Differential Geometry

A mini-course

Marcelo Epstein
University of Calgary

Contents

0.1 Preface

In subtitling this booklet a "mini course", I am reminded of the old joke known as "We sell fresh fish here". But the connection between the subtitle and the proverbial sign at the market is more one of contrast than of analogy. Indeed, in the latter it was the superfluity of the sign what counted, whereas in the former it is its inaccuracy. Both words, "mini" and "course", are patently inaccurate. That this is anything but "mini" should be clear from the number of printed pages, closer to eight dozen than to two. Moreover, this is not a "course" but a haphazard collection of geometrically sounding topics. To make matters worse, I am not a differential geometer but an occasional user of the discipline. Should I then take the sign down and, outdoing the fishmonger, close the fish stand altogether? While considering the proper course of action, I find myself contemplating in despair the frightening scenario of the possible stature of the audience, consisting of active scientists gathered for a workshop on the applications of differential geometry to continuum mechanics. I should take the sign down, close shop and become a potential buyer rather than a seller.

Naturally, nothing of the sort is going to happen and I will do my duty. To enliven somewhat the enterprise, I decided to organize the material in a somewhat non-standard fashion, ascending first rather fast to increasingly complex algebraic and topologic structures without involving differentiability. This approach allows me to talk about fibre bundles and groupoids and their applications in a qualitative manner invoking only topological concepts. This material occupies the first two chapters. The last two chapters are addressed to the survivors of the first two. It is hoped that the motivation will be by then strong enough to endure the details brought about by the differential structure. Much of the material is

recycled from previous work, so that the fish is not even fresh! But I have attempted to make up for this peccadillo by preserving only the essential elements and reordering the material in the manner alluded to above.

Chapter 1

Topological constructs

1.1 Topological spaces

Physical theory has imposed on us the need for the notion of a *continuum*. Indeed, more than two millennia ago, the great magister said that space is "infinitely divisible into parts, themselves infinitely divisible". He also established that space has "three dimensions, by which all bodies are bounded". Whatever else may be wrong in Aristotle, at least when it comes to this basic notion he got it right. Fast-forward to the early twentieth century to find mathematicians (Riesz, Hausdorff, Kuratowski) tackling and solving the following question: What is the minimal structure that a set $\mathcal{S}$ must possess to sustain the notions of *nearness* and *continuity*? The answer is, of course, a *topology*. The key is provided by the identification of certain subsets of $\mathcal{S}$ as being *open*. This identification must satisfy three conditions, which are ultimately abstracted from the common notion of open set in $\mathbb{R}^n$:

(1) The null set $\emptyset$ and the total set $\mathcal{S}$ are open.

(2) Any arbitrary union of open subsets is open.

(3) The intersection of a finite number of open subsets is open.

The set $\mathcal{S}$ is then said to have been endowed with a topology or to be a *topological space*.

A function $f : \mathcal{S} \to \mathcal{T}$ is *continuous* if inverse images of open subsets in $\mathcal{T}$ are open in $\mathcal{S}$. This definition can be shown to be equivalent to the usual definition when particularized to a function $f : \mathbb{R}^m \to \mathbb{R}^n$. But does a topological space look like $\mathbb{R}^n$ in all respects? Not at all. In particular, the second feature of Aristotle's conception of space is missing, namely its fixed local dimensionality. To introduce this notion, we can get inspiration from Cartography: although the surface of the Earth cannot be mapped continuously in one piece onto the plane $\mathbb{R}^2$, it can be so represented in a piecewise fashion. To formalize this general idea, we need a few straightforward definitions.

Two topological spaces, $\mathcal{S}$ and $\mathcal{T}$, are said to be *homeomorphic*, if there exists a continuous bijection $\phi : \mathcal{S} \to \mathcal{T}$ (called a *homeomorphism*) whose inverse is continuous. A *neighbourhood* of a point $p \in \mathcal{S}$ is any open subset $\mathcal{U}_p \subset \mathcal{S}$ such that $p \in \mathcal{U}_p$. A topological space $\mathcal{S}$ is said to be a *Hausdorff space* if for any two points $a, b \in \mathcal{S}$ there exist respective disjoint neighbourhoods, namely: $\mathcal{U}_a \cap \mathcal{U}_b = \emptyset$.

1.2 Topological manifolds

We are now in a position to provide the following definition of a *topological manifold*:[1]

Definition 1.2.1 *An n-dimensional topological manifold $\mathcal{M}$ is a Hausdorff space each of whose points has a neighbourhood homeomorphic to (an open set in) $\mathbb{R}^n$.*

An *open cover* of a topological space $\mathcal{S}$ is a collection of open sets whose union is the total space $\mathcal{S}$. It follows from Definition 1.2.1 that for a topological manifold there exists an open cover each of whose elements is homeomorphic to an open set of $\mathbb{R}^n$. If we denote by $\mathcal{U}_\alpha$ the generic constituent of the open cover and by ϕ_α the corresponding homeomorphism, where α denotes a running index, we can identify the pair $(\mathcal{U}_\alpha, \phi_\alpha)$ with a *coordinate chart*. The collection of all these pairs is called an *atlas* of the topological manifold.

The terminology of coordinate charts arises from the fact that a chart introduces a *local coordinate system*. More specifically, the homeomorphism $\phi_\alpha : \mathcal{U}_\alpha \to \phi_\alpha(\mathcal{U}_\alpha) \subset \mathbb{R}^n$ assigns to each point $p \in \mathcal{U}_\alpha$ an ordered n-tuple $(x^1(p), ..., x^n(p))$, called the *local coordinates* of p.

Whenever two charts, $(\mathcal{U}_\alpha, \phi_\alpha)$ and $(\mathcal{U}_\beta, \phi_\beta)$, have a non-empty intersection, we define the *transition function* $\phi_{\alpha\beta}$ as:

$$\phi_{\alpha\beta} = \phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha(\mathcal{U}_\alpha \cap \mathcal{U}_\beta) \to \phi_\beta(\mathcal{U}_\alpha \cap \mathcal{U}_\beta). \tag{2.1}$$

Each transition function is a homeomorphism between open sets of $\mathbb{R}^n$. The inverse of $\phi_{\alpha\beta}$ is the transition function $\phi_{\beta\alpha} = \phi_{\alpha\beta}^{-1} = \phi_\alpha \circ \phi_\beta^{-1}$. Denoting by x^i and y^i $(i = 1, ..., n)$, respectively, the local coordinates of $\mathcal{U}_\alpha$ and $\mathcal{U}_\beta$, a transition function boils down to the specification of n continuous and continuously invertible real functions of the form:

$$y^i = y^i(x^1, ..., x^n), \qquad i = 1, ..., n. \tag{2.2}$$

1.2.1 Maps and their representations

If $\mathcal{M}$ and $\mathcal{N}$ are topological manifolds of dimensions m and n, respectively, a map $f : \mathcal{M} \to \mathcal{N}$ is *continuous* if it is a continuous map between the underlying topological spaces. A nice feature of topological manifolds, as opposed to general topological spaces, is the possibility of representing continuous maps locally as real functions of real variables. Let $p \in \mathcal{M}$ and denote $q = f(p) \in \mathcal{N}$, as shown in Figure 1.1. By continuity, we can always choose a chart $(\mathcal{U}, \phi)$ containing p such that its image $f(\mathcal{U})$ is contained in a chart $(\mathcal{V}, \psi)$ containing q. The map:

$$\hat{f} = \psi \circ f \circ \phi^{-1} : \phi(\mathcal{U}) \to \psi(\mathcal{V}) \tag{2.3}$$

maps an open set in $\mathbb{R}^m$ to an open set in $\mathbb{R}^n$. This continuous map $\hat{f}$ is the *local coordinate representation* of f in the coordinate charts chosen.

1.3 Topological groups

Recall that a *group* is a set $\mathcal{G}$ endowed with a binary associative internal operation, called *group multiplication* or *group product*, which is usually indicated by simple apposition, namely: if $g, h \in \mathcal{G}$ then the

[1] We are leaving out an important technical detail, namely, the existence of a countable basis for the topology.

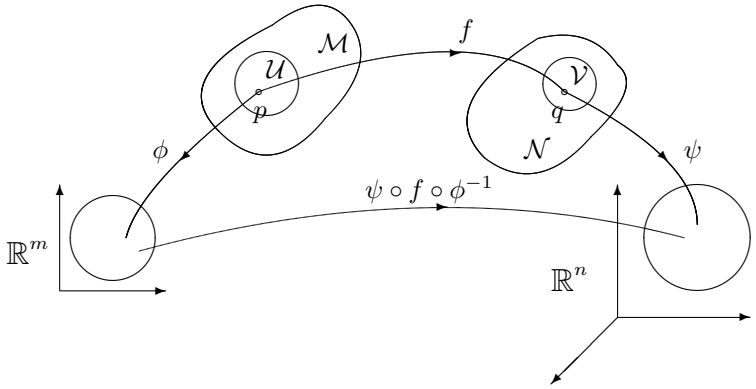

Figure 1.1: Representation of a map in two charts

product is $gh \in \mathcal{G}$. Associativity means that $(gh)k = g(hk)$, for all $g, h, k \in \mathcal{G}$. Moreover, there exists an *identity element* $e \in \mathcal{G}$ such that $eg = ge = g$ for all $g \in \mathcal{G}$. Finally, for each $g \in \mathcal{G}$ there exists an *inverse* $g^{-1} \in \mathcal{G}$ such that $gg^{-1} = g^{-1}g = e$. The identity can be shown to be unique, and so is also the inverse of each element of the group. If the group operation is also *commutative*, namely, if $gh = hg$ for all $g, h \in \mathcal{G}$, the group is said to be *commutative* or *Abelian*. In this case, it is customary to call the operation *group addition* and to indicate it as: $g + h$. The identity is then called the *zero element* and is often denoted as 0. Finally, the inverse of g is denoted as $-g$. This notation is easy to manipulate as it is reminiscent of the addition of numbers, which is indeed a particular case.

A *subgroup* of a group $\mathcal{G}$ is a subset $\mathcal{H} \subset \mathcal{G}$ closed under the group operations of multiplication and inverse. Thus, a subgroup is itself a group.

Given two groups, $\mathcal{G}_1$ and $\mathcal{G}_2$, a *group homomorphism* is a map $\phi : \mathcal{G}_1 \to \mathcal{G}_2$ that preserves the group multiplication, namely:

$$\phi(gh) = \phi(g)\ \phi(h) \quad \forall\ g, h \in \mathcal{G}_1, \tag{3.1}$$

where the multiplications on the left and right-hand sides are, respectively, the group multiplications of $\mathcal{G}_1$ and $\mathcal{G}_2$.

The group structure is a purely algebraic concept, whereby nothing is assumed as far as the nature of the underlying set is concerned. The concept of *topological group* arises from making such an assumption. More specifically, a *topological group* is a topological manifold $\mathcal{G}$ with a group structure that is compatible with the topological structure, namely, such that the multiplication $\mathcal{G} \times \mathcal{G} \to \mathcal{G}$ and the inversion $\mathcal{G} \to \mathcal{G}$ are continuous maps.

1.3.1 Group actions

Let $\mathcal{G}$ be a group (not necessarily a topological group) and let X be a set (not necessarily a topological manifold). We say that the group $\mathcal{G}$ *acts on the right* on the set X if for each $g \in \mathcal{G}$ there is a map $R_g : X \to X$ such that: (i) $R_e(x) = x$ for all $x \in X$, where e is the group identity; (ii) $R_g \circ R_h = R_{hg}$ for all $g, h \in \mathcal{G}$. When there is no room for confusion, we also use the notation xg for $R_g(x)$. Each of the maps R_g is a bijection of X. Moreover, $R_{g^{-1}} = (R_g)^{-1}$. The *orbit* through $x \in X$ is the subset $x\mathcal{G}$ of X consisting of all the elements of X of the form xg, where $g \in \mathcal{G}$.

The action of $\mathcal{G}$ on X is said to be *effective* if the condition $R_g(x) = x$ for every $x \in X$ implies $g = e$. The action is *free* if $R_g(x) = x$ for *some* $x \in X$ implies $g = e$. Finally, the action is *transitive* if for every $x, y \in X$ there exists $g \in \mathcal{G}$ such that $R_g(x) = y$.

In a completely analogous manner, we can say that $\mathcal{G}$ *acts on the left* on X if for each $g \in \mathcal{G}$ there is a map $L_g : X \to X$ such that: (i) $L_e(x) = x$ for all $x \in X$, where e is the group identity; (ii) $L_g \circ L_h = L_{gh}$ for all $g, h \in \mathcal{G}$. The order of the composition is the essential difference between a right and a left action. We may also use the notation gx for $L_g(x)$.

The notion of group action can naturally be applied when $\mathcal{G}$ is a topological group. In this instance, a case of particular interest is that for which the set on which $\mathcal{G}$ acts is a topological manifold and the induced bijections are transformations of this manifold. A *transformation* of a manifold $\mathcal{M}$ is a homeomorphism $\phi : \mathcal{M} \to \mathcal{M}$. The definition of the action is then supplemented with a continuity condition. More explicitly: A topological group $\mathcal{G}$ is said to *act on the right* on a manifold $\mathcal{M}$ if:

(1) Every element $g \in \mathcal{G}$ induces a transformation $R_g : \mathcal{M} \to \mathcal{M}$.

(2) $R_g \circ R_h = R_{hg}$, namely: $(ph)g = p(hg)$ for all $g, h \in \mathcal{G}$ and $p \in \mathcal{M}$.

(3) The *right action* $R : \mathcal{G} \times \mathcal{M} \to \mathcal{M}$ is a continuous map. In other words, $R_g(p)$ is continuous in *both* variables (g and p).

With these conditions, the topological group $\mathcal{G}$ is also called a *topological group of transformations of* $\mathcal{M}$. Just as in the general case, we have used the alternative notation pg for $R_g(p)$, with $p \in \mathcal{M}$, wherever convenient. A similar definition can be given for the *left action* of a topological group on a manifold.

Condition 1 is equivalent to the fact that R_e (and L_e) are the identity transformation of $\mathcal{M}$. Indeed, since a transformation is an invertible map, every point $p \in \mathcal{M}$ can be expressed as qg for some $q \in \mathcal{M}$ and some $g \in \mathcal{G}$. Using Property (2) of the right action we have: $R_e(p) = pe = (qg)e = q(ge) = qg = p$, with a similar proof for the left action.

It is convenient to introduce the following (useful, though potentially confusing) notation. We denote the right action *as a map from* $\mathcal{G} \times \mathcal{M}$ *to* $\mathcal{M}$ by the symbol R. Thus, $R = R(g, p)$ has two arguments, one in the group and the other in the manifold. Fixing, therefore, a particular element g in the group, we obtain a function of the single variable x which we have already denoted by $R_g : \mathcal{M} \to \mathcal{M}$. But we can also fix a particular element p in the manifold and thus obtain another function of the single variable g. We will denote this function by $R_p : \mathcal{G} \to \mathcal{M}$. A similar scheme of notation can be adopted for a left action L. Notice that the image of R_p (respectively L_p) is nothing but the orbit $p\mathcal{G}$ (respectively $\mathcal{G}p$). The potential for confusion arises when the manifold $\mathcal{M}$ happens to coincide with the group $\mathcal{G}$, as described below. Whenever an ambiguous situation arises, we will resort to the full action function of two variables.

Recall that a topological group can be both a group and a manifold. Thus, it is not surprising that every topological group $\mathcal{G}$ induces two canonical groups of transformations on itself, one by right action and one by left action, called, respectively, *right translations* and *left translations* of the group. They are defined, respectively, by: $R_g(h) = hg$ and $L_g(h) = gh$, with $g, h \in \mathcal{G}$, where the right-hand sides are given by the group multiplication itself. For this reason, it should be clear that these actions are both free (and, hence, effective) and transitive.

1.4 Topological fibre bundles

Topological fibre bundles are topological manifolds with extra structure. The points of a fibre bundle have, as it were, a double allegiance: not only to the manifold itself but also to a smaller entity called a *fibre*.

A topological fibre bundle generalizes the idea of a product topological manifold, which we use as a point of departure. Let two topological manifolds $\mathcal{B}$ and $\mathcal{F}$ be given. Consider the Cartesian product

$$\mathcal{C} = \mathcal{B} \times \mathcal{F}. \tag{4.1}$$

The *total space* $\mathcal{C}$ can be shown to be itself a topological manifold,[2] whose dimension is the sum of the dimensions of the *base manifold* $\mathcal{B}$ and the *typical fibre* $\mathcal{F}$.

We call $\mathcal{C}$ a *product bundle*. It is endowed with two natural *projection maps*, namely:

$$pr_1 : \mathcal{C} \longrightarrow \mathcal{B}\ , \tag{4.2}$$

and

$$pr_2 : \mathcal{C} \longrightarrow \mathcal{F}\ , \tag{4.3}$$

which assign to any given pair $(b, f) \in \mathcal{C}$ its first and second components, b and f, respectively. It is clear that these two maps, in addition to being surjective, are continuous in the product topology. Indeed, let $\mathcal{U}$ be an open subset of $\mathcal{B}$. We have, by definition of pr_1, that $pr_1^{-1}(\mathcal{U}) = \mathcal{U} \times \mathcal{F}$, which is open. A similar reasoning applies to pr_2.

For each $b \in \mathcal{B}$ the set $pr_1^{-1}(b)$ is called the *fibre at* b, denoted as $\mathcal{C}_b$, which in this case is simply a copy of $\mathcal{F}$. From this definition it follows that fibres at different points are disjoint sets and that each point $c \in \mathcal{C}$ necessarily belongs to a fibre, namely, to $\mathcal{C}_{pr_1(c)}$. The fibres can, therefore, also be seen as the equivalence classes corresponding to the equivalence relation of "having the same first projection".

Given an atlas in $\mathcal{B}$ and an atlas in $\mathcal{F}$, the product of these atlases is an atlas of the product bundle $\mathcal{C}$. Naturally, as a manifold, $\mathcal{C}$ may be endowed with other atlases. Nevertheless, we will always restrict attention to *product atlases*, namely, those that emphasize the product nature of the bundle $\mathcal{C}$. In a product chart, the projections acquire a particularly simple form. Indeed, let $x^i, i = 1, ..., m = \dim(\mathcal{B})$ and $y^\alpha, \alpha = 1, ..., n = \dim(\mathcal{F})$ be coordinate systems for some charts in the base and the fibre, respectively. Then we have:

$$\begin{array}{rcl} pr_1 : \mathcal{C} & \longrightarrow & \mathcal{B} \\ (x^i, y^\alpha) & \mapsto & (x^i)\ , \end{array} \tag{4.4}$$

and

$$\begin{array}{rcl} pr_2 : \mathcal{C} & \longrightarrow & \mathcal{F} \\ (x^i, y^\alpha) & \mapsto & (y^\alpha)\ . \end{array} \tag{4.5}$$

It follows from this coordinate representation that the projections are continuous surjections.

[2] Strictly speaking, fibre bundles can be defined using topological spaces, rather than the more specialized topological manifolds. The topology of the Cartesian product consists of all the possible unions of all possible Cartesian products of open sets of $\mathcal{B}$ with open sets of $\mathcal{F}$, taken in this order.

We want to proceed now to a more general concept than that of a product bundle just discussed. In essence, what we want to achieve is the loss of the second projection (pr_2), while preserving the first (pr_1). A clear example of the convenience of such a generalization is provided by the concept of *space-time*. The ancients (particularly, but not only, Aristotle) seemed to have thought that space and time were absolute entities existing, as it were, independently of each other. In our terminology, therefore, we would say that for them space-time was a product bundle with, say, time as the base manifold and space as the fibre. The physical meaning of pr_1 would thus be that of providing the time of occurrence of an event, such as a collision. Accordingly, two events p and q are *simultaneous* if $pr_1(p) = pr_1(q)$. Similarly, since space is absolute in this vision of the world, the second projection provides information about the location of an event in absolute space. Thus, two events p and q can be said to have occurred at the same place (regardless of their times of occurrence) if $pr_2(p) = pr_2(q)$. The principle of Galilean relativity can be said to have demolished the second projection. Indeed, unless two events happen to be simultaneous (a concept only questioned much later by Einstein's relativity principle), it is impossible, according to Galileo and Newton, to compare in an absolute way the places at which they occurred. While it is true that a given observer can make such a judgment, a different observer will in general legitimately disagree. Herein lies the clue to our generalization, namely, that although the total manifold $\mathcal{C}$ is no longer a product it still looks like a product (albeit a different one) to each observer. This intuitive idea leads to the following definition.

