to the component of a Lie-algebra valued one-form $A = A_\mu dx^\mu = A_\mu^a T_a \otimes dx^\mu \in \Omega^1(M, \mathfrak{g})$, where the T_a form a basis of \mathfrak{g} . The Yang-Mills action is then written as:

$$S_{YM} = \frac{1}{2\alpha} \int_M \text{tr}(F \wedge \star F) \quad (\text{B.7})$$

where the F is the field strength associated to A : $F^a = dA^a + \frac{1}{2}[A, A]^a$. Usually, \mathfrak{g} is a semi-simple matrix Lie algebra so that the trace is the usual trace on matrices, however in the more general case, one should think of the trace as symbolizing the Killing form κ on \mathfrak{g} ⁴³. The Lie bracket is that of \mathfrak{g} , while the differential form component of A is wedged. More precisely:

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + [A_\mu, A_\nu]^a \quad (\text{B.8})$$

Moreover, the notation $F \wedge \star F$ means that the wedge acts with respect to the forms, whereas the Lie algebra components of F is composed with that of $\star F$ (via the adjoint action, say). Much more details can be found in Chapter 3 of Part 2 of [Baez and Muniain, 1994]. It can be seen as a sigma model via the observation that the gauge field is a Lie algebroid morphism $A : TM \rightarrow \mathfrak{g}$. The source manifold is thus TM while the target manifold is \mathfrak{g} .

Exercise B.8. Show that, decomposing $F = dA + \frac{1}{2}[A, A]$ as $\frac{1}{2}F_{\mu\nu}^a T_a \otimes dx^\mu \wedge dx^\nu$, one indeed finds Equation (B.8).

The Yang-Mills action can be rewritten by introducing a $n-2$ differential form $X = X_\mu^a T_a \otimes dx^\mu$ taking values in \mathfrak{g} :

$$S_{YM} = \int_M \text{tr}(X \wedge F + \frac{\alpha}{2} X \wedge \star X) \quad (\text{B.9})$$

Then the Euler-Lagrange equation on X is $X = \frac{1}{\alpha} \star F$ (at least when M is a Lorentzian four-dimensional manifold, see Equation (A.28)), so that we retrieve the original action (B.7), upon replacing X by its value. Something particular occur when the manifold M is two dimensional (we will call it Σ), because in that case X is a function and $\star X = X\omega$ where $\omega = \sqrt{|g|} dx^1 \wedge dx^2$ is the normalized volume form on M , as defined in Equation (B.1). In the case where \mathfrak{g} is a finite dimensional semi-simple matrix Lie algebra, the 2-dimensional Yang-Mills action (B.9) becomes (up to some scalar factor):

$$\int_\Sigma \kappa_{ab} X^a (dA^b + \frac{1}{2}[A, A]^b) + \frac{\alpha}{2} \kappa_{ab} X^a X^b \omega$$

where $\kappa_{ab} = \text{tr}(\text{ad}_{T_a} \circ \text{ad}_{T_b})$ are the components of the Killing form on \mathfrak{g} . Since in this nice situation, the Killing form is a non-degenerate bilinear form on \mathfrak{g} , from now on we will use contracted indices instead. Upon integrating by part the term $X^a dA^a$ (assuming, e.g., that the source manifold Σ has no boundary), we obtain:

$$\int_\Sigma A^a \wedge dX_a + \frac{1}{2} X_a [A, A]^a + \frac{\alpha}{2} X_a X^a \omega \quad (\text{B.10})$$

Now, observe that \mathfrak{g} is the linear dual of the Poisson vector space \mathfrak{g}^* (see Example 3.4 for more details on linear Poisson structures). In other words, the smooth function X and the

⁴³For finite dimensional semi-simple matrix algebras such as \mathfrak{sl}_n , \mathfrak{so}_n , \mathfrak{su}_n , \mathfrak{sp}_{2n} , the Killing form $\kappa(u, v)$ – for any two elements $u, v \in \mathfrak{g}$ – is proportional to $\text{tr}(u \circ v)$, where u, v are in the latter case seen as endomorphisms of \mathbb{R}^n (or \mathbb{R}^{2n} for the symplectic algebra).

differential one-form A take values in $\mathfrak{g}^{**} \simeq \mathfrak{g}$ (because \mathfrak{g} is finite dimensional). This is true for any dimension of the source manifold Σ , but what is characteristic of the 2-dimensional case is that the expression $X_a X^a$ in the last term resembles the quadratic Casimir element of semi-simple Lie algebras (which usually form the kind of Lie algebras used in gauge theories). More precisely, for X^a a smooth function on M , the element $\sum_{a=1}^{\dim(\mathfrak{g})} X^a X^a T_a \odot T_a$ of the symmetric algebra of \mathfrak{g} can be seen as a polynomial function on \mathfrak{g}^* , which actually turns out to be a Casimir element in the sense of Poisson algebras. Thus, we have shown that the 2-dimensional Yang-Mills theory can be reformulated in terms of a sigma model involving a linear Poisson structure (that of \mathfrak{g}^*). A natural generalization is then to weaken that condition and allow this theory to be defined on any Poisson manifold:

Definition B.9. *The Poisson-sigma model is a sigma model defined by the following data:*

1. *the source Σ is a 2-dimensional oriented smooth manifold (possibly with boundary);*
2. *the target M is a finite dimensional Poisson manifold, with Poisson bivector π ;*
3. *the maps defining the model is a Lie algebroid morphism $(X, A) : T\Sigma \longrightarrow T^*M$;*

and by the following action functional:

$$S_{PSM}(X, A) = \int_{\Sigma} \langle A, X_* \rangle + \frac{1}{2} \langle A \wedge A, X^! \pi \rangle \quad (\text{B.11})$$

where $\langle \cdot, \cdot \rangle$ denotes the pairing between TM and T^*M , and where $C \in \mathcal{C}^\infty(M)$ is any Casimir function of π .

Let us explain each term in details. The pushforward $X_* : T\Sigma \longrightarrow X^!TM$ can be seen as a one form on Σ taking values in $\Gamma(X^!TM)$. In local coordinates, it can be written as $X_* = dX^i \frac{\partial}{\partial x^i}$ where d is the de Rham differential on Σ and where the x^i are coordinates on M . Then, since A takes values in $\Gamma(X^!T^*M)$, the pairing in the first term is indeed well defined, so that it becomes: $\langle A, X_* \rangle = A_i \wedge dX^i$. In the second term, the notation $X^! \pi$ symbolizes that we evaluate the Poisson bivector π on the image of X in M . In other words, $X^! \pi$ is a section of the pullback vector bundle $X^! \wedge^2 T^*M$. This is justified by the fact that the differential 2-form $A \wedge A$ takes values in $\Gamma(X^! \wedge^2 T^*M)$. Then the second term becomes in coordinates: $\frac{1}{2} \langle A \wedge A, X^! \pi \rangle = \frac{1}{2} \pi^{ij}(X) A_i \wedge A_j$ (because π contains $\frac{1}{2} \pi^{ij}$). Then, Equation (B.11), the action functional of the PSM, can be rewritten as:

$$S_{PSM}(X, A) = \int_{\Sigma} A_i \wedge dX^i + \frac{1}{2} \pi^{ij}(X) A_i \wedge A_j$$

It is usual to add an additional term in the Poisson-sigma model that plays the same role as $\frac{\alpha}{2} X_a X^a \omega$ in 2-dimensional Yang-Mills theory. Any choice of Casimir function $C \in \mathcal{C}^\infty(M)$ (relatively to the Poisson bivector π) can be added to the action functional, which then becomes:

$$S_{PSM}(X, A) = \int_{\Sigma} A_i \wedge dX^i + \frac{1}{2} \pi^{ij}(X) A_i \wedge A_j + \star(C(X))$$

As for the other terms, the Casimir function is evaluated on $\text{Im}(X) \subset M$. The constant $\frac{\alpha}{2}$ that was appearing in Yang-Mills action functional is not apparent in the above formula because it can be absorbed in the Casimir C . Obviously, if $M = \mathfrak{g}^*$ (where \mathfrak{g} is a finite dimensional semi-simple matrix Lie algebra, say), and if the Casimir function is the quadratic Casimir element of \mathfrak{g} , then the Poisson-sigma model with Casimir corresponds to the 2-dimensional Yang-Mills action functional (B.10), under the following considerations: 1. the map $X : \Sigma \longrightarrow \mathfrak{g}^*$ is considered to

take values in \mathfrak{g} by using the non-degenerate Killing form on \mathfrak{g} which allows to identify \mathfrak{g} and \mathfrak{g}^* ; 2. the fiber of $T^*\mathfrak{g}^*$ is identified with \mathfrak{g} so that the differential 1-form $A : T\Sigma \rightarrow T^*\mathfrak{g}^*$ is actually seen as taking values in \mathfrak{g} . This can be made explicit by realizing that A is actually a vector bundle morphism $T\Sigma \rightarrow X^!T^*M$ covering the identity map on Σ , then, evaluating the differential one-form A on a tangent vector at a point x gives an element of the fiber of $T_{X(x)}^*M$, i.e. an element of \mathfrak{g} , as required.

If $C(X) = 0$ then the Poisson-sigma model becomes a topological field theory, called a *BF-theory*. These are characterized by the following action functional:

$$S_{BF} = \int_{\Sigma} \text{tr}(B \wedge F)$$

where Σ is a n -dimensional oriented manifold, F is the field strength associated to the gauge potential A (taking values in some Lie algebra, say), while B is a \mathfrak{g} valued differential $n-2$ -form. The Euler-Lagrange equations of such topological field theories are:

$$F = 0 \quad \text{and} \quad d_A B = 0$$

where d_A is the covariant derivative associated to the connection A . The solutions of the equations are purely topological: B is a closed 2-form, while the field strength of A vanishes so A does not propagate. Under appropriate assumptions (e.g. Σ is compact without boundary), the Poisson-sigma model is a 2-dimensional BF-theory, since Equation (B.11) can be rewritten as:

$$S_{PSM} = \int_{\Sigma} X^i \wedge F_i$$

where summation on contracted indices is implicit. Thus, the Poisson-sigma model is a topological field theory that, under the addition of a Casimir function, can encode some physical model such as 2-dimensional Yang-Mills gauge theory.

Exercise B.10. Check that the action functional of the Poisson-sigma model is invariant under the following gauge transformations:

$$\delta_{(\epsilon, \lambda)} X^i = \lambda_j \pi^{ij} \quad \text{and} \quad \delta_{(\epsilon, \lambda)} A_i = d\lambda_i + \frac{\partial \pi^{kl}}{\partial x^i} A_k \epsilon_l$$

where $\epsilon_1, \dots, \epsilon_n$ are smooth functions on Σ and $\lambda = \lambda_i dx^i$ is a differential 1-form on Σ taking values in T^*M . They are obviously nonlinear generalizations of standard gauge transformations.

Another application of the Poisson-sigma model (and actually its original motivation) is to describe 2-dimensional gravity (without matter field). Let Σ be an oriented, 2-dimensional Lorentzian manifold, with metric g (of signature $(1, 1)$ then). Let us denote by x^0 and x^1 local coordinates on Σ . Recall that in two dimensions, the symmetries of the Riemann tensor impose that:

$$R_{\mu\nu\alpha\beta} = -\frac{R}{2} \epsilon_{\mu\nu} \epsilon_{\alpha\beta} \tag{B.12}$$

where R is some scalar identified with the Ricci scalar, and $\epsilon_{\mu\nu}$ and $\epsilon_{\alpha\beta}$ are antisymmetric Levi-Civita tensors on two indices, i.e. $\epsilon_{01} = \sqrt{|g|}$ and $\epsilon_{10} = -\sqrt{|g|}$. Due to Equation (B.12) and to the identity $\epsilon_{\mu\nu} \epsilon_{\alpha\beta} = g_{\mu\beta} g_{\nu\alpha} - g_{\mu\alpha} g_{\nu\beta}$, the Einstein tensor:

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu}$$

identically vanish on Σ . This is problematic since the vacuum Einstein field equation is $G_{\mu\nu} = 0$. The fact that it is automatically satisfied in 2-dimensional gravity shows that 2-dimensional

gravity without matter does not yield propagating gravitational modes. That is why physicists usually allow the Einstein-Hilbert Lagrangian to take more intricate forms in 2-dimensions. One of particular importance is a the $f(R)$ -gravity, in which the Ricci scalar is replaced by a well-behaved function:

$$\int_{\Sigma} \star(f(R)) = - \int_{\Sigma} \sqrt{|g|} f(R) dx^0 dx^1$$

Then one may show that under rather common assumptions, this action can be rewritten in terms of an auxiliary field ϕ called the *dilaton* and another well-behaved function $V(\phi)$:

$$\int_{\Sigma} \sqrt{|g|} (\phi R - V(\phi)) dx^0 dx^1 \quad (\text{B.13})$$

See for example Section 7 of [Schmidt, 1999] for an explicit treatment of this replacement.

Exercise B.11. Show that, for $f(R) = R^2$, we have the usual Gaussian integral:

$$\frac{1}{2} \int_{\Sigma} \sqrt{|g|} R^2 dx^0 dx^1 = \int_{\Sigma} \sqrt{|g|} (\phi R - \frac{1}{2} \phi^2) dx^0 dx^1$$

Let us now rewrite the $f(R)$ -lagrangian using zweibein and a spin connection, à la Palatini (see Chapter 3 Part III of [Baez and Muniain, 1994] for a treatment of Palatini formalism in n dimensions). The idea (in two dimensions) is the following: the metric g is locally diagonalizable, and even better, by a diagonal matrix of the form:

$$g \sim \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Then, in the neighborhood of every point, there exist two locally defined vector fields e_+ and e_- defining a frame of $T\Sigma$, such that in the local coordinate defined by this frame g takes the above diagonal form. We call the pair (e_+, e_-) a *zweibein* – the 2-dimensional analogs of tetrads in 4-dimensions and of vielbein in n dimensions. In particular, noting e^a for the differential one-form on Σ dual to e_a , where $a = \pm$, we have $e^+ \wedge e^- = \sqrt{|g|} dx^0 dx^1$. While the Einstein-Hilbert Lagrangian in n dimensions is invariant under diffeomorphisms, its reformulation in terms of vielbein is only invariant under the gauge group $SO(n-1, 1)$ (encoding every possible Lorentz rotations of the orthonormal frame). In two dimensions, this group is one dimensional, hence abelian. The gauge invariance under the Lorentz group is encoded by a connection ω called the *spin connection*. It is a differential 1-form on Σ taking values in $\mathfrak{so}(1, 1)$ (or $\mathfrak{so}(1, 1)$ when working on a n -dimensional space-time), satisfying the following compatibility condition:

$$De^a \equiv de^a + \omega_b^a e^b = 0 \quad (\text{B.14})$$

This condition implies that ω is uniquely expressed in terms of the zweibein and its dual. Since the gauge group $SO(1, 1)$ is abelian, the curvature of the spin connection reduces to $d\omega$, so that we have:

$$R \sqrt{|g|} dx^0 dx^1 = -2d\omega$$

To implement the constraint (B.14) in the $f(R)$ -lagrangian, one has to introduce two Lagrange multipliers X_+, X_- , so that the action (B.13) can be rewritten as:

$$\int_{\Sigma} \phi d\omega + X_a De^a + \frac{1}{2} V(\phi) e^+ \wedge e^- = \int_{\Sigma} \underbrace{\omega \wedge d\phi + e^a \wedge dX_a}_{A_i dX^i} + \underbrace{X_a \omega_b^a e^b + \frac{1}{2} V(\phi) e^+ \wedge e^-}_{\frac{1}{2} \pi^{ij}(X) A_i \wedge A_j}$$

The expression on the right-hand side corresponds to a Poisson-sigma model, where the Poisson manifold M is \mathbb{R}^3 , where the scalar function $X : \Sigma \rightarrow M$ is the triplet (X^+, X^-, ϕ) and where the differential one-form $A \in \Omega^1(\Sigma, X^! T^* M)$ is the triplet (e^+, e^-, ω) . This shows that the action functional of $f(R)$ 2-dimensional gravity (without matter) can be expressed as a Poisson-sigma model. There are additional applications of this model to other topological field theories.

B.3 The symplectic structure of phase space and canonical transformations

In this section we will provide some material non-related to canonical Hamiltonian formalism but which however is relevant for those interested into the geometrization of classical mechanics. The tangent bundle of the cotangent bundle $T(T^*Q)$ is a $4n$ -dimensional manifold and a rank $2n$ vector bundle over T^*Q . Over a local trivializing open set $U \subset Q$, the cotangent bundle T^*Q can be trivialized as the product of U (with coordinate functions q^i) with the fiber \mathbb{R}^n (with coordinate functions p_i), so that $T(T^*Q)$ is locally isomorphic to $TU \times T\mathbb{R}^n$. The same holds for the cotangent bundle of the cotangent bundle $T^*(T^*Q)$. A differential form on T^*Q is a section of $\wedge^\bullet T^*(T^*Q)$; since T^*Q can be locally seen as the product of a trivializing open set U and the fiber \mathbb{R}^n , one understands that a differential form on T^*Q is locally generated by products of covector fields dq^i on Q with covector fields dp_i on the fiber. Here the de Rham differential is the de Rham differential on T^*Q , so that dq^i should indeed be seen as a constant section of $T^*(T^*Q)$, although its action is trivial on the fiber.

With these conventions in hand, we observe that the cotangent bundle T^*Q is a symplectic manifold: there exists a closed non-degenerate 2-form $\omega \in \Omega^2(T^*Q)$ called the *Poincaré 2-form*. This two form is canonical (see below) and can be written in local coordinates q^i, p_i as:

$$\omega = \sum_{i=1}^n dp_i \wedge dq^i \quad (\text{B.15})$$

The Poincaré 2-form has the particularity of being exact and this represents a defining feature: it is the de Rham derivative $\omega = d\theta$ of the so-called *Liouville-Poincaré* – or *tautological* – *one-form*. This differential form is the unique globally defined one-form whose expression in local coordinates is $\theta = \sum_{i=1}^n p_i dq^i$. It admits a coordinate-free definition that we will now provide.

Definition B.12. *Let $\pi_Q : T^*Q \rightarrow Q$ be the projection associated to the bundle structure on the cotangent bundle of Q . Then, the tautological Liouville-Poincaré one-form is the unique differential one-form on T^*Q defined pointwise as follow:*

$$\text{for every } (q, p) \in T^*Q \quad \theta_{(q,p)} = p \circ (\pi_Q)_*|_{(q,p)}$$

where $(\pi_Q)_*|_{(q,p)} : T_{(q,p)}(T^*Q) \rightarrow T_qQ$ is the push-forward of π_Q at the point (q, p) .

We can understand this 1-form from the *double vector bundle* perspective [Mackenzie, 1992]. The tangent bundle of the cotangent bundle of Q is a double vector bundle over the cotangent bundle $T^*Q \xrightarrow{\pi_Q} Q$ and the tangent bundle $TQ \xrightarrow{\pi_Q} Q$, as it sits in the following commutative square [Schätz, 2009]:

$$\begin{array}{ccc}
 & T(T^*Q) & \\
 \pi_{T^*Q} \swarrow & & \searrow (\pi_Q)_* \\
 T^*Q & & TQ \\
 \pi_Q \searrow & & \swarrow \pi_Q \\
 & Q &
 \end{array}$$

where $\pi_{T^*Q} : T^*(T^*Q) \rightarrow T^*Q$ is the projection associated to the bundle structure on the cotangent bundle of T^*Q . Then, the tautological Liouville-Poincaré one-form is the differential one-form on T^*Q defined as:

$$\begin{aligned} \theta : T(T^*Q) &\longrightarrow \mathbb{R} \\ X &\longmapsto \langle \pi_{T^*Q}(X), (\pi_Q)_*(X) \rangle \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ denotes the fibre pairing between T^*Q and TQ . This definition is equivalent to that of Definition B.12.

The tautological one-form is the unique one-form that ‘cancels’ pullback: any differential one-form $\sigma \in \Omega^1(M)$ can be seen as a smooth section $\sigma : M \rightarrow T^*Q$. Then, the push-forward of σ is a vector bundle map $\sigma_* : TQ \rightarrow T(T^*Q)$, as well as the pull-back $\sigma^* : T^*(T^*Q) \rightarrow T^*Q$. Then, pulling back a differential one form on T^*Q via σ gives a differential one form on Q . Then, the tautological one-form is the unique one-form on T^*Q such that:

$$\sigma^*(\theta) = \sigma$$

Another way of saying this is the following: the tautological one-form is the only differential one-form θ on T^*Q such that:

$$\theta = \pi_Q^* \circ \pi_{T^*Q}(\theta)$$

where $(\pi_Q)^* : T^*Q \rightarrow T^*(T^*Q)$ is the pull-back of π_Q .

$$\begin{array}{ccc} \theta \in T^*(T^*Q) & \xrightarrow{\text{id}} & \theta \in T^*(T^*Q) \\ & \searrow \pi_{T^*Q} & \nearrow (\pi_Q)^* \\ & T^*Q & \end{array}$$

The tautological one-form can be used to characterize the Legendre transform between TQ and T^*Q in a more general and abstract way [Tulczyjew, 1977, Yoshimura and Marsden, 2006a, Yoshimura and Marsden, 2006b].

The canonical symplectic form on the phase space T^*Q is defined as the de Rham derivative of the tautological one form:

$$\omega = d\theta \tag{B.16}$$

The diffeomorphisms of T^*Q leaving the symplectic form invariant are called *symplectomorphisms*. In our case, they coincide with the Poisson isomorphisms of the corresponding (non-degenerate) Poisson structure. A symplectic manifold is always even-dimensional, so there is a neighboring notion for odd-dimensional smooth manifolds:

Definition B.13. A contact manifold is an odd-dimensional smooth manifold M equipped with a closed 2-form ω of maximal rank. An exact contact manifold consists of a $2n + 1$ -dimensional smooth manifold M and a 1-form θ on M such that $\theta \wedge (d\theta)^n$ is a volume form on M .

Example B.14. Let $M = T^*Q \times \mathbb{R}$ and $\theta_M = \theta + dt$ where θ is the tautological one-form on the cotangent bundle of Definition B.12. Then, (M, θ_M) is an exact contact manifold. There are similar results as Darboux theorem 3.54 for contact geometry, e.g. every exact contact manifold locally looks like $T^*Q \times \mathbb{R}$ (see Theorem 5.1.5 in [Abraham and Marsden, 1987]).

Remark B.15. Let $L \in C^\infty(TQ)$ be a Lagrangian and let $\mathcal{L} : TQ \rightarrow T^*Q$ be the Legendre transform defined from it. Then the canonical symplectic form on T^*Q can be pulled back to a 2-form on TQ through the pull-back of the Legendre transform:

$$\omega_{TQ} = \mathcal{L}^*(\omega)$$

This 2-form is closed hence *presymplectic*, and it is symplectic if and only if \mathcal{L} is invertible [Blohmann, 2023].

Contact geometry is a good setup to work with time-(in)dependent hamiltonians. We will only work with cotangent bundles and set $\pi : T^*Q \times \mathbb{R} \rightarrow T^*Q$ to be the projection on the first factor. For any smooth function $H \in C^\infty(T^*Q \times \mathbb{R})$ (possibly not depending on time), one defines:

$$\theta_H = \sum_{i=1}^n p_i dq^i - H dt \quad \text{and} \quad \omega_H = d\theta_H = \sum_{i=1}^n dp_i \wedge dq^i - dH \wedge dt$$

Then by construction $T^*Q \times \mathbb{R}$ is a contact manifold, but may not be exact, unless the following condition holds:

Proposition B.16. *The contact manifold $T^*Q \times \mathbb{R}$ is an exact contact manifold if and only if the smooth function $L = p_i \frac{\partial H}{\partial p_i} - H$ is nowhere vanishing.*

Proof. It is a straightforward calculation to show that:

$$\theta_H \wedge (d\theta_H)^n = n \left(H - p_i \frac{\partial H}{\partial p_i} \right) dq^1 \wedge \dots \wedge dq^n \wedge dp_1 \wedge \dots \wedge dp_n \wedge dt$$

It can be computed for $n = 2$ and then generalized to every n . □

If the function H is understood to be the Hamiltonian of the system then the function L obviously has the role of the Lagrangian. Contact geometry is useful to handle some kinds of transformations used by physicists:

Definition B.17. *Physicists call canonical transformation⁴⁴ a local transformation of the phase space preserving the form of Hamilton's equations of motion.*

What they mean is the following: assume that we work in local coordinates and that we have a symplectomorphism sending a time t a chart $V \subset T^*Q$ equipped with coordinates (Q^i, P_i) to the chart $U \subset T^*Q$ with coordinates (q^i, p_i) . It may not be the same chart, so one has to understand this map to send bijectively V onto U , and points labelled (Q, P) to points labelled (q, p) ⁴⁵. In the set of coordinates (q^i, p_i) , Hamilton's equations of motion are the following:

$$\dot{q}^i = \{q^i, H_t\} \quad \text{and} \quad \dot{p}_i = \{p_i, H_t\} \tag{B.17}$$

Here, H_t denotes the smooth function on U at time t , i.e. $H_t(q, p) = H(q, p, t)$ (notice that H may not depend of time) and the Poisson bracket is thus well-defined. Saying that the transformation $(Q, P) \mapsto (q, p)$ is a canonical transformation means that there exists a smooth function K defined on $V \times \mathbb{R}$ such that the coordinate functions (Q^i, P_i) satisfy the following equations:

$$\dot{Q}^i = \{Q^i, K_t\} \quad \text{and} \quad \dot{P}_i = \{P_i, K_t\} \tag{B.18}$$

⁴⁴There are several, non equivalent notions of canonical transformations, see this [quick review](#).