A fibre bundle with base $\mathcal{B}$, typical fibre $\mathcal{F}$ and structure group $\mathcal{G}$ (acting effectively on the left on $\mathcal{F}$), is a manifold $\mathcal{C}$ and a continuous surjective bundle-projection map $\pi : \mathcal{C} \to \mathcal{B}$ such that there exists an open covering $\mathcal{U}_\alpha$ of $\mathcal{B}$ and respective *local trivializations*:

$$\psi_\alpha : \pi^{-1}(\mathcal{U}_\alpha) \longrightarrow \mathcal{U}_\alpha \times \mathcal{F} \tag{4.6}$$

with the property $\pi = pr_1 \circ \psi_\alpha$, as illustrated in the following commutative diagram:

$$\begin{array}{ccc} \pi^{-1}(\mathcal{U}_\alpha) & \xrightarrow{\psi_\alpha} & \mathcal{U}_\alpha \times F \\ & \searrow^{\pi} & \downarrow{\scriptstyle pr_1} \\ & & \mathcal{U}_\alpha \end{array} \tag{4.7}$$

Moreover, as illustrated in Figure 1.2, whenever $b \in \mathcal{U}_\alpha \bigcap \mathcal{U}_\beta \neq \emptyset$, the transition maps $\tilde{\psi}_{\beta,\alpha}(b) := \tilde{\psi}_{\alpha,b} \circ \tilde{\psi}^{-1}_{\beta,b}$ belong to the structure group $\mathcal{G}$ and depend continuously on position throughout the intersection.

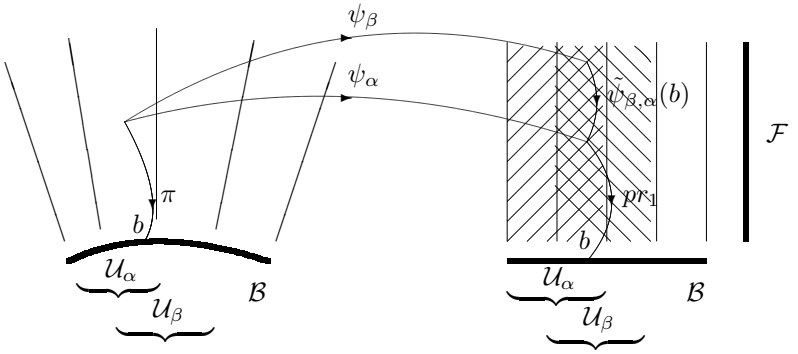

Figure 1.2: A general fibre bundle

Consider now, for the same $\mathcal{C}$, $\mathcal{B}$, $\mathcal{F}$, π and $\mathcal{G}$, a different open covering $\mathcal{V}_\beta$ with local trivializations ϕ_β. We say that it defines the same fibre bundle as before if, on non-vanishing intersections, the transition

maps $\tilde{\psi}_{\alpha,b} \circ \tilde{\phi}^{-1}_{\beta,b}$ belong to the structure group $\mathcal{G}$ and depend continuously on position b throughout the intersection. The two trivializations are said to be compatible. In this sense, we can say that the union of the two trivialization coverings becomes itself a new trivialization covering of the fibre bundle. When there is no room for confusion, a fibre bundle is denoted as a pair $(\mathcal{C}, \pi)$ indicating just the total space and the projection. An alternative notation is $\pi : \mathcal{C} \to \mathcal{B}$. A more complete notation would be $(\mathcal{C}, \pi, \mathcal{B}, \mathcal{F}, \mathcal{G})$.

According to the general definition, a fibre bundle is always locally trivializable. In other words, every point of the base manifold has a neighbourhood such that the restriction of the bundle to this neighbourhood is homeomorphic to its product with the typical fibre. If this property can be extended to the whole base manifold, that is, if the whole bundle happens to be homeomorphic to the product of the base times the typical fibre, we say that the bundle is *trivial* or, more precisely, (globally) *trivializable*.

The *fundamental existence theorem* of fibre bundles states that given the manifolds $\mathcal{B}$ and $\mathcal{F}$ and a topological group $\mathcal{G}$ acting effectively to the left on $\mathcal{F}$, and given, moreover, an open covering $\mathcal{U}_\alpha$ of $\mathcal{B}$ and a continuous assignment of an element of $\mathcal{G}$ to each point in every non-vanishing intersection of the covering, then there exists a fibre bundle $(\mathcal{C}, \pi)$ with local trivializations based upon that covering and with the assigned elements of $\mathcal{G}$ as transition maps. Furthermore, any two bundles with this property are equivalent.

An important application of the fundamental existence theorem is that given a bundle $(\mathcal{C}, \pi, \mathcal{B}, \mathcal{F}, \mathcal{G})$, we can associate to it other bundles with the same base manifold and the same structure group, but with different typical fibre $\mathcal{F}'$, in a precise way. Indeed, we can choose a trivialization covering of the given bundle, calculate the transition maps, and then define the *associated bundle* $(\mathcal{C}', \pi', \mathcal{B}, \mathcal{F}', \mathcal{G})$, modulo an equivalence, by means of the assertion of the fundamental theorem. A case of particular interest is that in which the new fibre is identified with the structure group. This gives rise to the so-called *associated principal bundle*.

1.4.1 Principal bundles

A *principal bundle* is a topological fibre bundle for which the typical fibre and the structure group coincide. As in any other topological bundle, the action of the structure group on the typical fibre is assumed to be a left action. Accordingly, in a principal bundle the action of the structure group is identified with the natural (or canonical) left action of the group on itself. The existence of a canonical *right* action of a group on itself, on the other hand, is a bonus that can be used to show that in a principal bundle $(\mathcal{P}, \pi, \mathcal{B}, \mathcal{G}, \mathcal{G})$ one can define a canonical right action of the structure group *on the total space* $\mathcal{P}$. By *canonical* we mean that this action is independent of the trivialization. Moreover, this canonical right action is free and fibre-preserving.

The existence of a free right action on a manifold is strong enough to provide an alternative definition of a principal fibre bundle which, although equivalent to the one just given, has the merit of being independent of the notion of transition maps. Moreover, once this more elegant and constructive definition has been secured, a subsidiary definition of the associated (non-principal) bundles becomes available, again without an explicit mention of the transition maps. Finally, this more abstract definition brings out intrinsically the nature and meaning of the associated bundles.

Let $\mathcal{P}$ be a topological manifold (the *total space*) and $\mathcal{G}$ a topological group (the *structure group*), and

let $\mathcal{G}$ act freely to the right on $\mathcal{P}$. This means that there exists a continuous map:

$$\begin{aligned} R_g : \mathcal{P} \times \mathcal{G} &\longrightarrow \mathcal{P} \\ (p, g) &\mapsto R_g p = pg, \end{aligned} \tag{4.8}$$

such that, for all $p \in \mathcal{P}$ and all $g, h \in \mathcal{G}$, we have:

$$\begin{aligned} R_{gh} p &= R_h R_g p = pgh, \\ R_e p &= p \end{aligned} \tag{4.9}$$

where e is the group identity. The fact that the action is free means that if, for some $p \in \mathcal{P}$ and some $g \in \mathcal{G}$, $R_g p = p$, then necessarily $g = e$. Define now the quotient $\mathcal{B} = \mathcal{P}/\mathcal{G}$ and check that $\mathcal{B}$ is a topological manifold and that the canonical projection $\pi_P : \mathcal{P} \to \mathcal{P}/\mathcal{G}$ is continuous. The set $\pi_P^{-1}(b)$ is called the fibre over $b \in \mathcal{B}$.

Recall that an element of the quotient $\mathcal{B} = \mathcal{P}/\mathcal{G}$ is, by definition of quotient, an equivalence class in $\mathcal{P}$ by the action of the group $\mathcal{G}$. In other words, each element b of the quotient (namely, of the base manifold $\mathcal{B}$) can be regarded as representing an orbit. The projection map assigns to each element of $\mathcal{P}$ the orbit to which it belongs. The fibre over b consists of all the elements of $\mathcal{P}$ that belong to the specific orbit represented by b.

To complete the definition of a principal bundle, we need only to add the condition that $\mathcal{P}$ be locally trivial, namely, that for each $b \in \mathcal{B}$, there exists a neighbourhood $\mathcal{U} \subset \mathcal{P}$ such that $\pi_P^{-1}(\mathcal{U})$ is isomorphic to the product $\mathcal{U} \times \mathcal{G}$. More precisely, there exists a fibre-preserving homeomorphism:

$$\begin{aligned} \psi : \pi_P^{-1}(\mathcal{U}) &\longrightarrow \mathcal{U} \times \mathcal{G} \\ p &\mapsto (b, \tilde{\psi}_b), \end{aligned} \tag{4.10}$$

where $b = \pi_P(p)$, with the additional property that it must be *consistent with the group action*, namely (see Figure 1.3):

$$\tilde{\psi}_b(pg) = \tilde{\psi}_b(p) g \qquad \forall p \in \pi_P^{-1}(\mathcal{U}), g \in \mathcal{G}. \tag{4.11}$$

This completes the definition of the principal bundle. The right action is fibre-preserving and every fibre is homeomorphic to $\mathcal{G}$. Moreover, every fibre coincides with an orbit of the right action of $\mathcal{G}$.

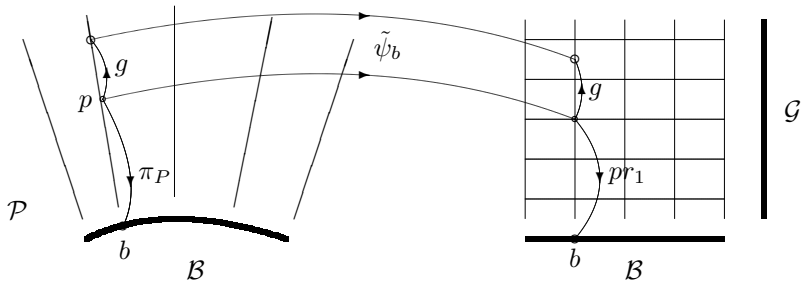

Figure 1.3: The group consistency condition

1.4.2 Equivalence of both definitions of a principal bundle

Consider a principal fibre bundle $\mathcal{P}$, with structure group $\mathcal{G}$, defined via the free right action of $\mathcal{G}$ on $\mathcal{P}$. Let two local trivializations, $(\mathcal{U}, \psi)$ and $(\mathcal{V}, \phi)$, be given and let $b \in \mathcal{U} \bigcap \mathcal{V}$. We want to show that for each

point $b \in \mathcal{B}$ the restricted map $\tilde{\phi}_b \circ \tilde{\psi}_b^{-1} : \mathcal{G} \to \mathcal{G}$ is in fact equivalent to the left action of a fixed element $a \in \mathcal{G}$. Let $g \in \mathcal{G}$. Since $\tilde{\phi}_b \circ \tilde{\psi}_b^{-1}(g)$ belongs to $\mathcal{G}$ and since the left action of $\mathcal{G}$ on itself is transitive, there exists a unique $a \in \mathcal{G}$ such that:

$$\tilde{\phi}_b \circ \tilde{\psi}_b^{-1}(g) = L_a g = ag. \tag{4.12}$$

We need to prove that this $a \in \mathcal{G}$ is independent of g. Let, therefore, $h \in \mathcal{G}$ be some other point in the typical fibre. By transitivity of the right action there exists a unique $c \in \mathcal{G}$ such that $h = gc$. Therefore,

$$\begin{aligned} \tilde{\phi}_b \circ \tilde{\psi}_b^{-1}(h) = \tilde{\phi}_b \circ \tilde{\psi}_b^{-1}(gc) = \tilde{\phi}_b[R_c \tilde{\psi}_b^{-1}(g)] \\ = R_c[\tilde{\phi}_b \circ \tilde{\psi}_b^{-1}(g)] = R_c L_a(g) = L_a R_c g = L_a h. \end{aligned} \tag{4.13}$$

The crucial step in this chain is the commutation of the trivializations $\tilde{\phi}_b$ and $\tilde{\psi}_b$ with the right action R_c, namely, the consistency condition. The last result clearly shows that transition maps can indeed be defined and that they belong to the structure group, as required by the standard definition. The equivalence of both definitions follows now from the fundamental existence theorem.

1.4.3 Cross sections

A *cross section* σ of a fibre bundle $(\mathcal{C}, \pi, \mathcal{B}, \mathcal{F}, \mathcal{G})$ is a (continuous) map:

$$\sigma : \mathcal{B} \longrightarrow \mathcal{C} \tag{4.14}$$

such that $\pi \circ \sigma = id_{\mathcal{B}}$, as shown in the following commutative diagram:

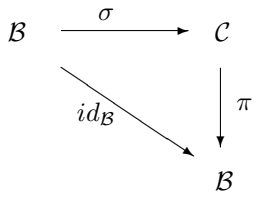

(4.15)

A cross section is thus nothing but a continuous assignment, to each point b in the base manifold, of an element of its fibre $\mathcal{C}_b$. At this point, it is appropriate to point out that one of the most useful physical interpretations of fibre bundles consists of regarding the typical fibre as a set of objects of some type (such as vectors, tensors, frames, group elements, and so on) which are paired with each of the points of an underlying continuum, namely the base manifold. With this interpretation in mind, we observe that a cross section corresponds precisely to the notion of a physical *field*. Indeed, a cross section is a continuous assignment of an element of the fibre (a vector, a tensor, a frame, a group element) to each point of the base (the underlying continuum).

It is important to realize that not all fibre bundles admit continuous cross sections. Globally trivializable fibre bundles, on the other hand, always do. Indeed, let $(\mathcal{B}, \psi)$ be a global trivialization of a globally trivializable fibre bundle and let a denote a fixed element of the typical fibre $\mathcal{F}$. Then, the function defined as $\sigma(b) = \tilde{\psi}_b^{-1}(a)$ is a cross section. One should not think, however, that this sufficient condition is also necessary for the existence of a cross section. As an intuitive example to the contrary, one can take the Moebius band, which is a non-trivializable bundle over the circle. Any continuous curve drawn on the original strip such that its end points are symmetrically arranged will do the job as a cross section,

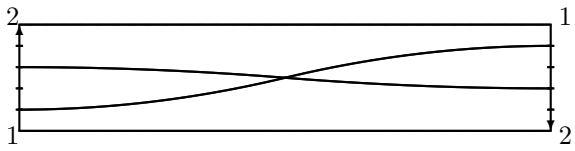

Figure 1.4: Cross sections of a Moebius band

as shown in Figure 1.4. (Notice, however, that any two such cross sections will end up having at least one point in common!)

The previous remark notwithstanding, we can prove that, in the case of principal bundles, global trivializability is equivalent to the existence of a cross section. We need to prove just the "only if" part, since the "if" part is always true, as we have shown in the previous remark. Assume, therefore, that a given principal bundle $(\mathcal{P}, \pi, \mathcal{B}, \mathcal{G}, \mathcal{G})$ admits a cross section $\sigma : \mathcal{B} \to \mathcal{P}$. The proof follows from the transitivity of the right action of $\mathcal{G}$ on $\mathcal{P}$. For, let $p \in \mathcal{P}$. By the transitivity of the right action on each fibre, there exists a unique $g(p) \in \mathcal{G}$ such that $p = R_{g(p)} \circ \sigma \circ \pi_P(p)$. Define the map:

$$\begin{aligned} \psi : \mathcal{P} &\longrightarrow \mathcal{B} \times \mathcal{G} \\ p &\mapsto (\pi_P(p), g(p)). \end{aligned} \tag{4.16}$$

It is not difficult to verify that this map is a global trivialization of the principal bundle, satisfying the group consistency condition. A nice way to picture this situation (Figure 1.5) is to imagine that the given cross section σ is translated by the right action of the group to give rise to a family of cross sections $R_g \circ \sigma$. Since the action is effective, no two such cross sections will intersect (provided just g is not the group identity). Moreover, since the action is transitive on fibres, every point of the fibre bundle will belong to one (and only one) translated cross section. In other words, we have a family of cross sections, parametrized by the structure group, that completely spans the total space. This is, naturally, tantamount to a global trivialization.

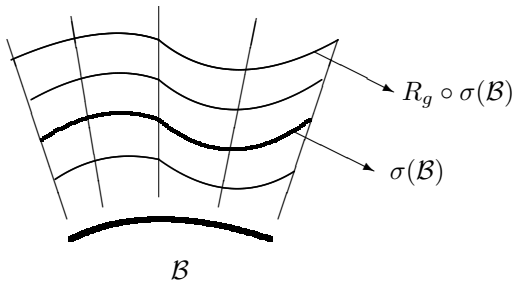

Figure 1.5: Translation of a cross section by the right action

An interesting corollary of this theorem is that, since the trivialization coverings of all associated bundles are the same, a fibre bundle is trivializable if, and only if, its associated principal bundle admits a cross section. This result is just one illustration of the assertion to the effect that working with the associated principal bundle often simplifies and helps to understand the mathematical picture.

1.5 Topological groupoids

The abstract notion of a *groupoid* emerges as the common structure underlying many constructions that arise naturally in a variety of apparently disconnected applications in algebra, topology, geometry, differential equations, numerical analysis and practically every branch of mathematics. In a restricted way, it can be seen as a generalization of the notion of group, but it is better to understand it as an important mathematical concept in its own right.[3]

A groupoid consists of a total set $\mathcal{Z}$, a base set $\mathcal{B}$, two ("projection") surjective maps:

$$\alpha : \mathcal{Z} \longrightarrow \mathcal{B} \quad \text{and} \quad \beta : \mathcal{Z} \longrightarrow \mathcal{B} \tag{5.1}$$

called, respectively, the *source* and the *target* maps, and a binary operation ("composition") defined only for those ordered pairs $(y, z) \in \mathcal{Z} \times \mathcal{Z}$ such that:

$$\alpha(z) = \beta(y). \tag{5.2}$$

This operation (usually indicated just by reverse apposition of the operands) must satisfy the following properties:

(1) Associativity:

$$(xy)z = x(yz), \tag{5.3}$$

whenever either product is defined;

(2) Existence of identities: for each $b \in \mathcal{B}$ there exists an element $id_b \in \mathcal{Z}$, called the *identity at* b, such that $z\ id_b = z$ whenever $\alpha(z) = b$, and $id_b\ z = z$ whenever $\beta(z) = b$;

(3) Existence of inverse: for each $z \in \mathcal{Z}$ there exists a (unique) *inverse* z^{-1} such that

$$zz^{-1} = id_{\beta(z)} \quad \text{and} \quad z^{-1}z = id_{\alpha(z)}. \tag{5.4}$$

It follows from this definition that to each ordered pair (a, b) of elements of $\mathcal{B}$ one can associate a definite subset $\mathcal{Z}_{ab}$ of $\mathcal{Z}$, namely the subset: $\{z \in \mathcal{Z} \mid \beta(z) = b, \ \ \alpha(z) = a\}$. It is clear that these sets (some of which may be empty) are disjoint and that their union is equal to $\mathcal{Z}$. It is also clear that the various identities are elements of subsets of the form $\mathcal{Z}_{bb}$. It is not difficult to show that each set of the form $\mathcal{Z}_{bb}$ is actually a group.

A useful way to think of a groupoid is as a collection of symbols $(a, b, c, ...)$ and arrows $(x, y, z, ...)$ connecting some of them. The symbols correspond to the elements of the base set $\mathcal{B}$, while the arrows correspond to the elements of the total set $\mathcal{Z}$. The tail and tip of an arrow z correspond to the source $\alpha(z)$ and the target $\beta(z)$, respectively. Two arrows z and y can be composed if, and only if, the tip of the first ends where the tail of the second begins. The result is an arrow yz whose tail is the tail of z and

[3] For a thorough treatment of groupoids see: MacKenzie K (1987), *Lie Groupoids and Lie Algebroids in Differential Geometry*, London Mathematical Society Lecture Note Series **124**, Cambridge University Press. An informal and illuminating explanation can be found in: Weinstein A (2000), Groupoids: Unifying Internal and External Symmetry. A tour through Examples, *Notices of the American Mathematical Society* **43**, 744-752.

whose tip is the tip of y:

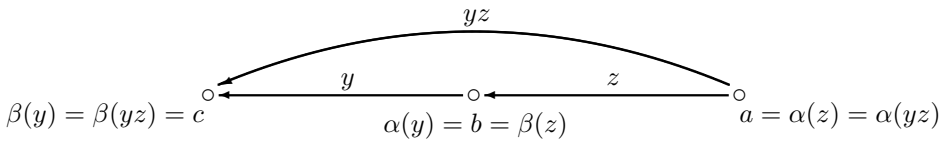

For this picture to correspond more or less exactly to the more formal definition of a groupoid, however, we have to add the proviso that for each arrow z connecting point a to point b, there exists an "inverse" arrow z^{-1} connecting point b with point a. It is also very important to bear in mind that there is no need for a given pair of points to be connected by one or more arrows. Some may be connected and some may not. In fact, an extreme case can occur whereby no two (different) points are thus connected. In this extreme case, the set $\mathcal{Z}$ becomes simply the disjoint union of the groups $\mathcal{Z}_{bb}$.