⁴⁵Usually physicists use the opposite transformation $(q^i, p_i) \mapsto (Q^i, P_i)$, but since a canonical transformation is invertible this is just a matter of convention. We have chosen a convention which makes sense of Hamilton's principal function in Example (B.27) as the generating function of the flow of the vector field $X_{-H} = \{ \cdot, H \}$.

Moreover, H_t may be defined on V as well, but since in general $H_t(q, p) \neq H_t(Q, P)$, the solutions of Equations (B.18) are *not* solutions of the following equations:

$$\dot{Q}^i = \{Q^i, H_t\} \quad \text{and} \quad \dot{P}_i = \{P_i, H_t\} \quad (\text{B.19})$$

where on both right-hand sides the bracket are assumed to be evaluated at (Q, P) . By assumption, Equations (B.18) are equivalent to Equations (B.17) because they result from performing a diffeomorphism on the phase space. However, Equations (B.19) are not equivalent to Equations (B.17) because they are defined with the same hamiltonian, on another chart V than U . The usefulness of canonical transformations come from the idea that it might be easier to solve Equations (B.18) than Equations (B.17), thus justifying that we look for an adapted set of coordinates and a new Hamiltonian. For simplicity we will often assume that the coordinates are globally defined, i.e. $U = T^*Q$ parametrized by the coordinates (q^i, p_i) and $V = T^*Q$ parametrized by the coordinates (Q^j, P_j) . Now let us give our own definition of canonical transformations, adapted from Definition 5.2.6 in [Abraham and Marsden, 1987] for the purpose of the present setup.

Definition B.18. *Mathematicians call canonical transformation a smooth family $(\varphi_t : T^*Q \rightarrow T^*Q)_{t \in \mathbb{R}}$ of symplectomorphisms such that the induced diffeomorphism of contact manifolds:*

$$\begin{aligned} \phi : T^*Q \times \mathbb{R} &\longrightarrow T^*Q \times \mathbb{R} \\ ((Q, P), t) &\longmapsto (\varphi_t(Q, P), t) \end{aligned}$$

satisfies the following condition:

for all $H \in \mathcal{C}^\infty(T^*Q \times \mathbb{R})$ there exists $K \in \mathcal{C}^\infty(T^*Q \times \mathbb{R})$ such that $\phi^*(\omega_H) = \omega_K$.

Remark B.19. As we will see in Example B.27, it does not mean that $K = \phi^*(H)$. Even though every symplectomorphism φ_t preserves Hamilton's equation in the sense that the integral curves of $X_{\varphi_t^*(H_t)}$ are the same of that of X_{H_t} (they are sent to one another via φ_t). This can be seen from the observation that any symplectomorphism ψ has the property that for every smooth function $f \in \mathcal{C}^\infty(T^*Q)$:

$$\psi_*(X_{\psi^*f}) = X_f$$

This result is Theorem 3.3.19 in [Abraham and Marsden, 1987] which itself is adapted from a result by Jacobi from 1837, and in turn implies that γ is an integral curve of $X_{\psi^*(f)}$ if and only if $\sigma = \psi \circ \gamma$ is an integral curve of X_f :

$$\dot{\sigma} = \frac{d(\psi \circ \gamma)}{dt} = \psi_*(\dot{\gamma}) = \psi_*(X_{\psi^*f}) = X_f$$

Let us deduce some consequence from Definition B.17. From here, we assume that that the coordinates q^i, p_i and Q^j, P_j are globally defined. We know that Hamilton's equations of motion descend from the Hamilton-Pontryagin action principle and its variation represented by the action (4.16). Since we assumed that there are no constraints, i.e. that the Legendre transform is invertible, we can solve the velocities v in terms of the momenta p and rewrite action (4.16) as:

$$S = \int_{t_1}^{t_2} (p_i \dot{q}^i - H(q, p, t)) dt = \int_{\sigma} p_i dq^i - H dt = \int_{\sigma} \theta_H \quad (\text{B.20})$$

where the integration is made over a path Σ in $T^*Q \times \mathbb{R}$. We apply the same argument to the Hamiltonian K (also called 'Kamiltonian'); because the coordinates (Q^i, P_i) satisfy Equations (B.18), the latter descend from the following action:

$$S = \int_{\Gamma} P_i dQ^i - K dt = \int_{\Gamma} \theta_K \quad (\text{B.21})$$

where Γ is the path in $T^*Q \times \mathbb{R}$ that is the inverse image of Σ through ϕ , i.e. $\phi(\Gamma) = \Sigma$. The diffeomorphism $\phi : T^*Q \times \mathbb{R} \rightarrow T^*Q \times \mathbb{R}$ preserves linear integrals (see subsection 7.1 in [Bamberg and Sternberg, 1988b]) so we have that:

$$\int_{\Sigma} \theta_H = \int_{\Gamma} \phi^*(\theta_H)$$

Now since the canonical transformation ϕ transforms Hamilton's equations (B.17) into (B.18), the two actions (B.20) and (B.21) should be equivalent, i.e. equal:

$$\int_{\Gamma} \phi^*(\theta_H) - \theta_K = 0$$

Equality between the two actions does not mean that the integrand $\phi^*(\theta_H) - \theta_K$ is zero, but rather than it is equal to a total differential dW . That is to say, there exists a smooth function $W \in C^\infty(T^*Q \times \mathbb{R})$ such that:

$$\phi^*(\theta_H) - \theta_K = dW \tag{B.22}$$

This equation was obtained from a physical argument, but it could be obtained from another mathematical argument. The condition that ϕ is a canonical transformation is that $\phi^*(\omega_H) = \omega_K$. While the left-hand side is $\phi^*(d\theta_H)$, the right-hand side is $d\theta_K$. Thus we have that:

$$\phi^*(\omega_H) = \omega_K \iff d(\phi^*(\theta_H) - \theta_K) = 0$$

that is to say the differential one-form $\phi^*(\theta_H) - \theta_K$ is closed, hence locally exact, which is the meaning of Equation (B.22). The fact that we find the same equation following a physical argument from Definition B.17 and a mathematical argument from Definition B.18 show that the two definition are more or less equivalent.

Definition B.20. *Mathematicians call generating function of the canonical transformation $(Q, P) \mapsto (q, p)$ the smooth function $W \in C^\infty(T^*Q \times \mathbb{R})$ satisfying Equation (B.22).*

We will see that the notion of generating function in physics is quite close to that. Now let us explore the mathematical consequences of Equation (B.22). It translates as:

$$\varphi_t^*(p_i dq^i) - \varphi_t^*(H)dt = P_i dQ^i - K dt + dW \tag{B.23}$$

The first term reads:

$$\varphi_t^*(p_i dq^i) = (p_i \circ \varphi_t) d(q^i \circ \varphi_t) = \varphi_{t,i} \left(\frac{\partial \varphi_t^i}{\partial Q^j} dQ^j + \frac{\partial \varphi_t^i}{\partial P_j} dP_j \right)$$

where we denoted $\varphi_t^i = q^i \circ \varphi_t$ and $\varphi_{t,i} = p_i \circ \varphi_t$; they are smooth function on V . From this we deduce that:

$$\varphi_{t,i} \frac{\partial \varphi_t^i}{\partial Q^j} - P_j = \frac{\partial W}{\partial Q^j} \tag{B.24}$$

$$\varphi_{t,i} \frac{\partial \varphi_t^i}{\partial P_j} = \frac{\partial W}{\partial P_j} \tag{B.25}$$

$$K - \varphi_t^*(H) = \frac{\partial W}{\partial t} \tag{B.26}$$

The first two equations are still too much coupled because $\varphi_{t,i}$ appears in both, so to simplify it we will assume soon that $\frac{\partial \varphi_t^i}{\partial Q^j} = 0$ for $1 \leq i, j \leq n$.

Remark B.21. As a side remark, notice that differentiating one more time Equations (B.24) and (B.25) with respect to Q^i, P_i induce consistency conditions that have to be satisfied by the double derivatives on the right-hand side, and thus by the left-hand sides as well. This leads to the introduction of what is called the *Lagrange bracket*, which can be considered as ‘inverse’ to the Poisson brackets, see pp. 36-37 of [Sudarshan and Mukunda, 1974]. In particular, we find that Q^i, P_i are canonical coordinates, i.e. $\{Q^i, Q^j\} = \{P_k, P_l\} = 0$ and $\{Q^i, P_j\} = \delta_j^i$, hence justifying the name of *canonical transformations*.

The $2n \times 2n$ Jacobian matrix representing $\varphi_t : T^*Q \mapsto T^*Q$ as a symplectic matrix is of the form:

$$\begin{pmatrix} A_t & B_t \\ C_t & D_t \end{pmatrix}$$

where each bloc is a $n \times n$ square matrix. The coefficients of A_t correspond to the partial derivatives $\frac{\partial \varphi_t^i}{\partial Q^j}$, i.e. the order of the lines and columns respects the order of the coordinates (q^i, p_i) and (Q^j, P_j) . The properties of symplectic matrices imply that $A^T C$ and $B^T D$ are symmetric and $A^T D - C^T B = I_n$. From now on, then, we will make the following assumption: that φ_t^i does not depend on Q^j , that is to say, $\frac{\partial \varphi_t^i}{\partial Q^j} = 0$ for $1 \leq i, j \leq n$. Equivalently, it means that $A_t = 0$, so $-C^T B = I_n$ i.e. B is invertible and its inverse is $-C^T$. The matrix B_t is thus a $n \times n$ matrix of maximal rank with (i, j) -th coefficient $B_t^{ij} = \frac{\partial \varphi_t^i}{\partial P_j}$. So (up to projection on the base manifold) it induces a local diffeomorphism ψ_t between the fibers of T^*Q parametrized by P_j and the base manifold Q parametrized by q^i . The assumption that the φ_t^i 's do not depend on the Q^j 's can be shown to hold in most general cases for regular Lagrangians, see pp. 63-64 in [Sudarshan and Mukunda, 1974].

From these assumptions, Equations (B.24) and (B.25) become:

$$P_j = -\frac{\partial W}{\partial Q^j} \quad \text{and} \quad \varphi_{t,i} = \frac{\partial \varphi_t^i}{\partial P_j} = \frac{\partial W}{\partial P_j}$$

and, multiplying the last one by B_t^{-1} , one obtains the following formula:

$$\varphi_{t,i} = \frac{\partial W}{\partial P_j} \frac{\partial P^j}{\partial \varphi_t^i} \tag{B.27}$$

Now the dependence of W on P can now be understood to become a dependence on q given the local diffeomorphism ψ_t between the fibers of T^*Q and the base via B_t . At the cost of using globally defined coordinates, we thus let $F \in C^\infty(Q \times Q \times \mathbb{R})$ to be the smooth function induced by W as:

$$F(q, Q, t) = W(Q, \psi_t^{-1}(q), t)$$

Then, one can now rewrite Equation (B.27) under the more famous form:

$$p_i = \frac{\partial F}{\partial q^i}$$

More generally, Equations (B.24)-(B.26) can be rewritten, with some abuse of notation, as:

$$p_i = \frac{\partial F}{\partial q^i} \tag{B.28}$$

$$P_i = -\frac{\partial F}{\partial Q^i} \tag{B.29}$$

$$K = H + \frac{\partial F}{\partial t} \tag{B.30}$$

where the last equation should be rigorously understood $K(Q, \psi_t^{-1}(q), t) = \varphi_t^*(H)(Q, \psi_t^{-1}(q), t) + \frac{\partial F}{\partial t}$. Then, F satisfies the following identity (non-rigorous rewriting of Equation (B.23)):

$$p_i dq^i - H dt = P_i dQ^i - K dt + dF \quad (\text{B.31})$$

and this characterizes what physicists call a generating function:

Definition B.22. *Physicists call generating function (of the first kind) of the canonical transformation $(Q, P) \mapsto (q, p)$ the smooth function $F \in C^\infty(Q \times Q \times \mathbb{R})$ satisfying Equation (B.31).*

Notice that in the logical steps leading to both Definition B.20 and Definition B.22, we have deduced the existence of the generating function from a choice of Hamiltonian and a canonical transformation. This resulted into a smooth function that links the Hamiltonian H and the ‘Kamiltonian’ K . Conversely, given a Hamiltonian and a Kamiltonian, a generating function (of the first kind), as its name indicates, is used to generate the corresponding canonical transformation between old and new coordinates. Then we want to make sense, geometrically, of the correspondence between canonical transformations and generating functions of the first kind. Notice that the choice of the sign in front of Equations (B.28)-(B.30) is arbitrary as one could define the generating function to be $-F$. The current sign comes from the fact that F sits on the right-hand side of Equation (B.31), and this is a convention. Eventually be aware that there exist several other kinds of generating functions, depending on the dynamical variables upon which it depends. They are related to one another through Legendre transforms. Depending on the canonical transformation, it is sometimes impossible to find a generating function of the first kind.