Box 1.5.1 Some examples of groupoids

To show the versatility of the concept of groupoid, we list a few examples drawn from different areas of Mathematics.

(1) The product groupoid: Given a set $\mathcal{B}$, the Cartesian product $\mathcal{B} \times \mathcal{B}$ is a groupoid with $\alpha = pr_1$ and $\beta = pr_2$.

(2) The general linear groupoid $GL(\mathbb{R})$: Take as the total set the collection of all non-singular square matrices of all orders. The base space will be taken as the natural numbers. The binary operation is matrix multiplication. We can see that this groupoid is nothing but the disjoint union of all the general linear groups $GL(n;\mathbb{R})$.

(3) The fundamental groupoid: Let $\mathcal{T}$ be a topological space. For each pair of points $a, b \in \mathcal{T}$ we consider the collection of all continuous curves starting at a and ending at b. We partition this set into equivalence classes, two curves being considered equivalent if the are homotopic[a], and we define $\mathcal{Z}_{ab}$ as the quotient set (namely, the set of these equivalence classes). The composition of curves is done just like with the arrows of our pictorial description. [Question: why is the partition into equivalence classes needed?]

[a] Two curves starting and ending at the same points are homotopic if, keeping these ends fixed, it is possible to transform continuously one curve into the other.

One can prove that if $\mathcal{Z}_{ab} \neq \emptyset$, then the groups $\mathcal{Z}_{aa}$ and $\mathcal{Z}_{bb}$ are conjugate, and the conjugation between them is achieved by any element of $\mathcal{Z}_{ab}$. Moreover, the set $\mathcal{Z}_{ab}$ is spanned completely by composing any one of its elements with $\mathcal{Z}_{aa}$ or with $\mathcal{Z}_{bb}$ (to the right or to the left, of course).

A groupoid is said to be *transitive* if for each pair of points $a, b \in \mathcal{B}$ there exists at least one element of the total set with a and b as the source and target points, respectively. In other words, a groupoid is transitive if, and only if, $\mathcal{Z}_{ab} \neq \emptyset \ \ \forall (a,b) \in \mathcal{B} \times \mathcal{B}$. In a transitive groupoid all the local groups $\mathcal{Z}_{bb}$ are mutually conjugate.

A groupoid is a *topological groupoid* if the total set $\mathcal{Z}$ and the base set $\mathcal{B}$ are topological manifolds, the projections α and β are continuous, and so are the operations of composition and of inverse. It follows from the definition that each of the sets $\mathcal{Z}_{bb}$ is a topological group.

1.5.1 From groupoids to principal bundles

Let $b \in \mathcal{B}$ be a fixed point in the base manifold of a transitive topological groupoid $\mathcal{Z}$. Consider the subset of the total set $\mathcal{Z}$ formed by the disjoint union $\tilde{\mathcal{Z}}_b$ of all the sets $\mathcal{Z}_{bx}, \forall x \in \mathcal{B}$. The elements $\tilde{z}$ of this set have the property $\alpha(\tilde{z}) = b$. The group $\mathcal{Z}_{bb}$ has a natural effective right action on $\tilde{\mathcal{Z}}_b$, as can be verified directly by composition. Moreover, two elements of $\tilde{\mathcal{Z}}_b$ that differ by the right action of an element of this group must have the same target. In other words, the equivalence classes corresponding to this action consist precisely of the sets $\mathcal{Z}_{bx}$ and, therefore, the quotient set is precisely the manifold $\mathcal{B}$. We are thus led to a principal bundle with total space $\tilde{\mathcal{Z}}_b$, structure group $\mathcal{Z}_{bb}$ and projection β (or, rather, the restriction of β to $\tilde{\mathcal{Z}}_b$).

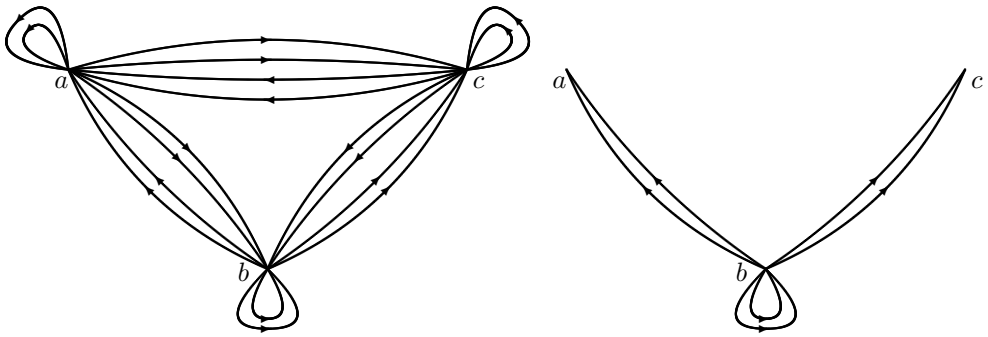

Figure 1.6: Schematic representation of a transitive groupoid $\mathcal{Z}$ (left) and the induced principal bundle $\tilde{\mathcal{Z}}_b$ (right), whose fibres are depicted as pairs of arrows projecting on the base manifold (a, b, c)

If we were to start from a different point, c say, of $\mathcal{B}$, the previous construction would lead to a principal bundle whose structure group $\mathcal{Z}_{cc}$ is conjugate to $\mathcal{Z}_{bb}$, and it is not difficult to show that the two principal bundles are isomorphic. We see, therefore, that giving a transitive topological groupoid is tantamount to giving an equivalence class of isomorphic principal bundles, each one conveying the same information as the groupoid. The choice of the reference point of departure is somewhat analogous to the choice of a basis in a vector space. No information is lost, but there is a certain loss of objectivity, in the sense that one is no longer working with the actual objects but rather with their representation in the chosen reference.

1.5.2 From principal bundles to groupoids

Somewhat more surprising than the previous passage from a groupoid to any one of its representative principal bundles is the fact that, given an arbitrary principal bundle $(\mathcal{P}, \pi_P, \mathcal{B}, \mathcal{G}, \mathcal{G})$, one can construct a groupoid of which it is a representative. Indeed, define the quotient space $\mathcal{Z} = (\mathcal{P} \times \mathcal{P})/\mathcal{G}$ as the total set of the intended groupoid. The reason for this choice is clear: we want to assign an arrow between two points of the base manifold to each diffeomorphism between their fibres which is consistent with the right action of the structure group. More precisely, let $a, b \in \mathcal{B}$. The diffeomorphisms $z : \pi_P^{-1}(a) \to \pi_P^{-1}(b)$ we are referring to are those satisfying:

$$z(pg) = z(p)g, \quad \forall p \in \pi_P^{-1}(a), \ g \in \mathcal{G}. \tag{5.5}$$

These are exactly the diffeomorphisms generated by pairing any two points $p \in \pi_P^{-1}(a),\ q \in \pi_P^{-1}(b)$ (i.e., one from each fibre) and then assigning to pg the point qg. As g runs through the group the diffeomorphism is generated. In other words, we assign the same arrow to the pair $(p, q) \in \mathcal{P} \times \mathcal{P}$ as we assign to (pg, qg). Hence the quotient $\mathcal{Z} = (\mathcal{P} \times \mathcal{P})/\mathcal{G}$.

The remainder of the construction is straightforward. The inverse of the arrow with the pair (p, q) as representative, is the arrow represented by the pair (q, p). The identity at $b \in \mathcal{B}$ is represented by any pair (p, p) with $\pi_P(p) = b$. The source and target maps are simply: $\alpha(p, q) = \pi_P(p)$ and $\beta(p, q) = \pi_P(q)$. Finally, the composition of arrows is effected by composition of maps. More carefully, let (p, q) be an arrow between a and b, and let (r, s) be an arrow between b and c, with $a, b, c \in \mathcal{B}$. We proceed to change the representative of the equivalence class (r, s) by applying the right action of the unique $g \in \mathcal{G}$ such that $rg = q$. We obtain the new representative of the *same* arrow as (q, sg). The successive application of the arrows is the arrow from a to c whose representative is the pair (p, sg) (the common intermediate element being cancelled out, as it were).

Chapter 2

Physical illustrations

2.1 The configuration space of a mechanical system

Lagrange's (1736-1813) conception of Mechanics was purportedly purely analytical. In the Preface to the first edition of his Mécanique Analytique[1] he explicitly states that: "On ne trouvera point de Figures dans cet Ouvrage. Les méthodes que j'y expose ne demandent ni constructions, ni raisonnements géométriques ou méchaniques, mais seulemnet des opérations algébriques, assujeties à une marche régulière et uniforme. Ceux qui aiment l'Analyse verront avec plaisir la Méchanique en devenir une nouvelle branche, et me sauront gré d'en avoir étendu ainsi le domain." Nevertheless, it is not an exaggeration to say that in laying down the foundations of Analytical Mechanics Lagrange was actually inaugurating the differential geometric approach to Physics. In Lagrange's view, a mechanical system was characterized by a finite number n of *degrees of freedom* to each of which a *generalized coordinate* is assigned. A *configuration* of the system is thus identified with an ordered n-tuple of real numbers. But, is this assignment unique? And, anyway, what are these numbers coordinates of?

Consider the classical example of a (rigid) double pendulum oscillating in a vertical plane under gravity, as shown in Figure 2.1. Clearly, this system can be characterized by two independent degrees of freedom. If we were to adopt as generalized coordinates the horizontal displacements, x_1 and x_2, of the two masses from, say, the vertical line through the point of suspension, we would find that to an arbitrary combination of these two numbers, there may correspond as many as 4 different configurations. If, to avoid this problem, we were to adopt as generalized coordinates the angular deviations θ_1 and θ_2, we would find that a given configuration can be characterized by an infinite number of combinations of values of these coordinates, due to the additive freedom of 2π. If we attempt to solve this problem by limiting the range of these coordinates to the interval $[0, 2\pi)$, we lose continuity of the representation (since two neighbouring configurations would correspond to very distant values of the coordinates).

Let us, therefore, go against Lagrange's own advice and attempt to draw a mental picture of the geometry of the situation. Since the first mass (attached to the main point of suspension) is constrained to move along a circle, thus constituting a simple pendulum, we conclude that its configurations can be homeomorphically mapped onto a circumference (or the perimeter of a square or of any other closed curve in the plane). We say that this circumference is the *configuration space* of a simple pendulum.

[1] Lagrange (1788), *Mécanique Analitique* [sic], chez la Veuve Desaint, Libraire, Paris.

Now, the second mass can describe a circumference around any position of the first. It is not difficult to conclude that the configuration space of the double pendulum is given by the surface of a torus. Now that this basic geometric (topological) question has been settled, we realize that an atlas of this torus must consist of several charts. But the central conceptual gain of the geometrical approach is that the configuration space of a mechanical system, whose configurations are defined with continuity in mind, can be faithfully represented by a unique topological manifold, up to a homeomorphism.

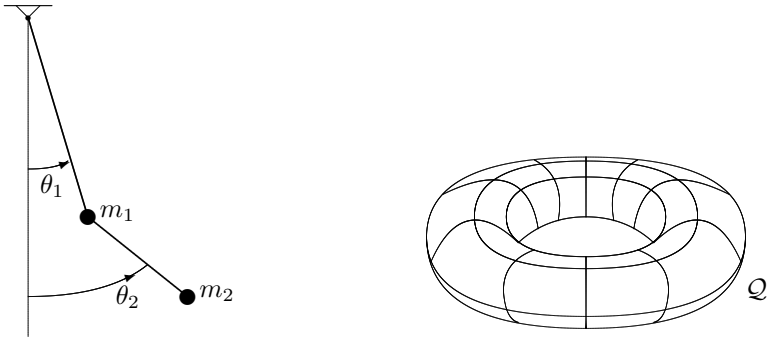

Figure 2.1: The plane double pendulum and its configuration manifold

2.2 Local symmetries of constitutive laws

Think of a material point as a small (infinitesimal) die in $\mathbb{R}^3$ that can be deformed into small arbitrary parallelepipeds by means of regular linear maps. These maps are, therefore, represented by non-singular matrices $\mathbf{F}$. Consider now a scalar function of state, or *constitutive function* ψ, such as a stored elastic energy, that depends exclusively on $\mathbf{F}$ via a *constitutive equation* or *constitutive law*

$$\psi = \psi(\mathbf{F}). \tag{2.1}$$

The *general linear group* $GL(3;\mathbb{R})$, that is, the (topological) group of all non-singular 3×3 real matrices $\mathbf{G}$, acts to the right on the collection $\mathcal{F}$ of all possible constitutive equations of the form (2.1) according to the following prescription:

$$R_{\mathbf{G}}\psi(\mathbf{F}) = \psi(\mathbf{FG}). \tag{2.2}$$

An element $\mathbf{G} \in GL(3;\mathbb{R})$ is a *symmetry* of the constitutive law ψ if

$$R_{\mathbf{G}}\psi = \psi, \tag{2.3}$$

or, more explicitly, if

$$\psi(\mathbf{FG}) = \psi(\mathbf{F}) \tag{2.4}$$

identically for all $\mathbf{F} \in GL(3;\mathbb{R})$. We also say that the constitutive equation $\psi = \psi(\mathbf{F})$ is *invariant* under the right action of the element $\mathbf{G}$ of the general linear group $GL(3;\mathbb{R})$. It is not difficult to show that the collection of symmetries of a given constitutive law ψ is a subgroup $\mathcal{G}_\psi$ of $GL(3;\mathbb{R})$, called the *material symmetry group* of ψ.

Clearly, the unit element of $GL(3;\mathbb{R})$, namely the unit matrix $\mathbf{I}$, is a trivial symmetry of all constitutive laws. On the other hand, it is not difficult to construct examples of constitutive laws that have non-trivial symmetries. For example, any function of the determinant of $\mathbf{F}$ is invariant under the action of any matrix $\mathbf{G}$ with unit determinant.

We ask now whether the right action of $GL(3;\mathbb{R})$ on $\mathcal{F}$ is transitive. The answer is clearly negative, since otherwise all constitutive equations would be identical to each other. Is the action free? Again, the answer is negative since, as shown above, there exist constitutive equations with non-trivial symmetries. Finally, one might have expected that the action be at least effective. That would have meant that if an element $\mathbf{G} \in GL(3;\mathbb{R})$ leaves all constitutive laws invariant then it must necessarily be the group unit $\mathbf{I}$. It can be shown, however, that if the set $\mathcal{F}$ is restricted to those constitutive equations abiding by the *general principle of material frame indifference*, then all constitutive laws are invariant under the action of $-\mathbf{I}$, in which case the action is not even effective.

2.3 Space-time

2.3.1 Aristotelian space-time

We may think separately of time as a 1-dimensional manifold $\mathcal{Z}$ (the time line) and of space as a 3-dimensional manifold $\mathcal{P}$. Nevertheless, as soon as we try to integrate these two entities into a single space-time manifold $\mathcal{S}$, whose points represent *events*, we realize that there are several possibilities. The first possibility that comes to mind is what we may call *Aristotelian space-time*, whereby time and space have independent and absolute meanings. Mathematically, this idea corresponds to the product:

$$\mathcal{S}_A = \mathcal{Z} \times \mathcal{P}, \tag{3.1}$$

where $\times$ denotes the Cartesian product. Recall that the *Cartesian product* of two sets is the set formed by all ordered pairs such that the first element of the pair belongs to the first set and the second element belongs to the second set. Thus, the elements s of $\mathcal{S}_A$, namely the events, are ordered pairs of the form (t, p), where $t \in \mathcal{Z}$ and $p \in \mathcal{P}$. In other words, for any given $s \in \mathcal{S}_A$, we can determine independently its corresponding temporal and spatial components. In mathematical terms, we say that the 4-dimensional (product) manifold $\mathcal{S}_A$ is endowed with two *projection maps*:

$$\pi_1 : \mathcal{S}_A \longrightarrow \mathcal{Z}, \tag{3.2}$$

and

$$\pi_2 : \mathcal{S}_A \longrightarrow \mathcal{P}, \tag{3.3}$$

defined, respectively, by:

$$\pi_1(s) = \pi_1(t, p) := t, \tag{3.4}$$

and

$$\pi_2(s) = \pi_2(t, p) := p. \tag{3.5}$$

2.3.2 Galilean space-time

The physical meaning of the existence of these two natural projections is that any observer can tell independently whether two events are simultaneous and whether or not (regardless of simultaneity) they

have taken place at the same location in space. According to the principle of *Galilean relativity*, however, this is not the case. Two different observers agree, indeed, on the issue of simultaneity. They can tell unequivocally, for instance, whether or not two light flashes occurred at the same time and, if not, which preceded which and by how much. Nevertheless, in the case of two non-simultaneous events, they will in general disagree on the issue of position. For example, an observer carrying a pulsating flashlight, will interpret successive flashes as happening always 'here', while an observer receding uniformly from the first will reckon the successive flashes as happening farther and farther away as time goes on. Mathematically, this means that we would like to get rid of the non-physical second projection (the spatial one) while preserving the first projection.

We would like, accordingly, to construct an entity that looks like $\mathcal{S}_A$ for each observer, but which is a different version of $\mathcal{S}_A$, so to speak, for different observers. This delicate issue can be handled as follows. We define space-time as a 4-dimensional manifold $\mathcal{S}$ endowed with a projection map:

$$\pi : \mathcal{S} \longrightarrow \mathcal{Z}, \tag{3.6}$$

together with a collection of smooth and (smoothly) invertible maps:

$$\phi : \mathcal{S} \longrightarrow \mathcal{S}_A, \tag{3.7}$$

onto the naive Aristotelian space-time $\mathcal{S}_A$. Each of these maps, called a *trivialization* and potentially representing an observer, cannot be completely arbitrary, in a sense that we will now explain.

Fix a particular point of time $t \in \mathcal{Z}$ and consider the inverse image $\mathcal{S}_t = \pi^{-1}(\{t\})$. We call $\mathcal{S}_t$ the *fibre of* $\mathcal{S}$ *at* t. Recall that the *inverse image* of a subset of the range of a function is the collection of all the points in its domain that are mapped to points in that subset. With this definition in mind, the meaning of $\mathcal{S}_t$ is the collection of all events that may happen at time t. We clearly want this *collection* to be the same for all observers, a fact guaranteed by the existence of the projection map π. Different observers will only differ in that they will attribute possibly different locations to events in this fibre. Therefore, we want the maps ϕ to be *fibre preserving* in the sense that each fibre of $\mathcal{S}$ is mapped to one and the same fibre in $\mathcal{S}_A$. In other words, we don't want to mix in any way whatsoever the concepts of space and time. We require, therefore, that the image of each fibre in $\mathcal{S}$ by each possible ϕ be exactly equal to a fibre of $\mathcal{S}_A$. More precisely, for each $t \in \mathcal{Z}$ we insist that:

$$\phi(\mathcal{S}_t) = \pi_1^{-1}(\{t\}). \tag{3.8}$$

A manifold $\mathcal{S}$ endowed with a projection π onto another manifold $\mathcal{Z}$ (called the *base manifold*) and with a collection of smooth invertible fibre-preserving maps onto a product manifold $\mathcal{S}_A$ (of the base times another manifold $\mathcal{P}$) is known as a *fibre bundle*. Note that the fibres of $\mathcal{S}_A$ by π_1 are all exact copies of $\mathcal{P}$. We say that $\mathcal{P}$ is the *typical fibre* of $\mathcal{S}$. A suggestive pictorial representation of these concepts is given in Figure 2.2.

Notice in Figure 2.2 how the fibres are shown hovering above (rather than touching) the base manifold. This device is used to suggest that, although each fibre is assigned to a specific point of the base manifold, the fibre and the base do not have any points in common, nor is there any preferential point in the fibre (such as a zero). Quite apart from the ability of Differential Geometry to elicit simple mental pictures to describe very complex objects, such as a fibre bundle, another important feature is that it uses the minimal amount of structure necessary. In the case of the space-time bundle, for instance, notice that we have not made any mention of the fact that there is a way to measure distances in space and a way to measure time intervals. In other words, what we have presented is what might be called a *proto-Galilean*

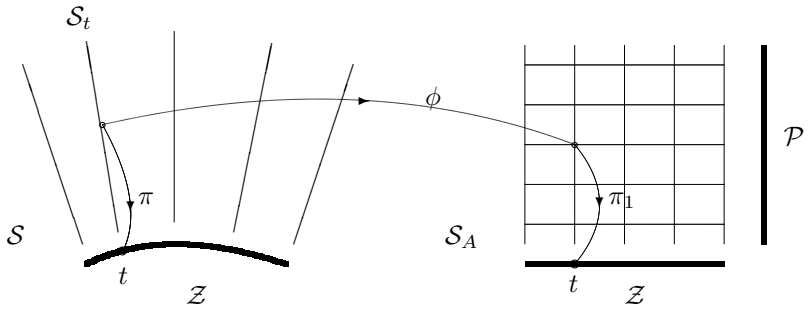

Figure 2.2: A fibre bundle

space-time, where the notion of simultaneity has a physical (and geometrical) meaning. Beyond that, we are now in a position to impose further structure either in the base manifold, or in the typical fibre, or in both. Similarly, restrictions can be placed on the maps ϕ (governing the change of observers). In classical *Galilean space-time*, the fibre $\mathcal{P}$ has the structure of an *affine space* (roughly a vector space without an origin). Moreover, this vector space has a distinguished *dot product*, allowing to measure lengths and angles. Such an affine space is called a *Euclidean space*. The time manifold $\mathcal{Z}$ is assumed to have a Euclidean structure as well, albeit one-dimensional. Physically, these structures mean that there is an observer-invariant way to measure distances and angles in space (at a given time) and that there is also an observer-invariant way to measure intervals of time. We say, accordingly, that Galilean space-time is an *affine bundle*. In such a fibre bundle, not only the base manifold and the typical fibre are affine spaces, but also the functions ϕ are limited to affine maps. These are maps that preserve the affine properties (for example, parallelism between two lines). In the case of Euclidean spaces, the maps may be assumed to preserve the metric structure as well.

2.3.3 Observer transformations

Having identified an observer with a trivialization ϕ, we can consider the notion of *observer transformation*. Let $\phi_1 : \mathcal{S} \to \mathcal{S}_A$ and $\phi_2 : \mathcal{S} \to \mathcal{S}_A$ be two trivializations. Since each of these maps is, by definition, invertible and fibre preserving, the composition:

$$\phi_{1,2} = \phi_2 \circ \phi_1^{-1} : \mathcal{S}_A \to \mathcal{S}_A, \tag{3.9}$$

is a well-defined fibre-preserving map from $\mathcal{S}_A$ onto itself. It represents the transformation from observer number 1 to observer number 2. Because of fibre preservation, the map $\phi_{1,2}$ can be seen as a smooth collection of time-dependent maps $\tilde{\phi}^t_{1,2}$ of the typical fibre $\mathcal{P}$ onto itself, as shown schematically in Figure 2.3. In Galilean space-time proper, we limit these maps to affine maps that preserve the orientation and the metric (Euclidean) structure of the typical fibre $\mathcal{P}$ (which can be seen as the usual 3-dimensional Euclidean space).

Among all such maps $\tilde{\phi}^t_{1,2} : \mathcal{P} \to \mathcal{P}$, it is possible to distinguish some that not only preserve the Euclidean structure but also represent changes of observers that travel with respect to each other at a fixed inclination (i.e., without angular velocity) and at a constant velocity of relative translation. Observers related in this way are said to be *inertially related*. It is possible, accordingly, to divide the collection of all observers into equivalence classes of inertially related observers. Of all these inertial classes, Isaac

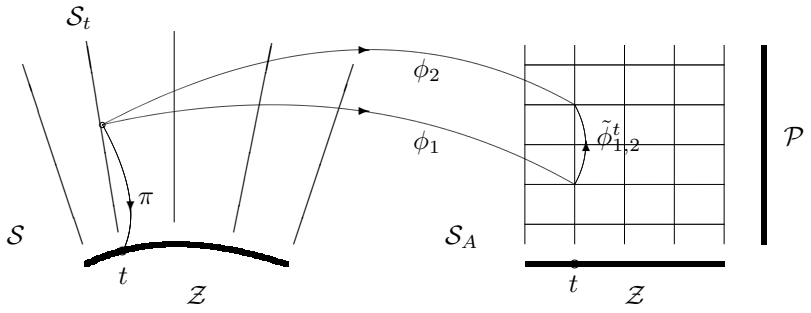

Figure 2.3: Observer transformation

Newton declared one to be privileged above all others[2]. This is the class of *inertial observers*, for which the laws of Physics acquire a particularly simple form.

Notice that in the case of Galilean space-time, a crosee section represents a *world line* or, more classically, a *trajectory* of a particle.

2.3.4 Relativistic space-time

The revolution brought about by the theory of Relativity (both in its special and general varieties) can be said to have destroyed the bundle structure altogether. In doing so, it in fact simplified the geometry of space-time, which becomes just a 4-dimensional manifold $\mathcal{S}_R$. On the other hand, instead of having two separate metric structures, one for space and one for time, Relativity assumes the existence of a space-time metric structure that involves both types of variables into a single construct. This type of metric structure is what Riemann had already considered in his pioneering work on the subject, except that Relativity (so as to be consistent with the Lorentz transformations) required a metric structure that could lead both to positive and to negative squared distances between events, according to whether or not they are reachable by a ray of light. In other words, the metric structure of Relativity is not positive definite. By removing the bundle structure of space time, Relativity was able to formulate a geometrically simpler picture of space time, although the notion of simplicity is in the eyes of the beholder.

To summarize, the theory of Relativity is simpler than its Classical counterpart from at least the following point of view: the structure of relativistic space-time is less involved than that of Galilean space-time. The extra structure in the latter is provided by the notion of *absolute simultaneity*. Starting from a 4-dimensional manifold of events, Classical Physics assumes that all observers agree on whether or not two events happened simultaneously, regardless of their locations. As a consequence, a time-projection operator arises naturally in this context. As a result, a Physics that abides by the principle of absolute simultaneity must of necessity be formulated upon a space-time manifold that has the structure of a fibre bundle, the base manifold being a 1-dimensional manifold. The typical fibre, representing space, is a 3-dimensional manifold. In the Galilean formulation, this typical fibre is the Euclidean space $\mathbb{E}^3$. The structure group is the group of Galilean transformations of $\mathbb{E}^3$ (those preserving Euclidean length). An observer is a bundle trivialization. In contrast, in Relativity, space-time is just a 4-dimensional manifold endowed with a special metric structure.

[2]This appears to be the meaning of Newton's first law of motion.

2.4 Microstructure

2.4.1 Shells

One of the many different ways to describe a *shell* in structural engineering is to regard it as the product bundle of a two-dimensional manifold $\mathcal{B}$ times the open (or sometimes closed) segment $\mathcal{F} = (-1, 1) \in \mathbb{R}$. The base manifold is known as the *middle surface* while the fibre conveys the idea of thickness, eventually responsible for the bending stiffness of the shell. The fact that this is a product bundle means that one can in a natural way identify corresponding locations throughout the thickness at different points of the middle surface. Thus, two points of the shell standing on different points of the middle surface can be said to correspond to each other if they have the same value of the second projection. This fact can be interpreted as being on the same side of the middle surface and at the same fraction of the respective thicknesses.

2.4.2 General microstructure

In a more general context, we can consider three-dimensional bodies for which the usual kinematic degrees of freedom are supplemented with extra (internal) degrees of freedom intended to describe a physically meaningful counterpart. This idea, going at least as far back as the pioneering work of the Cosserat brothers[3], applies to diverse materials, such as liquid crystals and granular media. The base manifold represents the matrix, or *macromedium*, while the fibres represent the *micromedium* (the elongated molecules or the grains, as the case may be). An example of this situation is provided by an everyday material such as concrete, which is formed by embedding in a cement matrix an aggregate consisting of stones whose size is relatively large when compared with the grains of cement. Each of these stones can then be considered as a *micromedium*. In a continuum model we expect to have these micromedia continuously assigned to each point of the matrix, thus generating a fibre bundle, whose typical fibre is the micromedium[4]. In contradistinction with the case of the shell, there is no canonical correspondence between points belonging to micromedia attached at different points of the macromedium.

2.5 Material uniformity

2.5.1 An imprecisely defined material body

In Section 2.2 we introduced the intuitive idea of a material point as a small die of material that can be subjected to linear deformations and whose constitutive response is governed by one or more constitutive equations. Following this imprecise intuitive line of thought, we can consider a sort of "continuous collection" of such material points[5] and regard the resulting entity as a material body. Each of the constituent material points is endowed with its own constitutive law and, if we cavalierly denote by $\mathbf{X}$ a running three-dimensional variable indicating the location of the body $\mathcal{B}$, we obtain the constitutive law

[3] Cosserat E, Cosserat F (1909), *Théorie des corps déformables*, Hermann et Fils, Paris.

[4] This example has been chosen for its graphical clarity. In the actual practice of Civil Engineering it is rare to find that concrete is treated in such a degree of detail. Instead, the contribution of the micromedium is averaged or *homogenized* into a supposedly equivalent ordinary macromedium

[5] This notion will be made more precise when we define a differentiable manifold and its tangent bundle.

of the body as some function

$$\psi = \psi(\mathbf{F}, \mathbf{X}). \tag{5.1}$$

2.5.2 Distant versus local symmetry

An illuminating example of distant, as opposed to local, symmetries is suggested by the tiling of a bathroom floor.[6] Each square tile has a symmetry group consisting of certain rotations and reflections. But it is also intuitively recognized that the floor as a whole has a repetitive pattern and, therefore, some extra symmetry. Because the floor is not infinite, however, we cannot describe all of these extra symmetries by means of a group of global transformations of the plane, such as translations. The notion of groupoid circumvents this problem. In the case of a material body, the fact that two distant points are made of the same material should be understood as an extra degree of symmetry that the body possesses, just as in the case of the bathroom floor, where distant tiles happen to have the same shape. This analogy should not be pushed too far, but it serves to trigger a useful picture and to understand the unifying role that the concept of groupoid plays in terms of encompassing all types of symmetries.

2.5.3 Material isomorphisms

We want to formalize the answer to the question: Are two material points $\mathbf{X}_1$ and $\mathbf{X}_2$ made of the same material? We reason that for this to be so, the only possible difference between the local constitutive equations $\psi(\mathbf{F}, \mathbf{X}_1)$ and $\psi(\mathbf{F}, \mathbf{X}_2)$ must be a fixed *transplant* represented by some matrix $\mathbf{P}_{12}$ such that

$$\psi(\mathbf{F}, \mathbf{X}_2) = \psi(\mathbf{F}\mathbf{P}_{12}, \mathbf{X}_1), \tag{5.2}$$

identically for all deformations $\mathbf{F}$. Indeed, in this case we would agree that the responses of the two points are exactly the same except for the fact that the die at point $\mathbf{X}_2$ is a rotated or otherwise distorted version of the die at point $\mathbf{X}_1$. In the standard Continuum Mechanics terminology, such a material transplant is known as a *material isomorphism*. It is not difficult to verify that material isomorphism is an equivalence relation. A body is said to be *materially uniform* if all its points are mutually materially isomorphic.

Notice that a material point is trivially isomorphic to itself (via the identity map), but it may also be non-trivially so (via a non-trivial material symmetry). We thus see that a material isomorphism is a generalization of the notion of local material symmetry to encompass what we may call distant material symmetries of a material body.

2.5.4 The material groupoid

Given a material body, whether uniform or not, we can imagine an arrow drawn for every material isomorphism between two points, including the material symmetries, namely, the cases whereby the source and target points coincide. In this way, without much further ado, we conclude that every material body with a specified constitutive law gives rise to a groupoid, which we shall call the *material groupoid* of the body. In case the body is materially uniform, we obtain a transitive groupoid. If, in addition, the constitutive equation is continuous in $\mathbf{X}$, we obtain a transitive topological groupoid.[7]

[6] See Weinstein A (2000), Groupoids: Unifying Internal and External Symmetry. A tour through Examples, *Notices of the American Mathematical Society* **43**, 744-752.

[7] For the use of groupoids in the theory of material uniformity see: Epstein M and de León M(1998), Geometrical theory of uniform Cosserat media, *Journal of Geometry and Physics* **26**, 127-170.

Consider now a non-uniform body. The material groupoid is still properly definable, except that it loses its transitivity. It may still preserve its continuity (namely, it may still be a topological groupoid). A good example of this last situation is provided by the so-called *functionally-graded materials*, which have continuously varying material properties tailored to specific applications. Under certain circumstances, however, the transitivity of the material groupoid of functionally graded materials can be restored by modifying the definition of material isomorphism[8].

2.5.5 Material principal bundles

A *material principal bundle* of a materially uniform body $\mathcal{B}$ as any one of the equivalent principal bundles that can be obtained from the material groupoid. Physically speaking, a material principal bundle is obtained by arbitrarily singling out a material point $\mathbf{X}_0$, called the *material archetype* and replacing the material transplants between arbitrary pairs of points by material *implants* $\mathbf{P}(\mathbf{X})$ from the archetype to each and every point $\mathbf{X}$, as shown in Figure 2.4. The constitutive equation of a uniform body thus conceived is given by:

$$\psi(\mathbf{F}, \mathbf{X}) = \bar{\psi}(\mathbf{F}\mathbf{P}(\mathbf{X})), \tag{5.3}$$

where we have indicate by $\bar{\psi}$ the constitutive law of the archetype.

We observe, however, that whereas the material groupoid always exists (whether or not the body is uniform), the material principal bundles can only be defined when the body is smoothly uniform. Then, and only then, we have a transitive topological groupoid to work with. In conclusion, although both geometrical objects are suitable for the description of the material structure of a body, the groupoid representation is the more faithful one, since it is unique and universal.

The structure group of a material principal bundle is, according to the previous construction, nothing but the material symmetry group of the archetype. As expected, it controls the degree of freedom available in terms of implanting this archetype at the points of the body.

A material principal bundle may, or may not, admit (global) cross sections. If it does, the body is said to be *globally uniform*. This term is slightly misleading, since uniformity already implies that *all* the points of the body are materially isomorphic. Nevertheless the term conveys the sense that the material isomorphisms can be prescribed smoothly in a single global chart of the body (which, by definition, always exists). Put in other terms, the existence of a global section implies (in a principal bundle, as we know) that the principal bundle is trivializable. A cross section of a principal bundle establishes, through the right action of the structure group, a global isomorphism between the fibres, also called a *distant parallelism*. In our context, this property will be called a *material parallelism*. If the structure group is discrete, the material parallelism is unique. Moreover, if the material symmetry group consists of just the identity, a uniform body must be globally uniform.

[8] See: Epstein M, de León M (2000), Homogeneity Without Uniformity: Toward a Mathematical Theory of Functionally Graded Materials, *International Journal of Solids and Structures* **37**, 7577-7591. Also: Epstein M, Elżanowski M (2007) *Material inhomogeneities and their evolution*, Springer.

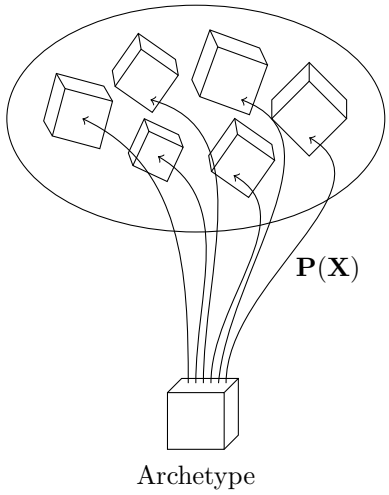

Figure 2.4: Generating a material principal bundle

Chapter 3

Differential constructs

3.1 Differentiable manifolds

Topological manifolds provide the most general arena for the definition of continuous functions. Continuity alone, however, may not be enough to formulate physical problems in a quantitative manner. Indeed, experience with actual dynamical and field theories of Mechanics and Electromagnetism, to name only the classical theories, has taught us to expect that the various phenomena are governed by ordinary or partial differential equations. These theories, therefore, must be formulated on a substratum that has more structure than a topological manifold, namely, an entity that allows for the definition of differentiable functions. Differentiable manifolds are the natural generalization of topological manifolds to handle differentiability.

Since a topological space does not possess in itself enough structure to sustain the notion of differentiability, the key to the generalization of a topological manifold is to be found in restrictions imposed upon the transition functions, which are clearly defined in $\mathbb{R}^n$. Two charts, $(\mathcal{U}_\alpha, \phi_\alpha)$ and $(\mathcal{U}_\beta, \phi_\beta)$, of a topological manifold $\mathcal{M}$ are said to be C^k*-compatible*, if the transition functions $\phi_{\alpha\beta}$ and $\phi_{\beta\alpha}$, as defined in Equation (2.1), are of class C^k. In terms of the representation (2.2), this means that all the partial derivatives up to and including the order k exist and are continuous. By convention, a continuous function is said to be of class C^0 and a *smooth* function is of class C^∞.

In a topological manifold, all charts of all possible atlases are automatically C^0-compatible. An *atlas of class C^k of a topological manifold $\mathcal{M}$* is an atlas whose charts are C^k-compatible. Two atlases of class C^k are *compatible* if each chart of one is compatible with each chart of the other. The union of compatible C^k-atlases is a C^k atlas. Given a C^k atlas, one can define the corresponding *maximal compatible atlas of class C^k* as the union of all atlases that are C^k-compatible with the given one. A maximal atlas, thus, contains all its compatible atlases.

Definition 3.1.1 *An n-dimensional differentiable manifold of class C^k is an n-dimensional topological manifold $\mathcal{M}$ together with a maximal atlas of class C^k. For $k = 0$ one recovers the topological manifold. The C^∞ case delivers a smooth manifold, or simply a manifold.*

A maximal C^k-atlas is also called a C^k*-differentiable structure*. Thus, a C^k-manifold is a topological manifold with a C^k-differentiable structure. For the particular case of $\mathbb{R}^n$, we can choose the *canonical*

atlas consisting of a single chart (the space itself) and the identity map. The induced C^∞-differentiable structure is the *standard differentiable structure of* $\mathbb{R}^n$.

A differentiable manifold is *oriented* if it admits an atlas, called an *oriented atlas*, such that all the transition functions have a positive Jacobian determinant. Two oriented atlases are either compatible or every transition function between charts of the two atlases has a negative determinant. An *oriented manifold* is an orientable manifold with and oriented maximal atlas. In other words, only those coordinate transformations that preserve the orientation are permitted.

Given two differentiable manifolds $\mathcal{M}$ and $\mathcal{N}$ of dimensions m and n, respectively, we define the $(m+n)$-dimensional *product manifold* by endowing the Cartesian product $\mathcal{M} \times \mathcal{N}$ with the atlas made of all the Cartesian products of charts of an atlas of $\mathcal{M}$ and an atlas of $\mathcal{N}$.

3.1.1 Differentiable maps

Let $\mathcal{M}$ and $\mathcal{N}$ be (smooth) manifolds of dimensions m and n, respectively. A continuous map $f : \mathcal{M} \to \mathcal{N}$ is *differentiable of class* C^k at a point $p \in \mathcal{M}$ if, using charts $(\mathcal{U}, \phi)$ and $(\mathcal{V}, \psi)$ belonging to the respective maximal atlases of $\mathcal{M}$ and $\mathcal{N}$, the local coordinate representation $\hat{f}$ of f, as defined in Equation (2.3), is of class C^k at $\phi(p) \in \mathbb{R}^m$. This definition is independent of chart, since the composition of differentiable maps in $\mathbb{R}^m$ is differentiable. Notice how the notion of differentiability within the manifolds has been cleverly deflected to the charts.

Maps of class C^∞ are said to be *smooth maps*, to which we will confine our analysis from now on. In the special case $\mathcal{N} = \mathbb{R}$, the map $f : \mathcal{M} \to \mathbb{R}$ is called a (real) *function*. When, on the other hand, $\mathcal{M}$ is an open interval $H = (a, b)$ of the real line, the map $\gamma : H \to \mathcal{N}$ is called a *(parametrized) curve in* $\mathcal{N}$. The name *diffeomorphism* is reserved for the case in which $\mathcal{M}$ and $\mathcal{N}$ are of the same dimension and both f and its (assumed to exist) inverse f^{-1} are smooth. Two manifolds of the same dimension are said to be *diffeomorphic* if there exists a diffeomorphism between them.