For simplicity let us now assume that q^i and Q^i are global coordinates on the configuration manifold Q (which can then be identified with \mathbb{R}^n). Then the canonical transformation $\phi : T^*Q \times \mathbb{R} \rightarrow T^*Q \times \mathbb{R}$ can be interpreted as a family of symplectomorphisms depending on time:

$$\begin{aligned} \varphi_t : T^*Q &\longrightarrow T^*Q \\ (Q, P) &\longmapsto (q, p) \end{aligned}$$

The graph of φ_t at time t is denoted $\text{Gr}(\varphi)$ and is an embedded submanifold of the product manifold $T^*Q \times T^*Q$. Here, it is understood that the first factor is parametrized by (q^i, p_i) and the second one by (Q^i, P_i) , so that a point on $T^*Q \times T^*Q$ can be generally written as (q, p, Q, P) (do not mistake that Q with the configuration manifold). This product manifold is a vector bundle of $Q \times Q$, with projection $(q, p, Q, P) \mapsto (q, Q)$. Then the following Proposition, first introduced in Example 3.89 will show itself useful:

Proposition B.23. *If M_1, M_2 are symplectic manifolds, then a diffeomorphism $\varphi : (M_1, \omega_1) \longrightarrow (M_2, \omega_2)$ is a symplectomorphism if and only if its graph $\text{Gr}(\varphi) \subset M_2 \times M_1$ is a Lagrangian submanifold of $(M_2 \times M_1, pr_{M_2}^*(\omega_2) - pr_{M_1}^*(\omega_1))$.*

Proof. See for example Proposition 1.1 in [Tulczyjew, 1977] or Proposition 5.2.1 in [Abraham and Marsden, 1987]. \square

Since the cotangent bundle T^*Q is a symplectic manifold with canonical symplectic form defined in Equation (B.15), the product vector bundle $T^*Q \times T^*Q$ canonically inherits a symplectic structure, given by the difference of the pullbacks of the canonical symplectic structure on T^*Q . More precisely, letting pr_1 (resp. pr_2) be the projection on the first (resp. second) factor of the product $T^*Q \times T^*Q$, we define:

$$\omega \ominus \omega = pr_1^*(\omega) - pr_2^*(\omega) = \sum_{i=1}^n dp_i \wedge dq^i - dP_i \wedge dQ^i$$

By Proposition B.23, the graph of φ_t is a Lagrangian submanifold of $(T^*Q \times T^*Q, \omega \ominus \omega)$ for every time t . Let us chose a Hamiltonian $H \in \mathcal{C}^\infty(T^*Q \times \mathbb{R})$ then because ϕ is a canonical transformation, there exists a ‘Kamiltonian’ $K \in \mathcal{C}^\infty(T^*Q \times \mathbb{R})$ satisfying $\phi^*(\omega_H) = \omega_K$. As a smooth function on $Q \times Q \times \mathbb{R}$, a generating function F of the canonical transformation ϕ with Hamiltonian H and ‘Kamiltonian’ K can be seen as a family of smooth functions $F_t \in \mathcal{C}^\infty(Q \times Q)$. Then its differential is a section of $T^*(Q \times Q)$, i.e. it can be seen as a smooth function $dF_t : Q \times Q \longrightarrow T^*(Q \times Q)$. The section dF_t being smooth, its image $\text{Im}(dF_t)$ is an embedded submanifold of $T^*(Q \times Q)$.

The cotangent bundle $T^*(Q \times Q)$ admits fiberwise linear coordinates on $T^*_{(q,Q)}(Q \times Q)$ denoted by \bar{p}_i and \bar{P}_i . In other words, a point in the cotangent bundle $T^*(Q \times Q)$ will be denoted (q, Q, \bar{p}, \bar{P}) . As a cotangent bundle, this manifold comes equipped with a canonical symplectic structure (B.15):

$$\omega \oplus \omega = \sum_{i=1}^n d\bar{p}_i \wedge dq^i + d\bar{P}_i \wedge dQ^i$$

which is the differential of the canonical tautological form:

$$\theta = \sum_{i=1}^n \bar{p}_i dq^i + \bar{P}_i dQ^i$$

Then, the restriction of the tautological form to the image of dF_t has the following property (many arguments of this discussion come from [Tulczyjew, 1977]):

Lemma B.24. *Let $\pi : T^*(Q \times Q) \rightarrow Q \times Q$ be the vector bundle projection. The tautological form θ satisfies:*

$$\theta|_{\text{Im}(dF_t)} = \pi^*(dF_t)$$

The fact that pullbacks commute with the de Rham differential together with Equation (B.16) have the following consequence: $\omega \oplus \omega|_{\text{Im}(dF_t)} = 0$. Thus, the image of the differential dF_t is a Lagrangian submanifold of $T^*(Q \times Q)$. Let us see how to parametrize it. The differential dF_t reads $dF_t = \frac{\partial F_t}{\partial q^i} dq^i + \frac{\partial F_t}{\partial Q^i} dQ^i$, and, by Equations (B.28) and (B.29), we deduce that:

$$dF_t = p_i dq^i - P_i dQ^i$$

where here p_i and P_i are the components of admissible momenta, i.e. those of the points $(q, p, Q, P) \in T^*Q \times T^*Q$ such that $(q, p) = \varphi(Q, P)$. Thus, the Lagrangian submanifold $\text{Im}(dF_t) \subset T^*(Q \times Q)$ admits the following parametrization:

$$\text{Im}(dF_t) = \left\{ (q, Q, \bar{p}, \bar{P}) = (q, Q, p, -P) \mid (q, p, Q, P) \in \text{Gr}(\varphi) \right\}$$

and by construction, it geometrically encodes Equations (B.28) and (B.29).

Then we observe that there exists a canonical symplectomorphism between $(T^*Q \times T^*Q, \omega \ominus \omega)$ and $(T^*(Q \times Q), \omega \oplus \omega)$ defined as:

$$\begin{aligned} \Phi : T^*Q \times T^*Q &\longrightarrow T^*(Q \times Q) \\ (q, p, Q, P) &\longmapsto (q, Q, p, -P) \end{aligned}$$

This is obviously a diffeomorphism, with inverse $\Psi : (q, Q, \bar{p}, \bar{P}) \mapsto (q, \bar{p}, Q, -\bar{P})$. Then by construction we have that $\Phi^*(\bar{p}_i) = p_i$ and $\Phi^*(\bar{P}_i) = -P_i$. It implies that Φ preserves the

respective symplectic forms, because:

$$\begin{aligned}
 \Phi^*(\omega \oplus \omega) &= \Phi^* \left(\sum_{i=1}^n d\bar{p}_i \wedge dq^i + d\bar{P}_i \wedge dQ^i \right) \\
 &= \sum_{i=1}^n d\Phi^*\bar{p}_i \wedge d\Phi^*q^i + d\Phi^*\bar{P}_i \wedge d\Phi^*Q^i \\
 &= \sum_{i=1}^n d\bar{p}_i \wedge dq^i - d\bar{P}_i \wedge dQ^i \\
 &= \omega \ominus \omega
 \end{aligned}$$

Thus Φ is a symplectomorphism. This observation allows us to conclude on the relationship between F_t and φ_t :

Proposition B.25. *The symplectomorphism Φ induces a diffeomorphism between the Lagrangian submanifold $\text{Gr}(\varphi_t) \subset T^*Q \times T^*Q$ and the Lagrangian submanifold $\text{Im}(dF_t) \subset T^*(Q \times Q)$.*




Figure 24: We usually represent the direct product of two vector bundles by the figure on the left-hand side. Then a Lagrangian submanifold is represented as a line. The symplectomorphism $\Phi : T^*Q \times T^*Q \rightarrow T^*(Q \times Q)$ sends $\text{Gr}(\varphi_t)$ onto $\text{Im}(dF_t)$.

We can even improve this result even further by involving contact geometry. Let us show how to make the product manifold $T^*Q \times T^*Q \times \mathbb{R}$ a contact manifold. Given the choice of Hamiltonian H and ‘Kamiltonian’ K , the following closed 2-form on $T^*Q \times T^*Q \times \mathbb{R}$:

$$\Omega_1 = \omega \ominus \omega - (dH - dK) \wedge dt$$

is of maximal rank, because $\omega \ominus \omega$ is a symplectic form on $T^*Q \times T^*Q$. This 2-form turns the product manifold $T^*Q \times T^*Q \times \mathbb{R}$ into a contact manifold. We can also make the product manifold $T^*(Q \times Q) \times \mathbb{R}$ a contact manifold when equipped with the following closed 2-form of maximal rank:

$$\Omega_2 = \omega \oplus \omega + d \left(\frac{\partial F}{\partial t} \right) \wedge dt$$

The two forms Ω_1 and Ω_2 descend, respectively, from the following canonical one-forms:

$$\Theta_1 = \sum_{i=1}^n p_i dq^i - P_i dQ^i - (H - K)dt \quad \text{and} \quad \Theta_2 = \sum_{i=1}^n \bar{p}_i dq^i + \bar{P}_i dQ^i + \frac{\partial F}{\partial t} dt$$

It is now straightforward to see that the symplectomorphism Φ induces a diffeomorphism:

$$\begin{aligned} \Psi : T^*Q \times T^*Q \times \mathbb{R} &\longrightarrow T^*(Q \times Q) \times \mathbb{R} \\ (q, p, Q, P, t) &\longmapsto (q, Q, p, -P, t) \end{aligned}$$

which satisfies the following property:

Proposition B.26. *The diffeomorphism Ψ is a contact transformation, i.e. $\Psi^*(\Theta_2) = \Theta_1$.*

Proof. We observe that $\Psi = \Phi \times \text{id}_{\mathbb{R}}$, and we already know that $\Phi^*(\Theta_2) = \Theta_1$. The only thing left is to show that the pullback by Ψ of $\frac{\partial F}{\partial t} dt$ is indeed $-(H - K)dt$. But this is the case since Equation (B.30) tells us that. Hence the property that Ψ is a contact transformation geometrically encodes that equation. \square

We have proved Proposition B.25 and Proposition B.26 under the assumption that the coordinates q^i and Q^i were globally defined over Q , thus identifying it with \mathbb{R}^n . However, in full generality, they would only be defined locally. This implies in turn that F is only a locally defined function and that the statement of Proposition B.25 is only a local one. This is not a problem because Definitions B.20 and B.22 involve in theory only locally defined generating functions. This is moreover consistent with the following observations: given a canonical transformation $\phi = (\varphi_t)_t$, the image through Φ of the graph $\text{Gr}(\varphi_t)$ is an embedded Lagrangian submanifold N_t of $T^*(Q \times Q)$. It means that on N_t , the symplectic form $\Omega = d\Theta$ is vanishing, turning $\Theta|_{N_t}$ into a closed form. By Poincaré Lemma, it is locally exact (on N_t), i.e. there exists a smooth function $\tilde{F}_t \in C^\infty(N_t)$ such that $\Theta|_{N_t} = d\tilde{F}_t$. If the canonical transformation has nice features, then N_t is diffeomorphic to $Q \times Q$ and the function \tilde{F}_t can be seen as a function F_t on $Q \times Q$, which is the legitimate generating function of the first kind. This is how one can obtain the local expression of the generating function F of any amenable canonical transformation ϕ .

Example B.27. The Hamilton principal function is the smooth function S that depends on three variables: an initial position q_0 , a final position q , and a given time t . If the time $0 < |t| < \epsilon$ is sufficiently small, it is understood that there is a unique path $\gamma :]-\epsilon, \epsilon[\rightarrow Q$ which satisfies the Euler-Lagrange equations and is such that $\gamma(0) = q_0$ and $\gamma(t) = q$. The latter condition indeed fixes the initial velocity v_0 , and hence initial momentum p_0 . Since the Legendre transform is non degenerate, there is a one-to-one correspondence between $T_{q_0}Q$ and $T_{q_0}^*Q$, i.e. between initial velocity and initial momentum. Thus there is a unique initial momentum p_0 such that $\gamma(0) = q_0$, $\gamma(t) = q$, and $\tau \mapsto \mathcal{L}(\gamma(\tau), \dot{\gamma}(\tau))$ satisfies Hamilton's equations of motion (B.17). The final momentum is moreover uniquely defined, as it corresponds to $\mathcal{L}(\gamma(t), \dot{\gamma}(t))$.