3.1.2 Smooth structures

Having introduced the concept of differentiable manifold, we are in a position to elevate all the topological constructs introduced so far from the status of continuity to a status of smoothness. Thus, a topological group whose underlying set is a differentiable (smooth) manifold will be called a *Lie group* if the operations of multiplication and inversion are smooth. A smooth fibre bundle is a topological fibre bundle in which the base manifold and the typical fibre are smooth manifolds and the structure group is a Lie group. The projection map is smooth (technically, a surjective submersion). An important smooth fibre bundle, canonically defined for any given smooth manifold, is the principal *bundle of linear frames*, which we will study later in some detail. Finally, a topological groupoid in which the total space and the base space are smooth manifolds, both projections are surjective submersions and all operations are smooth is called a *Lie groupoid*.

3.1.3 Tangent vectors

Let H be an open interval of the real line and, without loss of generality, assume that $0 \in H$. Consider the collection of all (smooth) curves $\gamma : H \to \mathcal{M}$ such that $\gamma(0) = p$. Our aim is to define the notion of

tangency of two such curves at p, an aim that we achieve by using the technique of deflecting to charts. Indeed, if $(\mathcal{U}, \phi)$ is a chart containing p, the composition $\hat{\gamma} = \phi \circ \gamma : H \to \mathbb{R}^m$ is a curve in $\mathbb{R}^m$, where m is the dimension of $\mathcal{M}$. The coordinate expression of $\hat{\gamma}$ is given by m smooth real functions $x^i = \gamma^i(t)$, where t is the natural coordinate of $\mathbb{R}$ and $i = 1, ..., m$. We say that two curves, γ_1 and γ_2, in our collection are *tangent at p* if:

$$\left.\frac{d\gamma_1^i}{dt}\right|_{t=0} = \left.\frac{d\gamma_2^i}{dt}\right|_{t=0}, \qquad i = 1, ..., m. \tag{1.1}$$

It is a simple matter to verify that this definition is independent of chart.

Noting that tangency at p is an equivalence relation, we define a *tangent vector at p* as an equivalence class of (smooth, parametrized) curves tangent at p. A tangent vector is thus visualized as what the members of a collection of tangent (parametrized) curves have in common. More intuitively, one may say that what these curves have in common is a small piece of a curve.

Let $f : \mathcal{M} \to \mathbb{R}$ be a (differentiable) function and let $\mathbf{v}$ be a tangent vector at $p \in \mathcal{M}$. Choosing any representative γ in the equivalence class $\mathbf{v}$, the composition $f \circ \gamma$ is a real-valued function defined on H. The *derivative of f at p along* $\mathbf{v}$ is defined as:

$$\mathbf{v}(f) = \left.\frac{d(f \circ \gamma)}{dt}\right|_{t=0}. \tag{1.2}$$

This notation suggests that a vector can be regarded as a linear operator on the collection of differentiable functions defined on a neighbourhood of a point. The linearity is a direct consequence of the linearity of the derivative. Not every linear operator, however, is a tangent vector because, by virtue of Equation (1.2), tangent vectors must also satisfy the Leibniz rule, namely, for any two functions f and g:

$$\mathbf{v}(fg) = f\mathbf{v}(g) + \mathbf{v}(f)g, \tag{1.3}$$

where, on the right-hand side, f and g are evaluated at p.

3.1.4 Tangent and cotangent spaces

The collection $T_p\mathcal{M}$ of all the tangent vectors at $p \in \mathcal{M}$ is called the *tangent space to the manifold at p*. It is not difficult to show that tangent vectors at a point p satisfy all the conditions of a vector space if we define their addition and multiplication by a scalar in the obvious way (for example, by using chart components). In other words, $T_p\mathcal{M}$ is a vector space. To find its dimension, we choose a local chart $(\mathcal{U}, \phi)$ with coordinates $x^1, ..., x^m$, such that the point p is mapped to the origin of $\mathbb{R}^m$. The inverse map ϕ^{-1}, when restricted to the natural coordinate lines of $\mathbb{R}^m$, delivers m curves at p. Each of these curves, called a *coordinate line in* $\mathcal{U}$, defines a tangent vector, which we suggestively denote by $(\partial/\partial x^i)_p$. It can be shown that these vectors constitute a basis of $T_p\mathcal{M}$, called the *natural basis associated with the given coordinate system*. The dimension of the tangent space at each point of a manifold is, therefore, equal to the dimension of the manifold itself. The *cotangent space at p*, denoted by $T_p^*\mathcal{M}$, is defined as the dual space of $T_p\mathcal{M}$.

3.1.5 The tangent and cotangent bundles

If we attach to each point p of an m-dimensional manifold $\mathcal{M}$ its tangent space $T_p\mathcal{M}$, we obtain, intuitively speaking, a $2m$-dimensional entity, which we denote by $T\mathcal{M}$ called the *tangent bundle* of $\mathcal{M}$. A crude

visualization of this entity can be gathered when $\mathcal{M}$ is a 2-sphere, such as a globe, at each point of which we have stuck a postal stamp or a paper sticker. The tangent bundle is not the globe itself but rather the collection of the stickers. This collection of tangent spaces, however, has the property that it *projects* on the original manifold. In our example, each sticker indicates the point at which it has been attached. In other words, the set $T\mathcal{M}$ is endowed, by construction, with a *projection map* τ on the *base manifold* $\mathcal{M}$. More explicitly, a typical point of $T\mathcal{M}$ consists of a pair $(p, \mathbf{v}_p)$, where $p \in \mathcal{M}$ and $\mathbf{v}_p \in T_p\mathcal{M}$. The projection map:

$$\tau : T\mathcal{M} \to \mathcal{M} \tag{1.4}$$

is given by the assignation:

$$\tau(p, \mathbf{v}_p) = p. \tag{1.5}$$

To see that the set $T\mathcal{M}$ can be regarded as a manifold, we construct explicitly an atlas out of any given atlas of the base manifold. Let $(\mathcal{U}, \phi)$ be a chart in $\mathcal{M}$ with coordinates $x^i, ..., x^m$. Adopting, as we may, the natural basis $(\partial/\partial x^i)_p$ of $T_p\mathcal{M}$ at each point $p \in \mathcal{U}$, we can identify each vector $\mathbf{v}_p$ with its components v^i_p. Put differently, we assign to each point $(p, \mathbf{v}_p) \in \tau^{-1}(\mathcal{U}) \subset T\mathcal{M}$ the $2m$ numbers $(x^1, ..., x^m, v^1, ..., v^m)$, namely, a point in $\mathbb{R}^{2m}$. We have thus obtained a coordinate chart on $\tau^{-1}(\mathcal{U})$. It is now a formality to extend this construction to a whole atlas of $T\mathcal{M}$ and to show that $T\mathcal{M}$ is a differentiable manifold of dimension $2m$. In the terminology of general *fibre bundles*, the set $T_p\mathcal{M} = \tau^{-1}(p)$ is called the *fibre* at $p \in \mathcal{M}$. Since each fibre is an m-dimensional vector space, we say that the *typical fibre* of $T\mathcal{M}$ is $\mathbb{R}^m$.

Upon a coordinate transformation represented by Equation (2.2), the components $\hat{v}^i$ of a vector $\mathbf{v}$ at p in the new natural basis $(\partial/\partial y^i)_p$ are related to the old components v^i in the basis $(\partial/\partial x^i)_p$ by the formula:

$$\hat{v}^i = \left(\frac{\partial y^i}{\partial x^j}\right)_p v^j, \tag{1.6}$$

while the base vectors themselves are related by the formula:

$$(\partial/\partial y^i)_p = \left(\frac{\partial x^j}{\partial y^i}\right)_p (\partial/\partial x^j)_p. \tag{1.7}$$

Comparing these two formulas, we conclude that the components of vectors behave *contravariantly*. In traditional treatments, it was customary to *define* tangent vectors as indexed quantities that transform contravariantly under coordinate changes.

A similar construction can be carried out by attaching to each point of a manifold $\mathcal{M}$ its cotangent space $T^*_p\mathcal{M}$ to obtain the set $T^*\mathcal{M}$, called the *cotangent bundle* of $\mathcal{M}$. A typical point of $T^*\mathcal{M}$ is a pair (p, ω_p), where $p \in \mathcal{M}$ and $\omega_p \in T^*_p\mathcal{M}$. The *projection map* $\pi : T^*\mathcal{M} \to \mathcal{M}$ is given by:

$$\pi(p, \omega_p) = p. \tag{1.8}$$

Given a chart, the local dual basis to the natural basis $(\partial/\partial x^i)_p$ is denoted by $(dx^i)_p$, with $i = 1, ..., m$. The covector $\omega_p \in T^*_p\mathcal{M}$ can be uniquely expressed as $\omega_p = \omega_i dx^i$, where the subscript p has been eliminated for clarity. Given a point $(p, \omega_p) \in \pi^{-1}(\mathcal{U}) \subset T^*\mathcal{M}$ we assign to it the $2m$ numbers $(x^1, ..., x^m, \omega_1, ..., \omega_m)$. In this way, it can be rigorously shown that $T^*\mathcal{M}$ is a manifold of dimension $2m$.

Upon a coordinate transformation, the components $\hat{\omega}_i$ of a covector ω at p transform according to:

$$\hat{\omega}_i = \left(\frac{\partial x^j}{\partial y^i}\right)_p \omega_j, \tag{1.9}$$

while the dual base vectors themselves are related by the formula:

$$dy^i = \left(\frac{\partial y^i}{\partial x^j}\right)_p dx^j. \tag{1.10}$$

The components of covectors behave *covariantly*.

3.1.6 The differential of a map

Given a differentiable map:

$$g : \mathcal{M} \to \mathcal{N} \tag{1.11}$$

between two manifolds, $\mathcal{M}$ and $\mathcal{N}$, of dimensions m and n, respectively, we focus attention on a particular point $p \in \mathcal{M}$ and its image $q = g(p) \in \mathcal{N}$. Let $\mathbf{v}_p \in T_p\mathcal{M}$ be a tangent vector at p and let $\gamma : H \to \mathcal{M}$ be one of its representative curves. The composite map:

$$g \circ \gamma : H \longrightarrow \mathcal{N} \tag{1.12}$$

is then a smooth curve in $\mathcal{N}$ passing through q. This curve (the image of γ by g) is, therefore, the representative of a tangent vector at q which we will denote $(g_*)_p(\mathbf{v}_p)$. The vector $(g_*)_p(\mathbf{v}_p)$ is independent of the representative curve γ chosen for $\mathbf{v}_p$. Moreover, $(g_*)_p$ is a linear map on vectors at p.

The map $(g_*)_p$ just defined is called the *differential of g at p*. It is a linear map between the tangent spaces $T_p\mathcal{M}$ and $T_{g(p)}\mathcal{N}$. Since this construction can be carried out at each and every point of $\mathcal{M}$, we obtain a map g_* between the tangent bundles, namely:

$$g_* : T\mathcal{M} \to T\mathcal{N}, \tag{1.13}$$

called the *differential of g*. Alternative notations for this map are: Dg and Tg, and it is also known as the *tangent map*. One should note that the map g_* includes the map g between the base manifolds, since it maps vectors at a point p linearly into vectors at the image point $q = g(p)$, and not just to any vector in $T\mathcal{N}$. It is, therefore, a *fibre-preserving map*. This fact is best illustrated in the following commutative diagram:

$$\begin{array}{ccc} T\mathcal{M} & \xrightarrow{g_*} & T\mathcal{N} \\ \Big\downarrow{\scriptstyle \tau_{\mathcal{M}}} & & \Big\downarrow{\scriptstyle \tau_{\mathcal{N}}} \\ \mathcal{M} & \xrightarrow{g} & \mathcal{N} \end{array} \tag{1.14}$$

where $\tau_{\mathcal{M}}$ and $\tau_{\mathcal{N}}$ are the projection maps of $T\mathcal{M}$ and $T\mathcal{N}$, respectively. The differential is said to *push forward* tangent vectors at p to tangent vectors at the image point $g(p)$.

In the particular case of a function $f : \mathcal{M} \to \mathbb{R}$, the differential f_* can be interpreted somewhat differently. Indeed, the tangent space $T_r\mathbb{R}$ can be trivially identified with $\mathbb{R}$ itself, so that f_* can be seen as a real-valued function on $T\mathcal{M}$. This function is denoted by $df : T\mathcal{M} \to \mathbb{R}$. The differential of a function satisfies the identity:

$$df(\mathbf{v}) = \mathbf{v}(f). \tag{1.15}$$

In local systems of coordinates x^i $(i = 1, ..., m)$ and y^α $(\alpha = 1, ..., n)$ around p and $g(p)$, respectively, the differential of g at p maps the vector with components v^i into the vector with components:

$$[(g_*)_p(\mathbf{v}_p)]^\alpha = \left(\frac{\partial g^\alpha}{\partial x^i}\right)_p v^i, \tag{1.16}$$

where $g^\alpha = g^\alpha(x^1, ..., x^n)$ is the coordinate representation of g in the given charts. The $(m \times n)$-matrix with entries $\{\left(\frac{\partial g^\alpha}{\partial x^i}\right)_p\}$ is the *Jacobian matrix* at p of the map g in the chosen coordinate systems. The *rank* of the Jacobian matrix is independent of the coordinates used. It is called the *rank of g at p*.

Let $f : \mathcal{N} \to \mathbb{R}$ be a differentiable function and let $g : \mathcal{M} \to \mathcal{N}$ be a differentiable map between manifolds. Then:

$$((g_*)_p \mathbf{v}_p)(f) = \mathbf{v}_p(f \circ g), \quad p \in \mathcal{M}. \tag{1.17}$$

The differential of a composition of maps is equal to the composition of the differentials. More precisely, if $g : \mathcal{M} \to \mathcal{N}$ and $h : \mathcal{N} \to \mathcal{P}$ are differentiable maps, then:

$$((h \circ g)_*)_p(\mathbf{v}_p) = (h_*)_{g(p)}((g_*)_p(\mathbf{v}_p)). \tag{1.18}$$

In coordinates, this formula amounts to the multiplication of the Jacobian matrices.

3.1.7 The linear frame bundle

The *bundle of linear frames*, $F\mathcal{B}$, of a base n-dimensional manifold $\mathcal{B}$ can be defined constructively in the following way. At each point $b \in \mathcal{B}$ we form the set $F_b\mathcal{B}$ of all ordered n-tuples $\{e\}_b = (e_1, ..., e_n)$ of linearly independent vectors e_i in $T_b\mathcal{B}$, namely, the set of all bases of $T_b\mathcal{B}$. Our total space will consist of all ordered pairs of the form $(b, \{e\}_b)$ with the obvious projection onto $\mathcal{B}$. The pair $(b, \{e\}_b)$ is called *a linear frame at b*. Following a procedure identical to the one used for the tangent bundle, we obtain that each basis $\{e\}_b$ is expressible uniquely as:

$$e_j = p^i{}_j \frac{\partial}{\partial x^i} \tag{1.19}$$

in a coordinate system x^i, where $\{p^i{}_j\}$ is a non-singular matrix. We conclude that the typical fibre in this case is $GL(n; \mathbb{R})$. But so is the structure group. Indeed, in another coordinate system, y^i, we have:

$$e_j = q^i{}_j \frac{\partial}{\partial y^i}, \tag{1.20}$$

where

$$q^i{}_j = \frac{\partial y^i}{\partial x^m} p^m_j = a^i{}_m p^m_j. \tag{1.21}$$

This is an instance of a *principal fibre bundle*, namely, a fibre bundle whose typical fibre and structure group coincide. The action of the group on the typical fibre is the natural left action of the group on itself. One of the interesting features of a principal bundle is that the structure group has also a natural *right* action *on the bundle itself*, and this property has been used in Section 1.4.1 to provide an alternative definition of principal bundles. In the case of $F\mathcal{B}$, for example, the right action is defined, in a given coordinate system x^i, by

$$R_a\{e\} = p^k_i a^i{}_j \frac{\partial}{\partial x^k}, \quad j = 1, ..., n, \tag{1.22}$$

which sends the basis (1.19) at b to another basis at b, i.e., the action is fibre-preserving. One can verify that this definition of the action is independent of the system of coordinates adopted. The principal bundle of linear frames of a manifold is associated to all the tensor bundles (see Section 3.3.3), including the tangent and the cotangent bundles, of the same manifold. By a direct application of the fundamental existence theorem, we know that the associated principal bundle is defined uniquely up to an equivalence. Many properties of bundles can be better understood by working first on the associated principal bundle.

3.1.8 Associated bundles

The concept of associated bundle has already been defined and used to introduce the notion of the principal bundle associated with any given fibre bundle. On the other hand, we have introduced an independent definition of principal bundles by means of the idea of a right action of a group on a given total manifold. We want now to show that this line of thought can be pursued to obtain another view of the collection of all (non-principal) fibre bundles associated with a given principal bundle.

As a more or less intuitive motivation for this procedure, it is convenient to think of the example of the principal bundle of linear frames $F\mathcal{B}$ of a manifold $\mathcal{B}$. We already know that this bundle is associated to the tangent bundle $T\mathcal{B}$. Consider now a pair (f, v), where $f \in F\mathcal{B}$ and $v \in T\mathcal{B}$, such that $\pi_P(f) = \pi(v) = b$. In other words, f and v represent, respectively, a basis and a vector of the tangent space at some point $b \in \mathcal{B}$. We can, therefore, identify v with its components on the linear frame f, namely, with an element of the typical fibre ($\mathbb{R}^n$) of $T\mathcal{B}$. If we consider now a pair $(\hat{f}, v)$, where v is the same as before but $\hat{f}$ is a new linear frame at b, the corresponding element of the typical fibre representing the *same* vector v changes. More explicitly, with an obvious notation, if $\hat{f}_j = a^i{}_j f_i$, then $v^i = a^i{}_j \hat{v}^j$ or: $\hat{v}^i = (a^{-1})^i{}_j v^j$. We conclude that to represent the *same object* under a change of frame, there needs to be some kind of compensatory action in the change of the components. The object itself (in this case, the tangent vector) can be identified with the collection (or equivalence class) of all pairs made up of a frame and a matrix related in this compensatory way. In terms of the group actions on the typical fibres, if $\hat{f} = R_a f$, then the representative r of the vector v in $\mathbb{R}^n$ changes according to $\hat{r} = L_{a^{-1}} r$. We may, therefore, think of a vector as an equivalence class of elements of the Cartesian product $\mathcal{G} \times \mathbb{R}^n$, corresponding to the following equivalence relation: $(g, r) \sim (\hat{g}, \hat{r})$ if, and only if, there exists $a \in \mathcal{G}$ such that $\hat{g} = ga$ and $\hat{r} = L_{a^{-1}} r$.

With the above motivation in mind, the following construction of a fibre bundle associated to a given principal bundle will seem less artificial than it otherwise would. We start from the principal bundle $(\mathcal{P}, \pi_P, \mathcal{B}, \mathcal{G}, \mathcal{G})$ and a manifold $\mathcal{F}$, which we want to construe as the typical fibre of a new fibre bundle $(\mathcal{C}, \pi, \mathcal{B}, \mathcal{F}, \mathcal{G})$ associated with $\mathcal{P}$. For this to be possible, we need to have an effective left action of $\mathcal{G}$ on $\mathcal{F}$, which we assume to have been given. To start off, we form the Cartesian product $\mathcal{P} \times \mathcal{F}$ and notice that the structure group $\mathcal{G}$ acts on it with a right action induced by its right action on $\mathcal{P}$ and its left action on $\mathcal{F}$. To describe this new right action, we will keep abusing the notation in the sense that we will use the same symbols for all the actions in sight, since the context should make clear which action is being used in each particular expression. Let (p, f) be an element of the product $\mathcal{P} \times \mathcal{F}$, and let $a \in \mathcal{G}$. We define the effective right action:

$$R_a(p, f) = (R_a p, L_{a^{-1}} f). \tag{1.23}$$

The next step towards the construction of the associated bundle with typical fibre $\mathcal{F}$ consists of taking the quotient space $\mathcal{C}$ generated by this action. In other words, we want to deal with a set whose elements

are equivalence classes in $\mathcal{C} \times \mathcal{F}$ by the equivalence relation: "$(p_1, f_1) \sim (p_2, f_2)$ if, and only if, there exists $a \in \mathcal{G}$ such that $(p_2, f_2) = R_a(p_1, f_1)$". The motivation for this line of attack should be clear from the introductory remarks to this section. Recalling that the right action of $\mathcal{G}$ on $\mathcal{P}$ is fibre preserving, it becomes obvious that all the pairs (p, f) in a given equivalence class have first components p with the same projection $\pi_P(p)$ on $\mathcal{B}$. This means that we have a perfectly well-defined projection π in the quotient space $\mathcal{C}$, namely: $\pi : \mathcal{C} \to \mathcal{B}$ is a map that assigns to each equivalence class the common value of the projection of the first component of all its constituent pairs.

Having a projection, we can now define the fibre of $\mathcal{C}$ over $b \in \mathcal{B}$ naturally as $\pi^{-1}(b)$. We need to show now that each such fibre is diffeomorphic to the putative typical fibre $\mathcal{F}$. More precisely, we want to show that for each local trivialization $(\mathcal{U}, \psi)$ of the original principal bundle $\mathcal{P}$, we can also construct a local trivialization of $\pi^{-1}(\mathcal{U})$, namely, a diffeomorphism $\rho : \pi^{-1}(\mathcal{U}) \to \mathcal{U} \times \mathcal{F}$. To understand how this works, let us fix a point $b \in \mathcal{U}$ and recall that, given the local trivialization $(\mathcal{U}, \psi)$, the map $\tilde{\psi}_b$ provides us with a diffeomorphism of the fibre $\pi_P^{-1}(b)$ with $\mathcal{G}$. We now form the product map of $\tilde{\psi}_b$ with the identity map of $\mathcal{F}$, namely: $(\tilde{\psi}_b, id_{\mathcal{F}}) : \pi_P^{-1}(b) \times \mathcal{F} \to \mathcal{G} \times \mathcal{F}$. Each equivalence class by the right action (1.23) is mapped by the product map $(\tilde{\psi}_b, id_{\mathcal{F}})$ into an orbit, as shown in Figure 3.1.

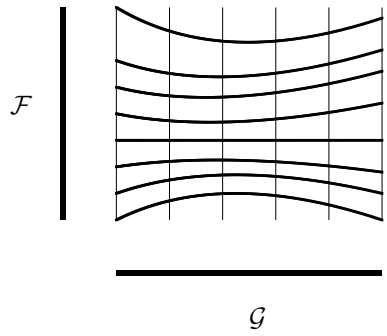

Figure 3.1: Images of equivalence classes

These orbits do not intersect with each other. Moreover, they can be seen as graphs of single-valued $\mathcal{F}$-valued functions of $\mathcal{G}$. Therefore, choosing any particular value $g \in \mathcal{G}$, we see that these orbits can be parametrized by $\mathcal{F}$. This provides the desired one-to-one and onto relation between the fibre $\pi^{-1}(b)$ and the manifold $\mathcal{F}$, which can now legitimately be called the typical fibre of $\mathcal{C}$. To complete the construction of the desired fibre bundle, we need to guarantee that the fibre-wise isomorphism that we have just constructed depends differentiably on b, a requirement that we assume fulfilled.

3.2 Vector fields and the Lie bracket

3.2.1 Vector fields

A *vector field* $\mathbf{V}$ on a manifold $\mathcal{M}$ is an assignment to each point $p \in \mathcal{M}$ of a tangent vector $\mathbf{V}_p = \mathbf{V}(p) \in T_p\mathcal{M}$. We restrict our attention to *smooth vector fields*, whose components are smooth functions in any given chart. A vector field is, therefore, a smooth map:

$$\mathbf{V} : \mathcal{M} \to T\mathcal{M}, \tag{2.1}$$

satisfying the condition:

$$\tau \circ \mathbf{V} = id_{\mathcal{M}}, \tag{2.2}$$

where $id_{\mathcal{M}}$ is the identity map of $\mathcal{M}$. The meaning of this last condition is that the vector assigned to the point p is a tangent vector at p, rather than at any other point.

A geometrically convenient way to look at a vector field is to regard it as a *cross section* of the tangent bundle. This terminology arises from the pictorial representation depicted in Figure 3.2, where the base manifold is represented by a shallow arc and the fibres (namely, the tangent spaces) by straight lines hovering above it Then, a cross section looks like a curve cutting through the fibres.

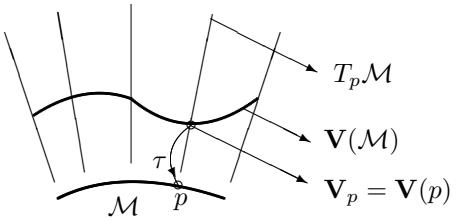

Figure 3.2: A vector field as a cross section of the tangent bundle

3.2.2 The Lie bracket

If $\mathbf{V}$ is a (smooth) vector field on a manifold $\mathcal{M}$ and $f : \mathcal{M} \to \mathbb{R}$ is a smooth function, then the map:

$$\mathbf{V}f : \mathcal{M} \to \mathbb{R}, \tag{2.3}$$

defined as:

$$p \mapsto \mathbf{V}_p(f) \tag{2.4}$$

is again a smooth map. It assigns to each point $p \in \mathcal{M}$ the directional derivative of the function f in the direction of the vector field at p. In other words, a vector field assigns to each smooth function another smooth function. Given, then, two vector fields $\mathbf{V}$ and $\mathbf{W}$ over $\mathcal{M}$, the iterated evaluation:

$$h = \mathbf{W}(\mathbf{V}f) : \mathcal{M} \to \mathbb{R}, \tag{2.5}$$

gives rise to a legitimate smooth function h on $\mathcal{M}$.

On the basis of the above considerations, one may be tempted to define a composition of vector fields by declaring that the composition $\mathbf{W} \circ \mathbf{U}$ is the vector field which assigns to each function f the function h defined by Equation (2.5). This wishful thinking, however, does not work. To see why, it is convenient to work in components in some chart with coordinates x^i. Let:

$$\mathbf{V} = V^i \frac{\partial}{\partial x^i} \qquad \mathbf{W} = W^i \frac{\partial}{\partial x^i}, \tag{2.6}$$

where the components V^i and W^i $(i = 1, ..., m)$ are smooth real-valued functions defined over the m-dimensional domain of the chart. Given a smooth function $f : \mathcal{M} \to \mathbb{R}$, the function $g = \mathbf{V}f$ is evaluated at a point $p \in \mathcal{M}$ with coordinates x^i $(i = 1, ..., m)$ as:

$$g(p) = V^i \, \frac{\partial f}{\partial x^i}. \tag{2.7}$$

Notice the slight abuse of notation we incur into by identifying the function f with its representation in the coordinate system.

We now apply the same prescription to calculate the function $h = \mathbf{W}g$ and obtain:

$$h(p) = W^i \frac{\partial g}{\partial x^i} = W^i \frac{\partial \left(V^j \frac{\partial f}{\partial x^j} \right)}{\partial x^i} = \left(W^i \frac{\partial V^j}{\partial x^i} \right) \frac{\partial f}{\partial x^j} + W^i V^j \frac{\partial^2 f}{\partial x^i \partial x^j}. \tag{2.8}$$

The last term of this expression, by involving second derivatives, will certainly not transform as the components of a vector should under a change of coordinates. Neither will the first. This negative result, on the other hand, suggests that the offending terms could perhaps be eliminated by subtracting from the composition $\mathbf{WV}$ the opposite composition $\mathbf{VW}$, namely:

$$(\mathbf{WV} - \mathbf{VW})(f) = \left(W^i \frac{\partial V^j}{\partial x^i} - V^i \frac{\partial W^j}{\partial x^i} \right) \frac{\partial f}{\partial x^j}. \tag{2.9}$$

The vector field thus obtained, is called the *Lie bracket* of $\mathbf{W}$ and $\mathbf{V}$ (in that order) and is denoted by $[\mathbf{W}, \mathbf{V}]$. More explicitly, its components in the coordinate system x^i are given by:

$$[\mathbf{W}, \mathbf{V}]^j = W^i \frac{\partial V^j}{\partial x^i} - V^i \frac{\partial W^j}{\partial x^i}. \tag{2.10}$$

Upon a coordinate transformation, these components transform according to the rules of transformation of a vector.

The following properties of the Lie bracket are worthy of notice:

(1) Skew symmetry:

$$[\mathbf{W}, \mathbf{V}] = -[\mathbf{V}, \mathbf{W}] \tag{2.11}$$

(2) Jacobi identity:

$$[[\mathbf{W}, \mathbf{V}], \mathbf{U}] + [[\mathbf{V}, \mathbf{U}], \mathbf{W}] + [[\mathbf{U}, \mathbf{W}], \mathbf{V}] = 0 \tag{2.12}$$

The collection of all vector fields over a manifold has the natural structure of an infinite dimensional vector space, where addition and multiplication by a scalar are defined in the obvious way. In this vector space, the Lie bracket operation is bilinear. A vector space endowed with a bilinear operation satisfying conditions (1) and (2) is called a *Lie algebra*.

Vector fields can be multiplied by functions to produce new vector fields. Indeed, for a given function f and a given vector field $\mathbf{V}$, we can define the vector field $f\mathbf{V}$ by:

$$(f\mathbf{V})_p = f(p)\mathbf{V}_p. \tag{2.13}$$

It can be shown that:

$$[g\mathbf{W}, f\mathbf{V}] = gf[\mathbf{W}, \mathbf{V}] + g\,(\mathbf{W}f)\ \mathbf{V} - f\,(\mathbf{V}g)\ \mathbf{W}, \tag{2.14}$$

where g, f are smooth functions and $\mathbf{W}, \mathbf{V}$ are vector fields over a manifold $\mathcal{M}$.

3.2.3 Push-forwards

We have seen that the differential of a map between manifolds carries tangent vectors to tangent vectors. This operation is sometimes called a *push-forward*. Does a map also push forward vector fields to vector fields? Let $\mathbf{V} : \mathcal{M} \to T\mathcal{M}$ be a vector field on $\mathcal{M}$ and let $g : \mathcal{M} \to \mathcal{N}$ be a smooth map. Since the differential of g is a map of the form $g_* : T\mathcal{M} \to T\mathcal{N}$, the composition $g_* \circ \mathbf{V}$ makes perfect sense, but it delivers a (well-defined) map $g_*\mathbf{V}$ from $\mathcal{M}$ (and *not* from $\mathcal{N}$) into $T\mathcal{N}$. This is not a vector field, nor can it in general be turned into one. If the dimension of $\mathcal{M}$ is larger than that of $\mathcal{N}$, points in $\mathcal{N}$ will end up being assigned more than one vector. If the dimension of the source manifold is less than that of the target, on the other hand, even if the function is one-to-one, there will necessarily exist points in $\mathcal{N}$ to which no vector is assigned. The only case in which the push-forward of a vector field can be regarded as a vector field on the target manifold is the case in which both manifolds are of the same dimension and the map is a diffeomorphism.

Notwithstanding the above remark, let $g : \mathcal{M} \to \mathcal{N}$ be a smooth map. We say that the vector fields $\mathbf{V} : \mathcal{M} \to T\mathcal{M}$ and $\mathbf{W} : \mathcal{N} \to T\mathcal{N}$ are *g-related* if:

$$g_*\mathbf{V}(p) = \mathbf{W}(g(p)) \qquad \forall p \in \mathcal{M}. \tag{2.15}$$

According to this definition, if g happens to be a diffeomorphism, then $\mathbf{V}$ and $g_*\mathbf{V}$ are automatically g-related. The pushed-forward vector field is then given by $g_* \circ \mathbf{V} \circ g^{-1}$.

Theorem 3.2.1 *Let $\mathbf{V}_1$ be g-related to $\mathbf{W}_1$ and let $\mathbf{V}_2$ be g-related to $\mathbf{W}_2$. Then the Lie bracket $[\mathbf{V}_1, \mathbf{V}_2]$ is g-related to the Lie bracket $[\mathbf{W}_1, \mathbf{W}_2]$, that is:*

$$[g_*\mathbf{V}_1, g_*\mathbf{V}_2] = g_*[\mathbf{V}_1, \mathbf{V}_2]. \tag{2.16}$$

3.2.4 The algebra of tensors on a vector space

Although a more general situation may be envisioned, we now consider the collection of all possible tensor products involving any finite number of factors, each factor being equal to a given vector space V or its dual V^*. The order of the factors, of course, matters, but it is customary to say that a tensor product is *of type (r,s)* if it is obtained by multiplying r copies of V and s copies of V^*, regardless of the order in which these copies appear in the product. An element of such tensor product is also called a tensor of type (r, s). Another common terminology is to refer to r and s, respectively, as the contravariant and covariant degrees of the tensor. Thus, a vector is a tensor of type $(1, 0)$, while a covector is of type $(0, 1)$. By convention, a tensor of type $(0, 0)$ is identified with a scalar. Since the field of scalars $\mathbb{R}$ has the natural structure of a vector space (whose elements are tensors of type $(0, 0)$), it makes sense to take its tensor product with a vector space. Note that $\mathbb{R} \otimes V = V$.

The tensor product of a tensor of type (r_1, s_1) with a tensor of type (r_2, s_2) is a tensor of type $(r_1 + r_2, s_1 + s_2)$. A map from a Cartesian product of vector spaces into a vector space is said to be *multilinear* if it is linear in each of the arguments. A tensor T of type (r, s) can be considered as a multilinear map such that $T(\omega_1, ..., \omega_r, \mathbf{v}_1, ..., \mathbf{v}_s) \in \mathbb{R}$, where $\mathbf{v}_i$ and ω_j belong, respectively, to V and V^*, for each $i = 1, ..., r$ and each $j = 1, ..., s$.

The collection of all tensors of all orders defined on a vector space V can be given the formal structure of an *algebra* (with the operations of *direct sum* and tensor product) known as the *algebra of tensors*

on V. Considering only tensors of covariant degree zero, namely, tensors of type $(r, 0)$, we obtain the *contravariant tensor algebra* of V. When written in components, all indices of tensors in this algebra are superscripts.

In a similar way, one can define the *covariant tensor algebra* by considering tensors of type $(0, s)$. On the other hand, considering V^* as a vector space in its own right, we could form its contravariant tensor algebra, and these two objects turn out to be the same. The contravariant and covariant algebras can be considered dual to each other in the sense that there exists a canonical way to evaluate an element of one over an element of the other to produce a real number linearly. Considering a tensor T of type $(k, 0)$ and a tensor S of type $(0, k)$ and using a basis in V, this evaluation reads:

$$\langle S, T\rangle = S_{i_1 \ldots i_k}\, T^{i_1 \ldots i_k}. \tag{2.17}$$

If the tensors are of different orders (that is, a tensor of type $(r, 0)$ and a tensor of type $(0, s)$ with $s \neq r$), we define the evaluation as zero.

A tensor T of type $(r, 0)$ can be seen as a multilinear map:

$$\begin{array}{rcl} T : V^*, ..., V^* & \longrightarrow & \mathbb{R} \\ (\omega_1, ..., \omega_r) & \mapsto & T(\omega_1, ..., \omega_r), \quad \omega_1, ..., \omega_r \in V^*. \end{array} \tag{2.18}$$

For tensors in the contravariant or covariant algebras it makes sense to speak about *symmetry* and *skew symmetry*.

A tensor of type $(r, 0)$ is said to be *(completely) symmetric* if the result of the operation (2.18) is independent of the order of the arguments. Put in other words, exchanging any two arguments with each other produces no effect in the result of the multilinear operator T. A similar criterion applies for completely symmetric tensors of order $(0, s)$, except that the arguments are vectors rather than covectors. Choosing a basis in V, symmetry boils down to indifference to index swapping.

Analogously, a tensor of type $(r, 0)$ is *(completely) skew-symmetric* if every mutual exchange of two arguments alters the sign of the result, leaving the absolute value unchanged. By convention, all tensors of type $(0, 0)$ (scalars), $(1, 0)$ (vectors) and $(0, 1)$ (covectors) are considered to be both symmetric and skew-symmetric. Notice that a completely skew-symmetric tensor of type $(r, 0)$ with r larger than the dimension of the vector space of departure must necessarily vanish.

The collections of all symmetric or skew symmetric tensors (whether contravariant or covariant) do not constitute a subalgebra of the tensor algebra, for the simple reason that the tensor multiplication of two symmetric (or skew-symmetric) tensors is not symmetric (skew-symmetric) in general. Nevertheless, it is possible, and convenient, to define algebras of symmetric and skew-symmetric tensors by modifying the multiplicative operation so that the results stay within the algebra. The case of skew-symmetric tensors is the most fruitful. It gives rise to the so-called *exterior algebra* of a vector space, which we will now explore. It will permit us to answer many intriguing questions such as: is there anything analogous to the cross-product of vectors in dimensions other than 3? What is an area and what is the meaning of flux?

3.2.5 Exterior algebra

The space of skew-symmetric contravariant tensors of type $(r, 0)$ will be denoted by $\Lambda^r(V)$. The elements of $\Lambda^r(V)$ will be also called *r-vectors* and, more generally, *multivectors*. The number r is the *order* of the

multivector. As before, the space $\Lambda^0(V)$ coincides with the scalar field $\mathbb{R}$, while $\Lambda^1(V)$ coincides with the vector space V.

Consider the ordered r-tuple of covectors $(\omega_1, \omega_2, ..., \omega_r)$ and let π denote a *permutation* of this set. Such a permutation is even (odd) if it is obtained by an even (odd) number of exchanges between pairs of elements in the original set. An even (odd) permutation π has a *signature*, denoted by $sign(\pi)$, equal to 1 (-1).

Given an arbitrary tensor T of type $(r, 0)$, we define its *skew-symmetric part* $\mathcal{A}^r(T)$ as the multilinear map defined by the formula:

$$\mathcal{A}^r(T)(\omega_1, \omega_2, ..., \omega_r) = \frac{1}{r!} \sum_{\pi} sign(\pi)\, T(\pi). \tag{2.19}$$

As an example, for the case of a contravariant tensor of degree 3, namely, $T = T^{ijk}\mathbf{e}_i \otimes \mathbf{e}_j \otimes \mathbf{e}_k$, where $\mathbf{e}_h$ $(h = 1, ..., n \geq r)$ is a basis of V, the skew symmetric part is obtained as:

$$\mathcal{A}^3(T) = \frac{1}{6}\left(T^{ijk} + T^{jki} + T^{kij} - T^{ikj} - T^{jik} - T^{kji}\right) \mathbf{e}_i \otimes \mathbf{e}_j \otimes \mathbf{e}_k. \tag{2.20}$$

Given two multivectors a and b, of orders r and s, respectively, we define their *exterior product* or *wedge product* as the multivector $a \wedge b$ of order $(r + s)$ obtained as:

$$a \wedge b = \mathcal{A}^{r+s}(a \otimes b). \tag{2.21}$$

What this definition in effect is saying is that in order to multiply two skew-symmetric tensors and obtain a skew-symmetric result, all we have to do is take their tensor product and then project back into the algebra (that is, skew-symmetrize the result)[1]. Since $\mathcal{A}$ is, by definition, a linear operator, the wedge product is linear in each of the factors.

We have seen that the tensor product is not commutative. But, in the case of the exterior product, exchanging the order of the factors can at most affect the sign. The general result is:

$$b \wedge a = (-1)^{rs} a \wedge b, \quad a \in \Lambda^r(V), b \in \Lambda^s(V). \tag{2.22}$$

Thus, for example, the wedge product with itself of a multivector of odd order must necessarily vanish. With some work, it is possible to show that the wedge product is associative, namely: $(a \wedge b) \wedge c = a \wedge (b \wedge c)$.

To calculate the dimension of $\Lambda^r(V)$, we note that, being a tensor, every element in $\Lambda^k(V)$ is expressible as a linear combination of the n^r tensor products $\mathbf{e}_{i_1} \otimes ... \otimes \mathbf{e}_{i_r}$, where $\mathbf{e}_i$, $i = 1, ..., n$, is a basis of V. Because of the skew-symmetry, however, we need to consider only products of the form $\mathbf{e}_{i_1} \wedge ... \wedge \mathbf{e}_{i_r}$. Two such products involving the same factors in any order are either equal or differ in sign, and a product with a repeated factor vanishes. This means that we need only count all possible combinations of n symbols taken r at a time without repetition. The number of such combinations is $\frac{n!}{(n-r)!r!}$. One way to keep track of all these combinations is to place the indices $i_1, ..., i_k$ in strictly increasing order. These combinations are linearly independent, thus constituting a basis. Therefore, the dimension of $\Lambda^r(V)$ is $\frac{n!}{(n-r)!r!}$.