Changing the initial position q_0 , the final position q , and the time t , defines another unique initial momentum p_0 and another unique final momentum p . Moreover, since the path γ and the Legendre transform are smooth, a smooth change in the position data q_0 and q induce a smooth change in the momentum data p_0 and p . We then have, for each fixed time t , initial position and final position, a diffeomorphism φ_t sending the pair (q_0, p_0) to the pair (q, p) , and this family $(\varphi_t)_t$ of diffeomorphisms is smoothly varying over t . By construction this family of diffeomorphisms is the flow of the vector field X_{-H} . The diffeomorphisms generated by the flow of a Hamiltonian vector field form a particular class of symplectomorphisms, called the Hamiltonian symplectomorphisms. Then the family $(\varphi_t)_t$ thus forms a family of symplectomorphisms φ_t sending a pair of initial data (q_0, p_0) to a pair of final data (q, p) at time t , reached through

an integral curve of X_{-H} , so every point on the curve satisfies Hamilton's equations of motion. This implies that the family of diffeomorphisms $(\varphi_t)_t$ preserves those equations, hence forming a canonical transformation.

The relationship between hamiltonian vector fields and canonical transformations is not innocent: actually *every canonical transformation is locally the flow of a Hamiltonian vector field*. This comes from the fact that any symplectic vector field – in particular the one induced by a canonical transformation – is locally a hamiltonian vector field. In the present example we have taken as this function precisely the Hamiltonian H , so that it will turn out that the generating function (of the first kind) corresponding to the canonical transformation ϕ is the Hamilton principal function, which is defined as:

$$S_t(q, q_0) = \int_0^t L(\gamma(\tau), \dot{\gamma}(\tau)) d\tau \quad (\text{B.32})$$

where γ in the integrand is the unique path $\gamma :]-\epsilon, \epsilon[\rightarrow Q$ which satisfies the Euler-Lagrange equations and is such that $\gamma(0) = q_0$ and $\gamma(t) = q$. We then see that the generating function is the action of the system. For each t the graph of the canonical transformation φ_t is diffeomorphic to the image of the differential dS_t . In particular in that context, Equations (B.28) and (B.29) become:

$$p_i = \frac{\partial S_t}{\partial q^i} \quad \text{and} \quad (p_0)_i = -\frac{\partial S_t}{\partial q_0^i} \quad (\text{B.33})$$

On the other hand, the particularity of the Hamilton principal function is that Equation (B.30) reduces to:

$$\frac{\partial S_t}{\partial t} + H\left(q, \frac{\partial S_t}{\partial q^i} dq^i, t\right) = 0 \quad (\text{B.34})$$

where the time-dependence is optional but we have included it in order to be complete. Indeed, the very definition of Hamilton's principal function, Equation (B.32), implies that $L(q, \dot{q}, t) = \frac{dS_t}{dt} = \frac{\partial S_t}{\partial t} + \frac{\partial S_t}{\partial q^i} \dot{q}^i$. However, since the Legendre transform is bijective, one can write the Hamiltonian in terms of the Lagrangian: $H(q, p, t) = p_i \dot{q}^i - L(q, \dot{q}, t)$, where the couple (q, \dot{q}) is understood to be $\mathcal{L}^{-1}(q, p)$. By Equation (B.33), one can replace p_i by the derivative of S_t with respect to q^i . Replacing the Lagrangian as well, one obtains:

$$H = \frac{\partial S_t}{\partial q^i} \dot{q}^i - \left(\frac{\partial S_t}{\partial t} + \frac{\partial S_t}{\partial q^i} \dot{q}^i \right) = -\frac{\partial S_t}{\partial t}$$

thus giving back Equation (B.34).

The particularity of the generating function S_t and the corresponding family of canonical transformations φ_t is then that the resulting new coordinates are such that the 'Kamiltonian' K in this system is zero. Let us explain in more details what is happening: since every point on the integral curve of H passing through the initial data (q_0, p_0) are the image of (q_0, p_0) through φ_t for some t , the coordinate functions $q_0^i, p_{0,i}$ are constant of motion (which is expected as an initial data is constant along the flow of time). Then it is not surprising that both left-hand sides of Equations (B.18) – once one has realized that $Q^i = q_0^i$ and $P_i = p_{0,i}$ – are zero. This is possible only if $K_t = 0$ for all t , which is precisely the meaning of Equation (B.34). To summarize, we do not lose information between Equations (B.17) and (B.18) because on the one-hand, either we have an Hamiltonian and Hamilton's equations (B.17) on the usual coordinates (q^i, p_i) , or we have the initial data (q_0, p_0) and the flow φ_t of the Hamiltonian, but no hamiltonian and differential equations. These two sets of data are equivalent, because obtaining the flow is equivalent to solving Hamilton's equations of motion. As a final remark, this example has shown that whereas the Hamiltonian describes the infinitesimal form of the canonical transformation generated by the motion in phase space, the time integral of the Lagrangian – Hamilton's

principal function, Equation (B.32) – describes the finite form of this canonical transformation. This reciprocal relationship has an analogy in quantum mechanics as well, and is most clearly seen in the formulation due to Feynman.

B.4 Canonical formalism in field theories

Contrary to classical physical systems with a finite number of degrees of freedom, field theories are considered to have an infinite number of such degrees of freedom. Indeed, while in the former case the physical variables are smooth functions denoted q^i and p_i and labelled by a finite index, in the latter case the physical variables are fields and the corresponding index which labels them is often understood as a continuous index symbolized by the physical position x in space-time at a given time. More precisely, a field φ evaluated at the position x at time t should be considered a different object than the same field evaluated at x' at time t (even if they respective values coincide at both points). In that context the position x is used as an infinite, continuous index which labels the fields, as was the case for the index i with the coordinate functions q^i and p_i . The generalization from a finite set of indices to an uncountable infinite set is due to the fact that the canonical Hamiltonian formalism in field theory treats space and time variables differently and describes classical fields as infinite-dimensional systems evolving in time. A thorough treatment of the field theoretic aspects of canonical Hamiltonian formalism can be found in [Gitman and Tyutin, 1990, Henneaux and Teitelboim, 1992, Rothe and Rothe, 2010]. A more covariant approach, treating space and time on an equal footing, is the De Donder-Weyl theory [Kanatchikov, 2001], and relies on multisymplectic or polysymplectic geometry.

To illustrate how Bergmann-Dirac algorithm and canonical Hamiltonian formalism apply to field theory, this subsection is entirely dedicated to one example: that of free Maxwell's theory of electromagnetism in flat space-time [Matschull, 1996, Blaschke and Gieres, 2021]. Additionally, this theory is a gauge theory so we should see many things appear that we studied earlier. Let $M = \mathbb{R}^4$ with metric η of signature (3, 1) i.e. $(-, +, +, +)$ in physical language. The action of Maxwell's electromagnetism in the void is:

$$S = -\frac{1}{4} \int_M d^4x F^{\mu\nu} F_{\mu\nu} \quad (\text{B.35})$$

Here, as explained in the discussion below Equations (1.38)-(1.41), $F = \frac{1}{2} F_{\mu\nu} dx^\mu \wedge dx^\nu$ is a differential 2-form on M , and is such that $F = dA$, where $A = A_\mu dx^\mu$ is the *gauge field* or *connection 1-form* associated to the covariant derivative $D = d + A$, needed to enforce the gauge symmetry of Maxwell's electromagnetism. In physical terms, A is the *vector potential*. The upper indices in Equation (B.35) are obtained through the procedure of raising indices via inverting the metric as was discussed in Section A.2. In particular, one has: $F^{\mu\nu} = \eta^{\mu\alpha} \eta^{\nu\beta} F_{\alpha\beta}$, where $\eta^{\mu\alpha}$ denotes the (μ, α) component of the matrix η^{-1} but luckily in the case of Minkowski space, it coincides with the original metric $\eta^{-1} = \eta$.

To proceed further, we need to distinguish space and time⁴⁶: space position will be used as a continuous index, while time will be used as the variable of integration, parametrizing the evolution of the fields, so that $S = \int L dt$. Thus, we denote by $t = x^0$ the time coordinate, will x^1, x^2, x^3 will correspond to the space coordinates. Accordingly, we can split the vector potential into a time component A_0 and space components A_1, A_2, A_3 (or A_x, A_y, A_z as in Equation (1.49)). Using the definition of the rule that raises and lowers the indices in non-euclidean space time,

⁴⁶Notice that this splitting into space and time coordinates breaks space-time covariance. This is an inherent problem of canonical Hamiltonian formalism and explains why it has not been quite met acclaim among the community of physicists, compared to e.g. the path-integral approach which preserves covariance.

the Lagrangian of Maxwell's electromagnetism can be written as:

$$L = \sum_{i,j=1,2,3} \int_{\mathbb{R}^3} d^3x \frac{1}{2} F_{0i} F_{0i} - \frac{1}{4} F_{ij} F_{ij} \quad (\text{B.36})$$

where the formula for the components of the field strength F can be found in Formula (1.44). In Maxwell's electromagnetism, we take the A_μ as the fields playing the roles of the "generalized coordinates", to which we can associate the velocities $V_\mu = \partial_0 A_\mu$. From Equation (B.36), and knowing that $F_{0i} = V_i - \partial_i A_0$, we then deduce the conjugate momenta associated to the A_μ :

$$\Pi^i(x) = \frac{\delta L}{\delta V_i(x)} = F_{0i}(x) = -E_i(x) \quad \text{while} \quad \Pi^0(x) = \frac{\delta L}{\delta V_0(x)} = 0 \quad (\text{B.37})$$

Notice how the space position x emerges in the present context: since the Lagrangian of Maxwell's theory is a integral of a lagrangian density, the correct way to compute a "partial derivative" of the Lagrangian is to use the variation with respect to a field, at a given position x .

From the first of Equations (B.37) and the Minkowski metric on space coordinates, we deduce that $\Pi_i = \Pi^i$ is minus the i -th component of the differential one-form $E = \vec{E}^b$ representing the electric field \vec{E} (see Equations (1.42) and (1.43)). In other words, in physical terms, the conjugate momenta of the vector potential is (minus) the electric field. On the other hand, the right-hand side of the second equation in (B.37) is automatically zero because L does not contain any term V_0 . Thus, this gives us an infinite set of constraints on the last conjugate momenta $\Pi_0 = -\Pi^0$:

$$\phi_x = -\Pi_0(x)$$

We chose to take $-\Pi_0$ as the constraint because the momenta associated to A_0 is $\Pi^0 = -\Pi_0$ because of the Minkowski metric $\eta_{00} = -1$. Notice that these primary constraints are labelled by the continuous index $x \in \mathbb{R}^3$ – thus illustrating the specificity of field theories – while before we had only a finite set of constraints. This constraint means that the velocity V_0 cannot be solved with respect to the conjugate momenta Π_μ . The fact that the Lagrangian density does not depend on V_0 reflects a degeneracy (related to gauge invariance) and means that A_0 does not really represent a dynamical variable. We will see later that it could be used as a gauge parameter.

By using the definition of F from both Formula (1.44) and Equation (1.48), we deduce that $F_{0i} \partial_0 A_i = F_{0i} F_{0i} + F_{0i} \partial_i A_0$, so that the Hamiltonian density is:

$$\begin{aligned} \mathcal{H}(x) &= \Pi^\mu(x) \partial_0 A_\mu(x) - \mathcal{L}(x) \\ &= -\Pi_0(x) \partial_0 A_0(x) + F_{0i}(x) \partial_0 A_i(x) - \frac{1}{2} F_{0i}(x) F_{0i}(x) + \frac{1}{4} F_{ij}(x) F_{ij}(x) \\ &= \frac{1}{2} F_{0i}(x) F_{0i}(x) + F_{0i}(x) \partial_i A_0(x) - \Pi_0(x) \partial_0 A_0(x) + \frac{1}{4} F_{ij}(x) F_{ij}(x) \end{aligned}$$

where the Lagrangian density is the integrand in Equation (B.36) and summation on repeated indices is implicit. Then, by replacing F_{0i} by $-E_i$ and $\frac{1}{4} F_{ij} F_{ij}$ by $\frac{1}{2} B_i B_i$ (summation implied), the canonical Hamiltonian is:

$$H_c = \sum_{i=1,2,3} \int_{\mathbb{R}^3} d^3x \frac{1}{2} E_i(x) E_i(x) + \frac{1}{2} B_i(x) B_i(x) - E_i(x) \partial_i A_0(x) - \Pi_0(x) \partial_0 A_0(x)$$

We recognize the first two terms as the density of electromagnetic energy $\frac{1}{2}(\vec{E}^2 + \vec{B}^2)$ while the third term can be integrated by part so that:

$$\int_{\mathbb{R}^3} d^3x -E_i \partial_i A_0 = - \int_{\mathbb{R}^3} d^3x \partial_i (E_i A_0) + \int_{\mathbb{R}^3} d^3x A_0 \partial_i E_i$$

Under appropriate physical assumptions – e.g. that the fields vanish at infinity – the first term on the right-hand side vanishes and we finally get, for the total Hamiltonian:

$$H_T = \int_{\mathbb{R}^3} d^3x \frac{1}{2} \left(\vec{E}^2(x) + \vec{B}^2(x) \right) + A_0(x) \operatorname{div}(\vec{E})(x) + u(x)\phi_x$$

We explicitly wrote $-\Pi_0(x)$ under the constraint symbol ϕ_x and $\partial_0 A_0(x)$ as a smooth parameter of time $u(x)$ for each position x because in the total Hamiltonian we let the coefficient $\partial_0 A_0$ multiplying the constraint to be arbitrary. As of now this is a formal parameter to simplify the expressions we manipulate and we will come back later to its meaning.