[1] In spite of the natural character of this definition of the wedge product, many authors adopt a definition that includes a combinatorial factor. Thus, the two definitions lead to proportional results. Each definition has some advantages, but both are essentially equivalent. Our presentation of exterior algebra follows closely that of Sternberg S (1983), *Lectures on Differential Geometry*, 2nd ed., Chelsea.

We note that the spaces of r-vectors and $(n-r)$-vectors have the same dimension. There is a kind of fusiform dimensional symmetry around the middle, the dimension starting at 1 for $r = 0$, increasing to a maximum towards $r = n/2$ (say, if n is even) and then going back down to 1 for $r = n$. This observation plays an important role in the identification (and sometimes confusion) of physical quantities. For example, an n-vector functions very much like a scalar, but with a subtle difference.

Let a skew-symmetric contravariant tensor $a \in \Lambda^r(V)$ be given by means of its components on the basis of $\mathcal{C}^r(V)$ inherited from a basis $\mathbf{e}_1, ..., \mathbf{e}_n$ of V as:

$$a = a^{i_1 ... i_r}\ \mathbf{e}_{i_1} \otimes ... \otimes \mathbf{e}_{i_r}. \tag{2.23}$$

Recalling that the skew-symmetry operator $\mathcal{A}^r$ is linear, we obtain:

$$a = \mathcal{A}^r(a) = a^{i_1 ... i_r}\ \mathcal{A}^r\left(\mathbf{e}_{i_1} \otimes ... \otimes \mathbf{e}_{i_r}\right) = a^{i_1 ... i_r}\ \mathbf{e}_{i_1} \wedge ... \wedge \mathbf{e}_{i_r}. \tag{2.24}$$

In these expressions, the summation convention is implied. We have obtained the result that, given a skew-symmetric tensor in components, we can substitute the wedge products for the tensor products of the base vectors. On the other hand, if we would like to express the r-vector a in terms of its components on the basis of $\Lambda^r(V)$ given by the wedge products of the base vectors of V taken in strictly increasing order of the indices, a coefficient of $r!$ will have to be included, namely:

$$a^{i_1, ... i_r}\ \mathbf{e}_{i_1} \wedge ... \wedge \mathbf{e}_{i_r} = r! \sum_{i_1 < ... < i_r} a^{i_1, ... i_r}\ \mathbf{e}_{i_1} \wedge ... \wedge \mathbf{e}_{i_r}. \tag{2.25}$$

This means that the components on the basis (with strictly increasing indices) of the skew symmetric part of a contravariant tensor of type $(k, 0)$ are obtained without dividing by the factorial $k!$ in the projection algorithm. This, of course, is a small advantage to be gained at the expense of the summation convention.

Consider the n-fold wedge product $a = \mathbf{v}_1 \wedge \mathbf{v}_2 \wedge ... \wedge \mathbf{v}_n$, where the $\mathbf{v}$'s are elements of an n-dimensional vector space V. Let $\{\mathbf{e}_1, \mathbf{e}_2, ..., \mathbf{e}_n\}$ be a basis of V. Since each of the $\mathbf{v}$'s is expressible uniquely in this basis we may write:

$$a = (v_1^{i_1}\mathbf{e}_{i_1}) \wedge (v_2^{i_2}\mathbf{e}_{i_2}) \wedge ... \wedge (v_n^{i_n}\mathbf{e}_{i_n}) = v_1^{i_1} v_2^{i_2} ... v_n^{i_n}\ \mathbf{e}_{i_1} \wedge \mathbf{e}_{i_2} \wedge ... \wedge \mathbf{e}_{i_n}, \tag{2.26}$$

where the summation convention is in full swing. Out of the possible n^n terms in this sum, there are exactly $n!$ that can survive, since each of the indices can attain n values, but repeated indices in a term kill it. However, since each of the surviving terms consists of a scalar coefficient times the exterior product of all the n elements of the basis, we can collect them all into a single scalar coefficient A multiplied by the exterior product of the base vectors arranged in a strictly increasing ordering of the indices, namely, we must have that $a = A\mathbf{e}_1 \wedge \mathbf{e}_2 \wedge ... \wedge \mathbf{e}_n$. This scalar coefficient consists of the sum of all the products $v_1^{i_1} v_2^{i_2} ... v_n^{i_n}$ with no repeated indices and with a minus sign if the superscripts form an odd permutation of $1, 2, ..., n$. This is precisely the definition of the determinant of the matrix whose entries are v_i^j. We conclude that, using in $\Lambda^n(V)$ the basis induced by a basis in V, the component of the exterior product of n vectors in an n-dimensional space is equal to the determinant of the matrix of the components of the individual vectors. Apart from providing a neat justification for the notion of determinant, this formula correctly suggests that the geometrical meaning of an n-vector is some measure of the ability of the (n-dimensional) parallelepiped subtended by the vectors to contain a volume. Since we have not yet introduced any metric notion, we cannot associate a number to this volume. Notice on the other hand that, although we cannot say how large a volume is, we can certainly tell that a given n-parallelepiped is, say, twice as large as another. Notice, finally, that changing the order of two factors, or reversing the

sense of one factor, changes the sign of the multivector. So, n-vectors represent *oriented* n-dimensional parallelepipeds.

The collection of all multivectors of all orders (up to the dimension of V), with the exterior product replacing the tensor product, constitutes the *exterior algebra* of V. In a similar way, starting from the dual space V^*, we can construct the algebra of multicovectors.

3.3 Forms and general tensor fields

3.3.1 1-forms

Let $f : \mathcal{U} \to \mathbb{R}$ be a smooth function defined in a neighbourhood $\mathcal{U} \subset \mathcal{M}$ of the point p, and let $df_p : T\mathcal{U} \to \mathbb{R}$ denote its differential at p. We can regard this differential as an element of $T_p^*\mathcal{M}$ by defining its value $\langle df_p, \mathbf{v}_p \rangle$ on any vector $\mathbf{v}_p \in T_p\mathcal{M}$ as:

$$\langle df_p, \mathbf{v}_p \rangle = df_p(\mathbf{v}_p) = \mathbf{v}_p(f). \tag{3.1}$$

In other words, the action of the evaluation of the differential of the function on a tangent vector is equal to the directional derivative of the function in the direction of the vector.

A smooth assignment of a covector ω_p to each point $p \in \mathcal{M}$ is called a *differential* 1*-form on the manifold*. It can be regarded as a cross section of the cotangent bundle, namely, a map:

$$\mathbf{\Omega} : \mathcal{M} \to T^*\mathcal{M}, \tag{3.2}$$

such that $\pi \circ \mathbf{\Omega} = id_{\mathcal{M}}$.

As we have seen, the differential of a function at a point defines a covector. It follows that a smooth scalar function $f : \mathcal{M} \to \mathbb{R}$ determines, by point-wise differentiation, a differential 1-form $\mathbf{\Omega} = df$. It is important to remark that *not all differential* 1*-forms can be obtained as differentials of functions*. The ones that can are called *exact*.

A differential 1-form $\mathbf{\Omega}$ (that is, a cross section of $T^*\mathcal{M}$) can be regarded as acting on vector fields $\mathbf{V}$ (cross sections of $T\mathcal{M}$) to deliver functions $\langle \mathbf{\Omega}, \mathbf{V} \rangle : \mathcal{M} \to \mathbb{R}$, by point-wise evaluation of a covector on a vector.

3.3.2 Pull-backs

Let $f : \mathcal{N} \to \mathbb{R}$ be a smooth function. We define its pull-back by g as the map $g^*f : \mathcal{M} \to \mathbb{R}$ given by the composition:

$$g^*f = f \circ g. \tag{3.3}$$

For a differential 1-form $\mathbf{\Omega}$ on $\mathcal{N}$, we define the pull-back $g^*\mathbf{\Omega} : \mathcal{M} \to T^*\mathcal{M}$ by showing how it acts, point by point, on tangent vectors:

$$\langle [g^*\mathbf{\Omega}](p), \mathbf{v}_p \rangle = \langle \mathbf{\Omega}(g(p)), (g_*)_p \mathbf{v}_p \rangle, \tag{3.4}$$

which can be more neatly written in terms of vector fields as:

$$\langle g^*\mathbf{\Omega}, \mathbf{V} \rangle = \langle \mathbf{\Omega} \circ g,\ g_*\mathbf{V} \rangle. \tag{3.5}$$

Expressed in words, this means that the pull-back by g of a 1-form in $\mathcal{N}$ is the 1-form in $\mathcal{M}$ that assigns to each vector the value that the original 1-form assigns to the image of that vector by g_*.

It is important to notice that the pull-backs of functions and differential 1-forms are always well-defined, regardless of the dimensions of the spaces involved. This should be contrasted with the push-forwards of vector fields, which fail in general to be vector fields on the target manifold.

3.3.3 Tensor bundles

Given a point p of a manifold $\mathcal{M}$, we may identify the vector space V with the tangent space $T_p\mathcal{M}$ and construct the corresponding spaces of tensors of any fixed type. Following the same procedure as for the tangent and cotangent bundles, which will thus become particular cases, one can define *tensor bundles* of any type by adjoining to each point of a manifold the tensor space of the corresponding type. A convenient notational scheme is: $\mathcal{C}^k(\mathcal{M}), \mathcal{C}_k(\mathcal{M})$, respectively, for the bundles of contravariant and covariant tensors of order k. Similarly, the bundles of k-vectors and of k-forms can be denoted, respectively, by: $\Lambda^k(\mathcal{M}), \Lambda_k(\mathcal{M})$. Each of these bundles can be shown (by a procedure identical to that used in the case of the tangent and cotangent bundles) to have a natural structure of a differentiable manifold of the appropriate dimension. A (smooth) section of a tensor bundle is called a *tensor field over* $\mathcal{M}$, of the corresponding type. A (smooth) section of the bundle $\Lambda_k(\mathcal{M})$ of k-forms is also called a *differential k-form*. A scalar function on a manifold is also called a *differential* 0*-form*.

In a chart of the m-dimensional manifold $\mathcal{M}$ with coordinates x^i, a contravariant tensor field $\mathbf{T}$ of order r is given as:

$$\mathbf{T} = T^{i_1,...,i_r} \frac{\partial}{\partial x^{i_1}} \otimes ... \otimes \frac{\partial}{\partial x^{i_r}}, \tag{3.6}$$

where $T^{i_1,...,i_r} = T^{i_1,...,i_r}(x^1, ..., x^m)$ are r^m smooth functions of the coordinates. Similarly, a covariant tensor field $\mathbf{U}$ of order r is given by:

$$\mathbf{U} = U_{i_1,...,i_r}\, dx^{i_1} \otimes ... \otimes dx^{i_r}, \tag{3.7}$$

and a differential r-form ω by:

$$\omega = \omega_{i_1,...,i_r}\, dx^{i_1} \wedge ... \wedge dx^{i_r}. \tag{3.8}$$

Notice that, in principle, the indexed quantity $\omega_{i_1,...,i_r}$ need not be specified as skew-symmetric with respect to the exchange of any pair of indices, since the exterior product of the base-forms will do the appropriate skew-symmetrization job. As an alternative, we may suspend the standard summation convention in (3.8) and consider only indices in ascending order. As a result, if $\omega_{i_1,...,i_r}$ *is* skew symmetric *ab initio*, the corresponding components are to be multiplied by $r!$.

Of particular interest for the theory of integration on manifolds are differential m-forms, where m is the dimension of the manifold. From our treatment of the algebra of r-forms, we know that the dimension of the space of m-covectors is exactly 1. In a coordinate chart, a basis for differential m-forms is, therefore, given by: $dx^1 \wedge ... \wedge dx^m$. In other words, the representation of a differential m-form ω in a chart is :

$$\omega = f(x^1, ..., x^m)\, dx^1 \wedge ... \wedge dx^m, \tag{3.9}$$

where $f(x^1, ..., x^m)$ is a smooth scalar function of the coordinates in the patch. Consider now another coordinate patch with coordinates $y^1, ..., y^m$, whose domain has a non-empty intersection with the domain of the previous chart. In this chart we have:

$$\omega = \hat{f}(y^1, ..., y^m)\, dy^1 \wedge ... \wedge dy^m. \tag{3.10}$$

We want to find the relation between the functions f and $\hat{f}$. Since the transition functions $y^i = y^i(x^1, ..., x^m)$ are smooth, we can write:

$$\omega = \hat{f}(y^1, ..., y^m)\, dy^1 \wedge ... \wedge dy^m = \hat{f}\, \frac{\partial y^1}{\partial x_{j_1}} ... \frac{\partial y^m}{\partial x_{j_m}}\, dx^{j_1} \wedge ... \wedge dx^{j_m} \tag{3.11}$$

or, by definition of determinant:

$$\omega = \det\{\frac{\partial y^1, ..., y^m}{\partial x^1, ..., x^m}\} \hat{f}\, dx^1 \wedge ... \wedge dx^m = J_{y,x}\, \hat{f}\, dx^1 \wedge ... \wedge dx^m, \tag{3.12}$$

where the Jacobian determinant $J_{y,x}$ does not vanish at any point of the intersection of the two coordinate patches. Comparing with Equation (3.9), we conclude that:

$$f = J_{y,x}\, \hat{f}. \tag{3.13}$$

A nowhere vanishing differentiable m-form on a manifold $\mathcal{M}$ of dimension m is called a *volume form* on $\mathcal{M}$. It can be shown that a manifold is orientable if, and only if, it admits a volume form.

The notion of pull-back can be naturally generalized for covariant tensors of any order. For a contravariant tensor field $\mathbf{U}$ of order r on $\mathcal{N}$ (and, in particular, for differential r-forms on $\mathcal{N}$), the pull-back by a smooth function $g : \mathcal{M} \to \mathcal{N}$ is a corresponding field on $\mathcal{M}$ obtained by an extension of the case $r = 1$, as follows:

$$g^*\mathbf{U}\ (\mathbf{V}_1, ..., \mathbf{V}_r) = (\mathbf{U} \circ g)\ (g_*\mathbf{V}_1, ..., g_*\mathbf{V}_r), \tag{3.14}$$

where $\mathbf{U}$ is regarded as a multilinear function of r vector fields $\mathbf{V}_i$.

3.3.4 The exterior derivative

The *exterior derivative* of differential forms is an operation that generalizes the gradient, curl and divergence operators of classical vector calculus. The exterior derivative of a differential r-form on a manifold $\mathcal{M}$ is a differential $(r+1)$-form defined over the same manifold. Instead of introducing, as one certainly could, the definition of exterior differentiation in an intrinsic axiomatic manner, we will proceed to define it in a coordinate system and show that the definition is, in fact, coordinate independent. Let, therefore, x^i $(i = 1, ..., m)$ be a coordinate chart and let ω be an r-form given as:

$$\omega = \omega_{i_1,...,i_r}\, dx^{i_1} \wedge ... \wedge dx^{i_r}, \tag{3.15}$$

where $\omega_{i_1,...,i_r} = \omega_{i_1,...,i_r}(x^1, ..., x^m)$ are smooth functions of the coordinates. We define the exterior derivative of ω, denoted by $d\omega$, as the differential $(r+1)$-form obtained as:

$$d\omega = d\omega_{i_1,...,i_r}\ \wedge dx^{i_1} \wedge ... \wedge dx^{i_r}, \tag{3.16}$$

where the d on the right-hand side denotes the ordinary differential of functions. More explicitly:

$$d\omega = \frac{\partial \omega_{i_1,...,i_r}}{\partial x^k}\, dx^k \wedge dx^{i_1} \wedge ... \wedge dx^{i_r}. \tag{3.17}$$

Note that for each specific combination of (distinct) indices $i_1, ..., i_r$, the index k ranges only on the remaining possibilities, since the exterior product is skew symmetric. Thus, in particular, if ω is a differential m-form defined over an m-dimensional manifold, its exterior derivative vanishes identically (as it should, being an $(m+1)$-form).

Let y^i $(i = 1, ..., m)$ be another coordinate chart with a non-empty intersection with the previous chart. We have:

$$\omega = \hat{\omega}_{i_1,...,i_r}\, dy^{i_1} \wedge ... \wedge dy^{i_r}, \tag{3.18}$$

for some smooth functions $\hat{\omega}_{i_1,...,i_r}$ of the y^i-coordinates. The two sets of components are related by:

$$\omega_{i_1,...,i_r} = \hat{\omega}_{j_1,...,j_r}\, \frac{\partial y^{j_1}}{\partial x^{i_1}} \, ... \, \frac{\partial y^{j_r}}{\partial x^{i_r}}. \tag{3.19}$$

Notice that we have not troubled to collect terms by, for example, prescribing a strictly increasing order. The summation convention is in effect. We now apply the prescription (3.16) and obtain:

$$d\omega = d\left(\hat{\omega}_{j_1,...,j_r}\, \frac{\partial y^{j_1}}{\partial x^{i_1}} \, ... \, \frac{\partial y^{j_r}}{\partial x^{i_r}}\right) \wedge dx^{i_1} \wedge ... \wedge dx^{i_r}. \tag{3.20}$$

The crucial point now is that the terms containing the second derivatives of the coordinate transformation will evaporate due to their intrinsic symmetry, since they are contracted with an intrinsically skew symmetric wedge product of two 1-forms. We have, therefore:

$$d\omega = \frac{\partial \hat{\omega}_{j_1,...,j_r}}{\partial y^m}\, \frac{\partial y^m}{\partial x^k}\, \frac{\partial y^{j_1}}{\partial x^{i_1}} \, ... \, \frac{\partial y^{j_r}}{\partial x^{i_r}}\, dx^k \wedge dx^{i_1} \wedge ... \wedge dx^{i_r}, \tag{3.21}$$

or, finally:

$$d\omega = \frac{\partial \hat{\omega}_{j_1,...,j_r}}{\partial y^m}\, dy^m \wedge dy^{j_1} \wedge ... \wedge dy^{j_r}, \tag{3.22}$$

which is exactly the same prescription in the coordinate system y^i as Equation (3.16) is in the coordinate system x^i. This completes the proof of independence from the coordinate system.

From this definition, we can deduce a number of important properties of the exterior derivative:

(1) Linearity: d is a linear operator, viz.:

$$d(a\,\alpha \,+\, b\,\beta) \;=\; a\,d\alpha \,+\, b\,d\beta \quad \forall a, b \in \mathbb{R} \quad \alpha, \beta \in \Lambda_r(\mathcal{M}). \tag{3.23}$$

(2) Quasi-Leibniz rule:

$$d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^r \alpha \wedge d\beta \quad \forall \alpha \in \Lambda_r(\mathcal{M}), \;\; \beta \in \Lambda_s(\mathcal{M}). \tag{3.24}$$

(3) Nilpotence:

$$d^2(.) = d(d(.)) = 0. \tag{3.25}$$

Moreover, it can be shown that the exterior derivative commutes with pull-backs. Finally, the exterior derivative of a 1-form has the following interesting interaction with the Lie bracket. If α is a differential 1-form and $\mathbf{u}$ and $\mathbf{v}$ are smooth vector fields on a manifold $\mathcal{M}$, then:

$$\langle d\alpha \mid \mathbf{u} \wedge \mathbf{v} \rangle = \mathbf{u}\,(\langle \alpha \mid \mathbf{v} \rangle) - \mathbf{v}\,(\langle \alpha \mid \mathbf{u} \rangle) - \langle \alpha \mid [\mathbf{u}, \mathbf{v}] \rangle. \tag{3.26}$$

A differential form ω is *closed* if $d\omega = 0$. Thus, all m-forms in an m-dimensional manifold are automatically closed. An r-form (with $r > 1$) is *exact* if there exists an $(r-1)$-form σ such that $\omega = d\sigma$. By Property 3 above, all exact forms are closed. The converse is true *locally*. In other words, for every point in a manifold there exists an open neighbourhood on which the restriction of a closed form is exact. But this property may fail globally. An example is the 1-form given by $\omega = (x\, dy - y\, dx)/(x^2 + y^2)$ defined on an annular region of $\mathbb{R}^2$ with centre at the origin $x = y = 0$. This form is closed but not exact. The existence of forms of this type reflects the presence of topological invariants (such as holes) in the manifold.

3.3.5 The flow of a vector field

Let $\mathbf{V} : \mathcal{M} \to T\mathcal{M}$ be a (smooth) vector field. A (parametrized) curve $\gamma : H \to \mathcal{M}$ is called an *integral curve* of the vector field if its tangent at each point coincides with the vector field at that point. In other words, denoting by s the curve parameter, the following condition holds:

$$\frac{d\gamma(s)}{ds} = \mathbf{V}(\gamma(s)) \qquad \forall s \in H \subset \mathbb{R}. \tag{3.27}$$

As a consequence of the fundamental theorem of existence and uniqueness of local solutions of systems of ordinary differential equations, it is possible to prove the following fundamental theorem for vector fields on manifolds.