Let us now construct the persistence equations associated to the set of primary (also first-stage) constraints $\phi_x = \phi_x^{(1)}$. The Poisson bracket between the canonical variables is:

$$\{A_\mu(x), \Pi_\nu(x')\} = \eta_{\mu\nu} \delta(x - x') \quad (\text{B.38})$$

where on the right-hand side we have the Minkowski metric for the discrete index, and a Dirac delta for the continuous index. We moreover set that the Poisson bracket $\{\Pi_\nu(x), A_\mu(x')\}$ is $-\eta_{\mu\nu} \delta(x - x')$. Since $\phi_x = -\Pi_0(x)$, the only variable it can interact with via the Poisson bracket (B.38) is $A_0(x)$. There are a priori two terms in A_0 : one multiplying $\operatorname{div}(\vec{E})$ while the other could be in $u(x)$ when $u(x) = \partial_0 A_0(x)$. But in the latter case, we have:

$$\{\phi_x, u(x')\phi_{x'}\} = \{\phi_x, u(x')\} \underbrace{\phi_{x'}}_{\approx 0} + u(x') \underbrace{\{\phi_x, \phi_{x'}\}}_{= 0} \approx 0$$

Hence, the Poisson bracket of ϕ_x with $u(x')\phi_{x'}$ is zero, whatever value $u(x')$ takes (be it $\partial_0 A_0$ or totally free). The persistence equations of ϕ_x are then given as follows:

$$\{\phi_x, \mathcal{H}(x')\} \approx \left\{ -\Pi_0(x), A_0(x') \operatorname{div}(\vec{E})(x') \right\} = \eta_{00} \delta(x - x') \operatorname{div}(\vec{E})(x') = -\operatorname{div}(\vec{E})(x)$$

We then have an infinite set of new constraints, which are the second-stage constraints:

$$\phi_x^{(2)} = -\operatorname{div}(\vec{E})(x)$$

Mathematically, we should rigorously write $\phi_x^{(2)} = \sum_{i=1,2,3} \partial_i \Pi_i(x)$ but physically, we would prefer to write $\operatorname{div}(\vec{E}) = 0$, since this constraint corresponds to the *Gauss law* (1.38) (in the absence of matter).

Now, we will see that the persistence equation for the second-stage constraints do not give any new constraints. First, the bracket of $\Pi_i = -E_i$ with \mathcal{H} can be computed thanks for the following observation: the only term in the Hamiltonian density with which it will interact is $\frac{1}{4} F_{ij} F_{ij}$ (summation implied). Indeed, this quadratic term can be written as $\frac{1}{2} \partial_i (A_j) F_{ij}$ which yields, upon integration by part:

$$\int_{\mathbb{R}^3} d^3x \frac{1}{4} F_{ij} F_{ij} = \int_{\mathbb{R}^3} d^3x \frac{1}{2} \partial_i A_j F_{ij} = \int_{\mathbb{R}^3} d^3x \frac{1}{2} \partial_i (A_j F_{ij}) - \int_{\mathbb{R}^3} d^3x \frac{1}{2} A_j \partial_i F_{ij} \quad (\text{B.39})$$

where as usual the summation and the presence of x is implied. The first term on the right-hand side vanishes under appropriate physical assumptions. Now, since $\operatorname{div}(\vec{E}) = \sum_i \partial_i E_i$ the bracket of (minus) $\phi_x^{(2)}$ with $\mathcal{H}(x')$ yields (summation implied):

$$\left\{ \partial_i E_i(x), \mathcal{H}(x') \right\} = \partial_i \left(\left\{ E_i(x), \mathcal{H}(x') \right\} \right) - \left\{ E_i(x), \partial_i \mathcal{H}(x') \right\} \quad (\text{B.40})$$

The first term on the right-hand side being exact, it will disappear upon integration. However in $\partial_i \mathcal{H}$ only $\partial_i(\frac{1}{4}F_{jk}F_{jk}) - \text{i.e.} -\frac{1}{2}\partial_i(A_k\partial_j F_{jk})(x')$ using Equation (B.39) – will interact with E_i in the Poisson bracket. So the last term of the right-hand side of Equation (B.40) reads:

$$-\{E_i(x), \partial_i \mathcal{H}(x')\} = \left\{E_i(x), \frac{1}{2}\partial_i(A_k\partial_j F_{jk})(x')\right\} = \left\{E_i(x), \frac{1}{2}A_k\partial_i(\partial_j F_{jk})(x')\right\} = \frac{1}{2}\partial_i\partial_j F_{ji}(x)$$

The summation over i, j and k is implicit through the computation. There is no minus sign in the last term because we use $\{E_i(x), A_k(x')\} = \{-\Pi_i(x), A_k(x')\} = \eta_{ki}\delta(x-x')$. Since F_{ij} is antisymmetric in i, j while the double derivative is symmetric, this term vanishes. Thus, we have the following automatically satisfied equation:

$$\{\phi_x^{(2)}, \mathcal{H}(x')\} = 0$$

There are no third-stage constraints and the secondary constraints are uniquely the second-stage ones.

We immediately see that they are first class because they only involve the electric field. The persistence equations do not enforce any restrictions on $u(x)$, and additionally the scalar potential $-A_0(x)$ can be somehow considered as a free parameter $v(x)$ as well, but in that case we obtain the extended hamiltonian:

$$H_E = \int_{\mathbb{R}^3} d^3x \frac{1}{2} \left(\vec{E}^2(x) + \vec{B}^2(x) \right) + u(x)\phi_x^{(1)} + v(x)\phi_x^{(2)}$$

The extended Hamiltonian differs from the total Hamiltonian in that v is considered a free parameter in the former while it supposedly coincides with $-A_0$ in the latter. This is the most general form the Hamiltonian can take, but then the solutions of the equations of motions it generates are much more wider than those generated by the total Hamiltonian (where $v = A_0$). The gauge transformations are generated by the constraints $\phi_x^{(1)}$ and $\phi_x^{(2)}$, that is to say:

$$\delta A_0(x) = \{u(x)\phi_x^{(1)} + v(x)\phi_x^{(2)}, A_0(x)\} = -u(x) \quad (\text{B.41})$$

$$\delta A_i(x) = \{u(x)\phi_x^{(1)} + v(x)\phi_x^{(2)}, A_i(x)\} = \partial_i v(x) \quad (\text{B.42})$$

The sign in the first line comes from the fact that $\{-\Pi_0(x), A_0(x')\} = \eta_{00}\delta(x-x')$ while on the second line we implicitly proceeded to an integration by part and developed the Poisson bracket: the only valuable part in $\{v(x)(-\partial_j E_j(x)), A_i(x')\}$ is $\partial_j v(x)\{E_j(x), A_i(x')\} = \partial_j v(x)\eta_{ij}\delta(x-x')$.

Notice that at first sight Equations (B.41) and (B.42) do not precisely correspond to what we know about gauge transformations in four-dimensional electromagnetism, since in the latter case $\delta A_\mu = \partial_\mu \lambda$ where λ is the gauge parameter. If we follow our intuition it would mean that $u = -\partial_0 v$, and this is precisely the case since we originally have used u to denote $\partial_0 A_0$, while we later used v to denote $-A_0$. So, implicitly, we always had $u = -\partial_0 v$ without saying. So when $v = -A_0$ and $u = \partial_0 A_0$ we indeed have $\delta A_\mu = \partial_\mu \lambda$, with $\lambda = -A_0$. However if one consider u and v as free parameters the relationship $u = -\partial_0 v$ does not necessarily holds, but Matschull explains that this is not too worrisome at the bottom of page 24 in [Matschull, 1996], so we obtain the same gauge transformations.

Eventually, notice that an adapted set of gauge fixing conditions single out one transversal to the gauge orbits (each set of gauge fixing conditions defines a different transversal) which, in geometric terms, represents the reduced phase space Σ_{ph} inside the constraint surface Σ . Since we have two first-class constraints $\phi_x = -\Pi_0(x)$ and $\phi_x^{(2)} = -\text{div}(\vec{E})(x) = \sum_{i=1,2,3} \partial_i \Pi_i(x)$, we need two gauge fixing conditions so that the four functions form a pure second-class system. For example, in subsection 9.3.1 in [Rothe and Rothe, 2010], the Coulomb gauge fixing condition is enforced via the addition of the following two gauge fixing conditions:

$$C_{1,x} = A_0(x) \quad \text{and} \quad C_{2,x} = \text{div}(\vec{A})(x) = \sum_{i=1,2,3} \partial_i \Pi_i(x)$$

The second one being the proper Coulomb gauge condition. These gauge fixing conditions define an (infinite dimensional) submanifold N in the constraint surface Σ , which is transversal to the gauge orbits and would be a representent of the reduced phase space Σ_{ph} .

The Poisson brackets are $\{\phi_x, C_{1,x'}\} = -\delta(x-x')$ and $\{\phi_x^{(2)}, C_{2,x'}\} = -\Delta_x \delta(x-x')$; the other Poisson brackets being zero. Then, the 2×2 -matrix of Poisson bracket is invertible and defines a Dirac bracket on the transversal N , that we denote with a small asterisk (see subsection 9.3.1 in [Rothe and Rothe, 2010] for the computation):

$$\{A_i(x), \Pi_j(x')\}^* = \left(\eta_{ij} - \frac{\partial_i \partial_j}{\Delta_x} \right) \delta(x-x')$$

The other brackets being zero – even $\{A_0(x), \Pi_0(x')\}^*$. Then, canonical quantization of this second-class system could be performed straightforwardly by adding $i\hbar$ on the right-hand side, and changing the Poisson bracket into the commutator of operators on the left-hand side. However, as shown in these [lecture notes](#), the Coulomb gauge has the disadvantage of breaking locality – the fact the Lagrangian or Hamiltonian can be expressed with only a finite sum of the fields and their derivatives – because the Green function of the Laplacian is not local. Moreover, the Coulomb gauge gives rise to so called Gribov ambiguities: the fact that the transversal to the gauge orbits crosses them several times, so there is no unique gauge fixing in some configuration (see p. 576 of [Itzykson and Zuber, 2005]). See [Gomes and Butterfield, 2022] for a philosophical interpretation of the Coulomb gauge; more generally, Gomes' work is worth looking for anyone interested in interpretations of gauge theories [Gomes, 2021].

This closes the treatment of electromagnetism in the canonical Hamiltonian formalism, but other field theories than electromagnetism have been also treated along these lines, among which many examples were done by Alberto Escalante (see his work on ArXiv).

References

- [Abraham and Marsden, 1987] Abraham, R. and Marsden, J. E. (1987). *Foundations of Mechanics*. Addison-Wesley Publishing Company, Inc., Redwood City, CA, 2 edition.
- [Agrachev et al., 2019] Agrachev, A., Barilari, D., and Boscain, U. (2019). *A Comprehensive Introduction to Sub-Riemannian Geometry*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge.
- [Ali and Engliš, 2005] Ali, S. T. and Engliš, M. (2005). Quantization methods: A guide for physicists and analysts. *Rev. Math. Phys.*, 17(04):391–490.
- [Anderson and Bergmann, 1951] Anderson, J. L. and Bergmann, P. G. (1951). Constraints in covariant field theories. *Phys. Rev.*, 83(5):1018–1025.
- [Baez and Muniain, 1994] Baez, J. and Muniain, J. P. (1994). *Gauge Fields, Knots and Gravity*, volume 4 of *Series on Knots and Everything*. World Scientific, Singapore.
- [Bamberg and Sternberg, 1988a] Bamberg, P. and Sternberg, S. (1988a). *A Course in Mathematics for Students of Physics*, volume 2. Cambridge University Press, Cambridge.
- [Bamberg and Sternberg, 1988b] Bamberg, P. and Sternberg, S. (1988b). *A Course in Mathematics for Students of Physics*, volume 1. Cambridge University Press, Cambridge.
- [Barnich and Del Monte, 2018] Barnich, G. and Del Monte, F. (2018). Introduction to classical Gauge Field Theory and to Batalin-Vilkovisky quantization.
- [Batalin and Vilkovisky, 1984] Batalin, I. A. and Vilkovisky, G. A. (1984). Closure of the gauge algebra, generalized lie equations and Feynman rules. *Nucl. Phys. B*, 234(1):106–124.
- [Batalin and Vilkovisky, 1985] Batalin, I. A. and Vilkovisky, G. A. (1985). Existence theorem for gauge algebra. *J. Math. Phys.*, 26(1):172–184.
- [Batchelor, 1979] Batchelor, M. (1979). The structure of supermanifolds. *Trans. Amer. Math. Soc.*, 253:329–338.
- [Batchelor, 1980] Batchelor, M. (1980). Two approaches to supermanifolds. *Trans. Amer. Math. Soc.*, 258(1):257–270.
- [Batlle et al., 1986] Batlle, C., Gomis, J., Pons, J. M., and Roman-Roy, N. (1986). Equivalence between the Lagrangian and Hamiltonian formalism for constrained systems. *J. Math. Phys.*, 27(12):2953–2962.
- [Berezin and Leites, 1975] Berezin, F. and Leites, D. (1975). Supermanifolds. *Dokl. Akad. Nauk SSSR*, 224(3):505–508.
- [Bergmann, 1949] Bergmann, P. G. (1949). Non-linear field theories. *Phys. Rev.*, 75(4):680–685.
- [Bergmann and Brunings, 1949] Bergmann, P. G. and Brunings, J. H. M. (1949). Non-linear field theories II. Canonical equations and Quantization. *Rev. Mod. Phys.*, 21(3):480–487.
- [Bergmann and Goldberg, 1955] Bergmann, P. G. and Goldberg, I. (1955). Dirac bracket transformations in phase space. *Phys. Rev.*, 98(2):531–538.
- [Blacker et al., 2022] Blacker, C., Miti, A. M., and Ryvkin, L. (2022). Reduction of L-infinity-algebras of observables on multisymplectic manifolds.

- [Blaschke and Gieres, 2021] Blaschke, D. N. and Gieres, F. (2021). On the canonical formulation of gauge field theories and Poincaré transformations. *Nucl. Phys. B*, 965:115366.
- [Blohmann, 2023] Blohmann, C. (2023). Lagrangian Field Theory:Diffeology, Variational Cohomology, Multisymplectic Geometry.
- [Brown, 2022] Brown, J. D. (2022). Singular Lagrangians, Constrained Hamiltonian Systems and Gauge Invariance: An Example of the Dirac–Bergmann Algorithm. *Universe*, 8(3):171.
- [Browning and McMullan, 1987] Browning, A. D. and McMullan, D. (1987). The Batalin, Fradkin, and Vilkovisky formalism for higher-order theories. *J. Math. Phys.*, 28(2):438–444.
- [Bursztyn et al., 2019] Bursztyn, H., Lima, H., and Meinrenken, E. (2019). Splitting theorems for Poisson and related structures. *J. Reine Angew. Math.*, 2019(754):281–312.
- [Cabo, 1986] Cabo, A. (1986). On Dirac’s conjecture for systems having only first-class constraints. *J. Phys. A: Math. Gen.*, 19(5):629–638.
- [Cabo and Louis-Martinez, 1990] Cabo, A. and Louis-Martinez, D. (1990). On Dirac’s conjecture for Hamiltonian systems with first- and second-class constraints. *Phys. Rev. D*, 42(8):2726–2735.
- [Calvo et al., 2010] Calvo, I., Falceto, F., and Zambon, M. (2010). Deformation of Dirac structures along isotropic subbundles. *Rep. Math. Phys.*, 65(2):259–269.
- [Candel and Conlon, 2000] Candel, A. and Conlon, L. (2000). *Foliations I*. Number 23 in Graduate Studies in Mathematics. American Mathematical Society, Providence, RI.
- [Carosso, 2022] Carosso, A. (2022). Quantization: History and problems. *Stud. Hist. Philos. Sci.*, 96:35–50.
- [Cattaneo et al., 2006] Cattaneo, A., Fiorenza, D., and Longoni, R. (2006). Graded Poisson algebras. In Françoise, J.-P., Naber, G. L., and Tsun, T. S., editors, *Encyclopedia of Mathematical Physics*, pages 560–567. Academic Press, Oxford.
- [Cattaneo and Schätz, 2011] Cattaneo, A. and Schätz, F. (2011). Introduction to supergeometry. *Rev. Math. Phys.*, 23(06):669–690.
- [Cattaneo and Zambon, 2009] Cattaneo, A. and Zambon, M. (2009). Coisotropic embeddings in Poisson manifolds. *Trans. Amer. Math. Soc.*, 361(7):3721–3746.
- [Crainic and Fernandes, 2004] Crainic, M. and Fernandes, R. L. (2004). Integrability of Poisson brackets. *J. Differ. Geom.*, 66(1):71–137.
- [Crainic et al., 2021] Crainic, M., Fernandes, R. L., and Marcut, I. (2021). *Lectures on Poisson Geometry*. Number 217 in Graduate Studies in Mathematics. American Mathematical Society, Providence, RI.
- [Curtright et al., 2013] Curtright, T. L., Fairlie, D. B., and Zachos, C. K. (2013). Historical survey. In *A Concise Treatise on Quantum Mechanics in Phase Space*, pages 1–28. World Scientific.
- [DeWitt, 1984] DeWitt, B. (1984). *Supermanifolds*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, New York.
- [Dirac, 1950] Dirac, P. A. M. (1950). Generalized Hamiltonian dynamics. *Canad. J. Math.*, 2:129–148.

- [Dirac, 1951] Dirac, P. A. M. (1951). The Hamiltonian Form of Field Dynamics. *Canad. J. Math.*, 3:1–23.
- [Dirac, 1958] Dirac, P. A. M. (1958). Generalized Hamiltonian dynamics. *Proc. R. Soc. A: Math. Phys. Eng. Sci.*, 246(1246):326–332.
- [Dirac, 1964] Dirac, P. A. M. (1964). *Lectures on Quantum Mechanics*. Number 2 in Monographs Series. Belfer Graduate School of Science, Yeshiva University, New York.
- [Dubois-Violette, 1987] Dubois-Violette, M. (1987). Systèmes dynamiques contraints : l’approche homologique. *Ann. Inst. Fourier*, 37(4):45–57.
- [Dufour and Zung, 2005] Dufour, J.-P. and Zung, N. T. (2005). *Poisson Structures and Their Normal Forms*. Number 242 in Progress in Mathematics. Birkhäuser, Basel.
- [Duncan and Duncan, 2012] Duncan, A. and Duncan, A. (2012). *The Conceptual Framework of Quantum Field Theory*. Oxford University Press, Oxford, New York.
- [Earman, 2003] Earman, J. (2003). Tracking down gauge: An ode to the constrained Hamiltonian formalism. In Castellani, E. and Brading, K., editors, *Symmetries in Physics: Philosophical Reflections*, pages 140–162. Cambridge University Press, Cambridge.
- [Faddeev and Jackiw, 1988] Faddeev, L. and Jackiw, R. (1988). Hamiltonian reduction of unconstrained and constrained systems. *Phys. Rev. Lett.*, 60(17):1692–1694.
- [Fairon, 2017] Fairon, M. (2017). Introduction to graded geometry. *Eur. J. Math.*, 3(2):208–222.
- [Falceto and Zambon, 2008] Falceto, F. and Zambon, M. (2008). An Extension of the Marsden–Ratiu Reduction for Poisson Manifolds. *Lett. Math. Phys.*, 85(2):203–219.
- [Fernandes, 2005] Fernandes, R. L. (2005). A note on proper Poisson actions.
- [Figueroa-O’Farrill and Kimura, 1991a] Figueroa-O’Farrill, J. M. and Kimura, T. (1991a). Geometric BRST quantization, I: Prequantization. *Commun. Math. Phys.*, 136(2):209–229.
- [Figueroa-O’Farrill and Kimura, 1991b] Figueroa-O’Farrill, J. M. and Kimura, T. (1991b). Homological approach to symplectic reduction.
- [Fisch et al., 1989] Fisch, J., Henneaux, M., Stasheff, J., and Teitelboim, C. (1989). Existence, uniqueness and cohomology of the classical BRST charge with ghosts of ghosts. *Commun. Math. Phys.*, 120(3):379–407.
- [Fuster et al., 2005] Fuster, A., Henneaux, M., and Maas, A. (2005). BRST-antifield quantization: A short review. *Int. J. Geom. Methods Mod. Phys.*, 02(05):939–963.
- [García and Pons, 1997] García, J. A. and Pons, J. M. (1997). Equivalence of Faddeev–Jackiw and Dirac Approaches for Gauge Theories. *Int. J. Mod. Phys. A*, 12(02):451–464.
- [García and Pons, 2000] García, J. A. and Pons, J. M. (2000). Rigid and gauge noether symmetries for constrained systems. *Int. J. Mod. Phys. A*, 15(29):4681–4721.
- [Gitman and Tyutin, 1990] Gitman, D. M. and Tyutin, I. V. (1990). *Quantization of Fields with Constraints*. Springer Series in Nuclear and Particle Physics. Springer-Verlag, Berlin, Heidelberg.
- [Gomes, 2021] Gomes, H. (2021). *Why Gauge? Conceptual Aspects of Gauge Theories*. PhD thesis, University of Cambridge, arXiv:2203.05339.

- [Gomes and Butterfield, 2022] Gomes, H. and Butterfield, J. (2022). How to Choose a Gauge? The Case of Hamiltonian Electromagnetism. *Erkenntnis*.
- [Gómez, 2016] Gómez, G. L. (2016). The Elegance of Cohomological Methods.
- [Gomis et al., 1995] Gomis, J., París, J., and Samuel, S. (1995). Antibracket, antifields and gauge-theory quantization. *Phys. Rep.*, 259(1):1–145.
- [Gotay, 1982] Gotay, M. J. (1982). On coisotropic imbeddings of presymplectic manifolds. *Proc. Amer. Math. Soc.*, 84(1):111–114.
- [Gotay, 1983] Gotay, M. J. (1983). On the validity of Dirac’s conjecture regarding first-class secondary constraints. *J. Phys. A: Math. Gen.*, 16(5):L141–L145.
- [Gotay, 1999] Gotay, M. J. (1999). On the Groenewold–Van Hove problem for R^{2n} . *J. Math. Phys.*, 40(4):2107–2116.
- [Gotay et al., 1996] Gotay, M. J., Grundling, H. B., and Tuynman, G. M. (1996). Obstruction results in quantization theory. *J. Nonlinear Sci.*, 6(5):469–498.
- [Gotay and Nester, 1984] Gotay, M. J. and Nester, J. M. (1984). Apartheid in the Dirac theory of constraints. *J. Phys. A: Math. Gen.*, 17(15):3063.
- [Gotay et al., 1978] Gotay, M. J., Nester, J. M., and Hinds, G. (1978). Presymplectic manifolds and the Dirac–Bergmann theory of constraints. *J. Math. Phys.*, 19(11):2388–2399.
- [Gràcia and Pons, 1988] Gràcia, X. and Pons, J. M. (1988). Gauge generators, Dirac’s conjecture, and degrees of freedom for constrained systems. *Ann. Phys.*, 187(2):355–368.
- [Groenewold, 1946] Groenewold, H. J. (1946). On the principles of elementary quantum mechanics. *Physica*, 12(7):405–460.
- [Gualtieri, 2003] Gualtieri, M. (2003). *Generalized Complex Geometry*. PhD thesis, University of Oxford, arXiv:math/0401221.
- [Henneaux, 1990] Henneaux, M. (1990). Lectures on the antifield-BRST formalism for gauge theories. *Nucl. Phys. B - Proc. Suppl.*, 18(1):47–105.
- [Henneaux and Teitelboim, 1988] Henneaux, M. and Teitelboim, C. (1988). BRST cohomology in classical mechanics. *Commun. Math. Phys.*, 115(2):213–230.
- [Henneaux and Teitelboim, 1992] Henneaux, M. and Teitelboim, C. (1992). *Quantization of Gauge Systems*. Princeton University Press, Princeton.
- [Hermann, 1962] Hermann, R. (1962). The differential geometry of foliations, II. *J. Math. Mech.*, 11(2):303–315.
- [Ikeda and Strobl, 2019] Ikeda, N. and Strobl, T. (2019). On the relation of Lie algebroids to constrained systems and their BV/BFV formulation. *Ann. Inst. Henri Poincaré A*, 20(2):527–541.
- [Itzykson and Zuber, 2005] Itzykson, C. and Zuber, J.-B. (2005). *Quantum Field Theory*. Dover, Mineola.
- [Kanatchikov, 2001] Kanatchikov, I. V. (2001). Precanonical Quantum Gravity: Quantization without the space–time decomposition. *Int. J. Theor. Phys.*, 40(6):1121–1149.

- [Karshon et al., 2022] Karshon, Y., Miyamoto, D., and Watts, J. (2022). Diffeological submanifolds and their friends.
- [Kimura, 1993] Kimura, T. (1993). Generalized classical BRST cohomology and reduction of Poisson manifolds. *Commun. Math. Phys.*, 151(1):155–182.
- [Kolar et al., 1993] Kolar, I., Slovák, J., and Michor, P. W. (1993). *Natural Operations in Differential Geometry*. Springer, Berlin, Heidelberg.
- [Kosmann-Schwarzbach and Rubtsov, 2010] Kosmann-Schwarzbach, Y. and Rubtsov, V. (2010). Compatible structures on Lie algebroids and Monge–Ampère operators. *Acta. Appl. Math.*, 109(1):101–135.
- [Kostant, 1977] Kostant, B. (1977). Graded manifolds, graded Lie theory, and prequantization. In Bleuler, K. and Reetz, A., editors, *Differential Geometrical Methods in Mathematical Physics*, number 570 in Lecture Notes in Mathematics, pages 177–306. Springer, Berlin, Heidelberg.
- [Kotov and Salnikov, 2021] Kotov, A. and Salnikov, V. (2021). The category of \mathbb{Z} -graded manifolds: What happens if you do not stay positive.
- [Kundt, 1966] Kundt, W. (1966). Canonical quantization of gauge invariant field theories. In Flügge, S., Stierstadt, K., and Kundt, W., editors, *Springer Tracts in Modern Physics, Volume 40*, number 40 in Springer Tracts in Modern Physics, pages 107–168. Springer, Berlin, Heidelberg.
- [Künzle, 1969] Künzle, H. P. (1969). Degenerate lagrangean systems. *Ann. Inst. Henri Poincaré A*, 11(4):393–414.
- [Laurent-Gengoux et al., 2020] Laurent-Gengoux, C., Lavau, S., and Strobl, T. (2020). The universal Lie infinity-algebroid of a singular foliation. *Doc. Math.*, 25:1571–1652.
- [Laurent-Gengoux et al., 2013] Laurent-Gengoux, C., Pichereau, A., and Vanhaecke, P. (2013). *Poisson Structures*. Grundlehren Der Mathematischen Wissenschaften. Springer-Verlag, Berlin, Heidelberg.
- [Lee, 2003] Lee, J. M. (2003). *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer-Verlag, New York.
- [Lee, 2009] Lee, J. M. (2009). *Manifolds and Differential Geometry*. Number 107 in Graduate Studies in Mathematics. American Mathematical Society, Providence, RI.
- [Leites, 1980] Leites, D. (1980). Introduction to the theory of Supermanifolds. *Russ. Math. Surv.*, 35(1):1–64.
- [Lepage and Godeaux, 1951] Lepage, T. H. J. and Godeaux, L. (1951). Sur certaines représentations unitaires d’un groupe infini de transformations. Mémoire présenté par Léon Van Hove. *Bull. Acad. R. Belg.*, 37(1):679–684.
- [Mac Lane and Moerdijk, 1994] Mac Lane, S. and Moerdijk, I. (1994). *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext. Springer-Verlag, New York.
- [Mackenzie, 1992] Mackenzie, K. (1992). Double Lie algebroids and second-order geometry, I. *Adv. Math.*, 94(2):180–239.

- [Mariño, 2021] Mariño, M. (2021). The phase space formulation of Quantum Mechanics. In *Advanced Topics in Quantum Mechanics*, pages 114–158. Cambridge University Press, Cambridge.
- [Marle, 1983] Marle, C.-M. (1983). Sous-variétés de rang constant d’une variété symplectique. In *IIIe Rencontre de Géométrie Du Schnepfenried, 10-15 Mai 1982*, volume 1 of *Astérisque*, No 107-108.
- [Matschull, 1996] Matschull, H.-J. (1996). Dirac’s canonical quantization programme.
- [McMullan, 1992] McMullan, D. (1992). Classical states and the BRST charge. *Commun. Math. Phys.*, 149(1):161–174.
- [Mehta, 2006] Mehta, R. A. (2006). *Supergroupoids, Double Structures, and Equivariant Cohomology*. PhD thesis, University of California, Berkeley.
- [Miyamoto, 2023] Miyamoto, D. (2023). Singular foliations through diffeology.
- [Mnev, 2017] Mnev, P. (2017). Lectures on Batalin-Vilkovisky formalism and its applications in topological quantum field theory.
- [Moerdijk and Mrcun, 2003] Moerdijk, I. and Mrcun, J. (2003). *Introduction to Foliations and Lie Groupoids*. Number 91 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge.
- [Molino, 1988] Molino, P. (1988). *Riemannian Foliations*. Number 73 in Progress in Mathematics. Birkhäuser, Boston, MA.
- [Nestruev, 2003] Nestruev, J. (2003). *Smooth Manifolds and Observables*. Graduate Texts in Mathematics. Springer-Verlag, New York.
- [Nirov and Razumov, 1993] Nirov, K. S. and Razumov, A. V. (1993). Equivalence between Lagrangian and Hamiltonian BRST formalisms. *J. Math. Phys.*, 34(9):3933–3953.
- [Pitts, 2022] Pitts, J. B. (2022). First-Class Constraints, Gauge Transformations, de-Ockhamization, and Triviality: Replies to Critics, Or, How (Not) to Get a Gauge Transformation from a Second-Class Primary Constraint.
- [Pons, 2005] Pons, J. M. (2005). On Dirac’s incomplete analysis of gauge transformations. *Stud. Hist. Philos. Sci. B*, 36(3):491–518.
- [Pooley and Wallace, 2022] Pooley, O. and Wallace, D. (2022). First-class constraints generate gauge transformations in electromagnetism (reply to Pitts).
- [Pradines, 1985] Pradines, J. (1985). How to define the differentiable graph of a singular foliation. *Cah. Topol. Géom. Différ. Catég.*, 26(4):339–380.
- [Rogers, 1980] Rogers, A. (1980). A global theory of supermanifolds. *J. Math. Phys.*, 21(6):1352–1365.
- [Rosenfeld, 1932] Rosenfeld, L. (1932). La théorie quantique des champs. *Ann. Inst. Henri Poincaré*, 2(1):25–91.
- [Rothe and Rothe, 2003] Rothe, H. J. and Rothe, K. D. (2003). Lagrange versus symplectic algorithm for constrained systems. *J. Phys. A: Math. Gen.*, 36(6):1671.

- [Rothe and Rothe, 2010] Rothe, H. J. and Rothe, K. D. (2010). *Classical and Quantum Dynamics of Constrained Hamiltonian Systems*, volume 81 of *World Scientific Lecture Notes in Physics*. World Scientific, Singapore.
- [Salisbury, 2012] Salisbury, D. C. (2012). Peter Bergmann and the Invention of Constrained Hamiltonian Dynamics. In Lehner, C., Renn, J., and Schemmel, M., editors, *Einstein and the Changing Worldviews of Physics*, Einstein Studies, pages 247–257. Birkhäuser, Boston.
- [Schätz, 2008] Schätz, F. (2008). BFV-complex and higher homotopy structures. *Commun. Math. Phys.*, 286(2):399.
- [Schätz, 2009] Schätz, F. (2009). *Coisotropic Submanifolds and the BFV-Complex*. PhD thesis, Universität Zürich, Zurich.
- [Schmidt, 1999] Schmidt, H.-J. (1999). The Classical solutions of two-dimensional gravity. *Gen. Relativ. Gravit.*, 31(8):1187–1210.
- [Schroeck, 1996] Schroeck, F. E. (1996). *Quantum Mechanics on Phase Space*. Springer Netherlands, Dordrecht.
- [Sears, 1963] Sears, F. W. (1963). A simplified simplification of Carathéodory’s treatment of Thermodynamics. *Am. J. Phys.*, 31(10):747–752.
- [Sears, 1966] Sears, F. W. (1966). Modified form of Carathéodory’s second axiom. *Am. J. Phys.*, 34(8):665–666.
- [Shanmugadhasan, 1963] Shanmugadhasan, S. (1963). Generalized canonical formalism for degenerate dynamical systems. *Math. Proc. Camb. Philos. Soc.*, 59(4):743–757.
- [Shanmugadhasan, 1973] Shanmugadhasan, S. (1973). Canonical formalism for degenerate Lagrangians. *J. Math. Phys.*, 14(6):677–687.
- [Shewell, 1959] Shewell, J. R. (1959). On the Formation of Quantum-Mechanical Operators. *Am. J. Phys.*, 27(1):16–21.
- [Śniatycki, 1974] Śniatycki, J. (1974). Dirac brackets in geometric dynamics. *Ann. Inst. Henri Poincaré A*, 20(4):365–372.
- [Stasheff, 1991] Stasheff, J. (1991). Homological (ghost) approach to constrained Hamiltonian systems.
- [Stasheff, 1997] Stasheff, J. (1997). Homological reduction of constrained Poisson algebras. *J. Differ. Geom.*, 45(1):221–240.
- [Stefan, 1974] Stefan, P. (1974). Accessible sets, orbits, and foliations with singularities. *Proc. London Math. Soc.*, s3-29(4):699–713.
- [Stefano, 1983] Stefano, R. D. (1983). Modification of Dirac’s method of Hamiltonian analysis for constrained systems. *Phys. Rev. D*, 27(8):1752–1765.
- [Sudarshan and Mukunda, 1974] Sudarshan, E. C. G. and Mukunda, N. (1974). *Classical Dynamics: A Modern Perspective*. Wiley, New York.
- [Sundermeyer, 1982] Sundermeyer, K. (1982). *Constrained Dynamics*. Number 169 in Lecture Notes in Physics. Springer-Verlag, Berlin/Heidelberg.

- [Tătar and Tătar, 1994] Tătar, L. and Tătar, R. (1994). Algebraic structure of the BRST symmetry. *J. Math. Phys.*, 35(11):5862–5868.
- [Tulczyjew, 1977] Tulczyjew, W. M. (1977). The Legendre transformation. *Ann. Inst. Henri Poincaré A*, 27(1):101–114.
- [Vaintrob, 1997] Vaintrob, A. (1997). Lie algebroids and homological vector fields. *Russ. Math. Surv.*, 52(2):428–429.
- [Vaisman, 1994] Vaisman, I. (1994). *Lectures on the Geometry of Poisson Manifolds*. Progress in Mathematics. Birkhäuser Basel.
- [Voronov, 2010] Voronov, T. (2010). Q-manifolds and higher analogs of Lie algebroids. *AIP Conference Proceedings*, 1307(1):191–202.
- [Warner, 1983] Warner, F. W. (1983). *Foundations of Differentiable Manifolds and Lie Groups*. Graduate Texts in Mathematics. Springer-Verlag, New York.
- [Weinstein, 1988] Weinstein, A. (1988). Coisotropic calculus and Poisson groupoids. *J. Math. Soc. Jpn.*, 40(4):705–727.
- [Woodhouse, 1997] Woodhouse, N. M. J. (1997). *Geometric Quantization*. Oxford Mathematical Monographs. Oxford University Press, Oxford, New York, second edition edition.
- [Yoshimura and Marsden, 2006a] Yoshimura, H. and Marsden, J. E. (2006a). Dirac structures in Lagrangian mechanics Part I: Implicit Lagrangian systems. *J. Geom. Phys.*, 57(1):133–156.
- [Yoshimura and Marsden, 2006b] Yoshimura, H. and Marsden, J. E. (2006b). Dirac structures in Lagrangian mechanics Part II: Variational structures. *J. Geom. Phys.*, 57(1):209–250.
- [Zambon, 2011] Zambon, M. (2011). Submanifolds in Poisson geometry: A survey. In Ebeling, W., Hulek, K., and Smoczyk, K., editors, *Complex and Differential Geometry*, Springer Proceedings in Mathematics, pages 403–420, Berlin, Heidelberg. Springer.
- [Zemansky, 1966] Zemansky, M. W. (1966). Kelvin and Caratheodory—A Reconciliation. *Am. J. Phys.*, 34(10):914–920.
- [Zia et al., 2009] Zia, R. K. P., Redish, E. F., and McKay, S. R. (2009). Making sense of the Legendre transform. *Am. J. Phys.*, 77(7):614–622.