Theorem 3.3.1 *If* $\mathbf{V}$ *is a vector field on a manifold* $\mathcal{M}$*, then for every* $p \in \mathcal{M}$ *there exists an integral curve* $\gamma(s,p) : I_p \to \mathcal{M}$ *such that: (i)* I_p *is an open interval of* $\mathbb{R}$ *containing the origin* $s = 0$*; (ii)* $\gamma(0,p) = p$*; and (iii)* I_p *is maximal in the sense that there exists no integral curve starting at* p *and defined on an open interval of which* I_p *is a proper subset. Moreover,*

$$\gamma(s, \gamma(s', x)) = \gamma(s + s', x) \qquad \forall s, s', s + s' \in I_p. \tag{3.28}$$

The map given by:

$$p, s \mapsto \gamma(s, p), \tag{3.29}$$

is called the *flow* of the vector field $\mathbf{V}$ whose integral curves are $\gamma(s,p)$. In this definition, the map is expressed in terms of its action on pairs of points belonging to two different manifolds, $\mathcal{M}$ and $\mathbb{R}$, respectively. Not all pairs, however, are included in the domain, since I_p is not necessarily equal to $\mathbb{R}$. Moreover, since the intervals I_p are point dependent, the domain of the flow is not even a product manifold. One would be tempted to take the intersection of all such intervals so as to work with a product manifold given by $\mathcal{M}$ times the smallest interval I_p. Unfortunately, as we know from elementary calculus, this (infinite) intersection may consist of a single point. All that can be said about the domain of the flow is that it is an open subset of the Cartesian product $\mathcal{M} \times \mathbb{R}$. When the domain is equal to this product manifold, the vector field is said to be *complete* and the corresponding flow is called a *global flow*. It can be shown that if $\mathcal{M}$ is compact, or if the vector field is smooth and vanishes outside a compact subset of $\mathcal{M}$, the flow is necessarily global.

3.3.6 One-parameter groups of transformations generated by flows

Given a point $p_0 \in \mathcal{M}$ it is always possible to find a small enough neighbourhood $U(p_0) \subset \mathcal{M}$ such that the intersection of all the intervals I_p with $p \in U(p_0)$ is an open interval J containing the origin. For each value $s \in J$, the flow $\gamma(s,p)$ can be regarded as a map:

$$\gamma_s : U(p_0) \longrightarrow \mathcal{M}, \tag{3.30}$$

defined as:

$$\gamma_s(p) = \gamma(s,p), \qquad p \in U(p_0). \tag{3.31}$$

This map is clearly one-to-one, since otherwise we would have two integral curves intersecting each other, against the statement of the fundamental theorem. Moreover, again according to the fundamental theorem, this is a smooth map with a smooth inverse over its image. The inverse is, in fact, given by:

$$\gamma_s^{-1} = \gamma_{-s}, \tag{3.32}$$

where γ_{-s} is defined over the image $\gamma_s(U(p_0))$. Notice that γ_0 is the identity map of $U(p_0)$. Finally, for the appropriate range of values of s and r, we have the composition law:

$$\gamma_r \circ \gamma_s = \gamma_{r+s}. \tag{3.33}$$

The set of maps γ_s is said to constitute the *one-parameter local pseudo-group* generated by the vector field (or by its flow). If the neighbourhood $U(p_0)$ can be extended to the whole manifold for some open interval J (no matter how small), each map γ_s is called a *transformation* of $\mathcal{M}$. In that case we speak of a *one-parameter pseudo-group of transformations* of $\mathcal{M}$. Finally, in the best of all possible worlds, if $J = \mathbb{R}$ the one-parameter subgroup of transformations becomes elevated to a *one-parameter group of transformations*.[2] This is an Abelian (i.e., commutative) group, as is clearly shown by the composition law (3.33). We may say that every complete vector field generates a one-parameter group of transformations of the manifold.

The converse construction, namely the generation of a vector field out of a given one-parameter pseudo-group of transformations, is also of interest. It can be shown that every one-parameter pseudo-group of transformations γ_s is generated by the vector field:

$$\mathbf{V}(p) = \frac{d\gamma_s(p)}{ds}|_{s=0}. \tag{3.34}$$

3.4 Lie derivatives, Lie groups

3.4.1 The Lie derivative

We have learned that a vector field determines at least a one-parameter pseudo-group in a neighbourhood of each point of the underlying manifold. For each value of the parameter s within a certain interval containing the origin, this neighbourhood is mapped diffeomorphically onto another neighbourhood. Having at our disposal a diffeomorphism, we can consider the pushed-forward or pulled-back versions of tensors of every type, including multivectors and differential forms. Physically, these actions represent how the various quantities are *convected* (or dragged) by the flow. To elicit a mental picture, we show in Figure 3.3 a vector $\mathbf{w}_p$ in a manifold as a small segment $\overrightarrow{pq}$ (a small piece of a curve, say), and we draw the integral curves of a vector field $\mathbf{V}$ emerging from each of its end points, p and q. These curves are everywhere tangent to the underlying vector field $\mathbf{V}$, which we do not show in the figure. If s denotes the (natural) parameter along these integral curves, an increment of Δs applied from each of these points along the corresponding integral curve, will result in two new points p' and q', respectively. The (small) segment $\overrightarrow{p'q'}$ can be seen as a vector $\mathbf{w}'$, which we regard as the convected counterpart of $\mathbf{w}_p$ as it is dragged by the flow of $\mathbf{V}$ by an amount Δs. If $\mathbf{w}_p$ happens to be part of a vector field $\mathbf{W}$ defined in a neighbourhood of p', so that $\mathbf{w}_p = \mathbf{W}(p)$, we have that at the point p' there is, in addition to the dragged vector $\mathbf{w}'$, a vector $\mathbf{W}(p')$. There is no reason why these two vectors should be equal. The difference $\mathbf{W}(p') - \mathbf{w}'$ (divided by Δs) gives us an idea of the meaning of the Lie derivative of $\mathbf{W}$ with respect to $\mathbf{V}$ at p'.

The idea behind the definition of the Lie derivative of a tensor field with respect to a given vector field at a point p is the following. We consider a small value s of the parameter and convect the tensor field back to $s = 0$ by using the appropriate pull-back or push-forward. This operation will, in particular,

[2] The general notion of group will be reviewed below.

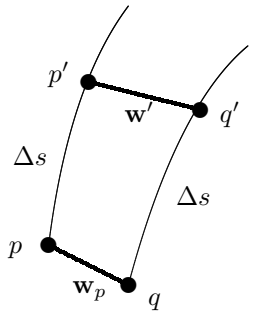

Figure 3.3: Dragging of a vector by a flow

provide a value of the convected tensor field at the point p. We then subtract from this value the original value of the field at that point (a legitimate operation, since both tensors operate on the same tangent and/or cotangent space), divide by s and compute the limit as $s \to 0$. To understand how to calculate a Lie derivative, it is sufficient to make the definition explicit for the case of functions, vector fields and one-forms. The general case is then inferred from these three basic cases, as we shall demonstrate. We will also prove that the term "derivative" is justified. Notice that a Lie derivative is defined with respect to a given vector field. It is not an intrinsic property of the tensor field being differentiated. The Lie derivative of a tensor field at a point is a tensor of the same type.

The Lie derivative of a scalar

Let $g : \mathcal{M} \to \mathcal{N}$ be a mapping between two manifolds and let $f : \mathcal{N} \to \mathbb{R}$ be a function. Recall that, according to Equation (3.3), the pull-back of f by g is the map $g^* f : \mathcal{M} \to \mathbb{R}$ defined as the composition:

$$g^* f = f \circ g. \tag{4.1}$$

Let a (time-independent, for now) vector field $\mathbf{V}$ be defined on $\mathcal{M}$ and let $\gamma_s : U \to \mathcal{M}$ denote the action of its flow on a neighbourhood of a point $p \in \mathcal{M}$. If a function $f : \mathcal{M} \to \mathbb{R}$ is defined, we can calculate:

$$\gamma_s^* f := f \circ \gamma_s. \tag{4.2}$$

The Lie derivative at the point p is given by:

$$L_V f(p) := \lim_{s \to 0} \frac{(\gamma_s^* f)(p) - f(p)}{s} = \lim_{s \to 0} \frac{f(\gamma_s(p)) - f(p)}{s} \tag{4.3}$$

Thus, we obtain

$$L_V f(p) = \mathbf{v}_p(f). \tag{4.4}$$

In simple words, the Lie derivative of a scalar field with respect to a given vector field coincides, at each point, with the directional derivative of the function in the direction of the field at that point.

The Lie derivative of a vector field

Vectors are pulled forward by mappings. Thus, given the map $g : \mathcal{M} \to \mathcal{N}$, to bring a tangent vector from $\mathcal{N}$ back to $\mathcal{M}$, we must use the fact that g is invertible and that the inverse is differentiable, such as when g is a diffeomorphism. Let $\mathbf{W} : \mathcal{N} \to T\mathcal{N}$ be a vector field on $\mathcal{N}$. The corresponding vector field

on $\mathcal{M}$ is then given by $g_*^{-1} \circ \mathbf{W} \circ g : \mathcal{M} \to T\mathcal{M}$. Accordingly, the Lie derivative of the vector field $\mathbf{W}$ with respect to the vector field $\mathbf{V}$, with flow γ_s, at a point $p \in \mathcal{M}$ is defined as:

$$L_V \mathbf{W}(p) = \lim_{s \to 0} \frac{\gamma_{s*}^{-1} \circ \mathbf{W} \circ \gamma_s(p) - \mathbf{W}(p)}{s}. \tag{4.5}$$

It can be shown that the Lie derivative of a vector field coincides with the Lie bracket:

$$L_V \mathbf{W} = [\mathbf{V}, \mathbf{W}]. \tag{4.6}$$

The Lie derivative of a one-form

Since one-forms are pulled back by a map, we define the Lie derivative of the one form $\omega : \mathcal{M} \to T^*\mathcal{M}$ at the point p as:

$$L_V \omega(p) = \lim_{s \to 0} \frac{\gamma_s^* \circ \omega \circ \gamma_s(p) - \omega(p)}{s}. \tag{4.7}$$

The Lie derivative of arbitrary tensor fields

It is clear that, by virtue of their definition by means of limits, the Lie derivatives defined so far are linear operators. To extend the definition of the Lie derivative to tensor fields of arbitrary order, we need to make sure that the Leibniz rule with respect to the tensor product is satisfied. Otherwise, we wouldn't have the right to use the term "derivative" to describe it. It is enough to consider the case of a monomial such as:

$$\mathbf{T} = \omega_1 \otimes ... \otimes \omega_m \otimes \mathbf{W}_1 \otimes ... \otimes \mathbf{W}_n, \tag{4.8}$$

where ω_i are m 1-forms and $\mathbf{W}_j$ are n vector fields. We define:

$$L_V \mathbf{T}(p) = \lim_{s \to 0} \frac{\gamma_s^* \circ \omega_1 \circ \gamma_s(p) \otimes ... \otimes \gamma_{s*}^{-1} \circ \mathbf{W}_1 \circ \gamma_s(p) \otimes ... - \mathbf{T}(p)}{s}. \tag{4.9}$$

Let us verify the satisfaction of the Leibniz rule for the case of the tensor product of a one-form by a vector:

$$L_V(\omega \otimes \mathbf{W})(p) = \lim_{s \to 0} \frac{\gamma_s^* \circ \omega \circ \gamma_s(p) \otimes \gamma_{s*}^{-1} \circ \mathbf{W} \circ \gamma_s(p) \otimes -\omega(p) \otimes \mathbf{W}(p)}{s}. \tag{4.10}$$

Subtracting and adding to the denominator the expression $\omega(p) \otimes \gamma_{s*}^{-1} \circ \mathbf{W} \circ \gamma_s(p)$ the Leibniz rule follows suit.

An important property of the Lie derivative is the following: The Lie derivative of a differential form (of any order) commutes with the exterior derivative, i.e.:

$$L_V(d\omega) = d(L_V \omega), \tag{4.11}$$

for all vector fields $\mathbf{V}$ and for all differential forms ω.

The Lie derivative in components

Taking advantage of the Leibniz rule, it is not difficult to calculate the components of the Lie derivative of a tensor in a given coordinate system x^i, provided the components of the Lie derivative of the base

vectors $\frac{\partial}{\partial x^i}$ and the base one-forms dx^i are known. A direct application of the formula (2.10) for the components of the Lie bracket, yields:

$$L_V\left(\frac{\partial}{\partial x^i}\right) = -\frac{\partial V^k}{\partial x^i}\ \frac{\partial}{\partial x^k}. \tag{4.12}$$

To obtain the Lie derivative of dx^i we recall that the action of dx^i (as a covector) on $\frac{\partial}{\partial x^j}$ is simply δ^i_j, whose Lie derivative vanishes. This action can be seen as the contraction of their tensor product. We obtain, therefore:

$$0 = L_V\langle dx^i, \frac{\partial}{\partial x^j}\rangle = \langle dx^i, -\frac{\partial V^k}{\partial x^j}\ \frac{\partial}{\partial x^k}\rangle + \langle L_v dx^i, \frac{\partial}{\partial x^j}\rangle, \tag{4.13}$$

whence:

$$L_V dx^i = \frac{\partial V^i}{\partial x^k} dx^k. \tag{4.14}$$

3.4.2 One-parameter subgroups of a Lie group

A topological group is a *Lie group* if its underlying set is smooth manifold and the group operations are smooth. A one-parameter subgroup of a Lie group $\mathcal{G}$ is a differentiable curve:

$$\begin{array}{rccl} \gamma : \mathbb{R} & \longrightarrow & \mathcal{G} \\ t & \mapsto & g(t), \end{array} \tag{4.15}$$

satisfying:

$$g(0) = e, \tag{4.16}$$

and

$$g(t_1)\ g(t_2) = g(t_1 + t_2) \qquad \forall t_1, t_2 \in \mathbb{R}. \tag{4.17}$$

If the group $\mathcal{G}$ acts (on the left, say) on a manifold $\mathcal{M}$, the composition of this action with a one-parameter subgroup determines a one-parameter group of transformations of $\mathcal{M}$, namely:

$$\gamma_t(p) = L_{g(t)}(p) \qquad p \in \mathcal{M}. \tag{4.18}$$

From Equation (3.34), we know that associated with this flow there exists a unique vector field $\mathbf{v}^\gamma$. More precisely, we have:

$$\mathbf{v}^\gamma(p) = \left.\frac{d\gamma_t(p)}{dt}\right|_{t=0} \tag{4.19}$$

Fixing the point p, we obtain the map L_p from the group to the manifold. The image of the curve γ under this map is obtained by composition as:

$$t \mapsto L_p(g(t)) = L(g(t), p) = L_{g(t)}(p) = \gamma_t(p), \tag{4.20}$$

where we have used Equation (4.18). In other words, the image of the curve γ (defining the one-parameter subgroup) by the map L_p is nothing but the integral curve of the flow passing through p. By definition of derivative of a map between manifolds, we conclude that the tangent $\mathbf{g}$ to the the one parameter subgroup γ at the group identity e is mapped by L_{p*} to the vector $\mathbf{v}^\gamma(p)$:

$$\mathbf{v}^\gamma(p) = \left(L_{p*}\right)_e\ \mathbf{g}. \tag{4.21}$$

This means that a one-parameter subgroup $g(t)$ appears to be completely characterized by its tangent vector $\mathbf{g}$ at the group identity. We will shortly confirm this fact more fully. The vector field induced on $\mathcal{M}$ by a one-parameter subgroup is called the *fundamental vector field* associated with the corresponding vector $\mathbf{g}$ at the group identity.

Let us now identify the manifold $\mathcal{M}$ with the group $\mathcal{G}$ itself. In this case, we have, as already discussed, two canonical actions giving rise to the left and right translations of the group. We want to reinterpret Equation (4.21) in this particular case. For this purpose, and to avoid the notational ambiguity alluded to above, we restore the fully fledged notation for the action as a function of two variables. We thus obtain:

$$\mathbf{v}^\gamma(h) = \left(\frac{\partial L(g,h)}{\partial g}\right)_{g=e} \mathbf{g}. \tag{4.22}$$

Notice that, somewhat puzzlingly, but consistently, this can also be written as:

$$\mathbf{v}^\gamma(h) = \left(R_{h*}\right)_e \mathbf{g}. \tag{4.23}$$

Thus, when defining the action of a one-parameter subgroup from the left, it is the right action whose derivative delivers the corresponding vector field, and viceversa.

3.4.3 Left and right invariant vector fields on a Lie group

A vector field $\mathbf{v} : \mathcal{G} \to T\mathcal{G}$ is said to be *left invariant* if:

$$\mathbf{v}(L_g h) = L_{g*}\mathbf{v}(h) \qquad \forall g, h \in \mathcal{G}. \tag{4.24}$$

In other words, vectors at one point are dragged to vectors at any other point by the derivative of the appropriate left translation. A similar definition, but replacing L with R applies to *right invariant vector fields*.

A vector field is left invariant if, and only if:

$$\mathbf{v}(g) = \left(L_{g*}\right)_e \mathbf{v}(e) \quad \forall g \in \mathcal{G}. \tag{4.25}$$

Another way of expressing this result is by saying that there exists a one-to-one correspondence between the set of left (or right) invariant vector fields on $\mathcal{G}$ and the tangent space $T_e\mathcal{G}$ at the group identity. This correspondence is linear. Moreover, one can show that the Lie bracket of two left- (right-) invariant vector fields is itself left- (right-) invariant. The set $\mathfrak{g}$ of left invariant vector fields (or, equivalently, the tangent space $T_e\mathcal{G}$) with the Lie bracket operation is called the *Lie algebra of the group* $\mathcal{G}$. From an intuitive point of view, the elements of the Lie algebra of a Lie group represent infinitesimal approximations, which Sophus Lie himself called *infinitesimal generators* of the elements of the group. Although the infinitesimal generators are in principle commutative (sum of vectors), the degree of non-commutativity of the actual group elements is captured, to first order, by the Lie bracket.

3.5 Integration

3.5.1 Integration of n-forms in $\mathbb{R}^n$

The simplest n-dimensional manifold is $\mathbb{R}^n$ itself with the standard topology and the standard notion of differentiability. Accordingly, we present the standard notion of integration over a domain of $\mathbb{R}^n$ in terms of differential forms so as to be able to extend this notion to arbitrary manifolds.

Let $x^1, ..., x^n$ be the standard global chart of $\mathbb{R}^n$, and let ω be a smooth n-form defined over some open set $\mathcal{D} \subset \mathbb{R}^n$. There exists, then, a smooth function $f : \mathcal{D} \to \mathbb{R}$ such that:

$$\omega = f \, dx^1 \wedge ... \wedge dx^n. \tag{5.1}$$

For any regular domain of integration $\mathcal{A} \subset \mathcal{D}$ we define:

$$\int_{\mathcal{A}} \omega := \underbrace{\int \int ... \int}_{\mathcal{A}} f dx^1 dx^2 ... dx^n, \tag{5.2}$$

where the right-hand side is the ordinary n-fold Riemann integral in $\mathbb{R}^n$.

It is important to check that this definition is independent of the coordinate system adopted in $\mathcal{D}$. For this purpose, let:

$$\phi : \mathcal{D} \longrightarrow \mathbb{R}^n, \tag{5.3}$$

be a coordinate transformation expressed in components as the n smooth functions:

$$x^1, ..., x^n \mapsto y^1(x^1, ..., x^n), ..., y^n(x^1, ..., x^n). \tag{5.4}$$

Recall that for (5.4) to qualify as a coordinate transformation, the Jacobian determinant:

$$J = \det \left[\frac{\partial(y^1, ..., y^n)}{\partial(x^1, ..., x^n)} \right], \tag{5.5}$$

must be non-zero throughout $\mathcal{D}$. For definiteness, we will assume that it is strictly positive (so that the change of coordinates is orientation preserving). According to the formulas of transformation of variables under a multiple Riemann integral, we must have:

$$\underbrace{\int \int ... \int}_{\mathcal{A}} f(x^i) \, dx^1 ... dx^n = \underbrace{\int \int ... \int}_{\mathcal{A}} f(x^i(y^j)) \, J^{-1} dy^1 ... dy^n. \tag{5.6}$$

But, according to Equation (3.13) the representation of ω in the new coordinate system is precisely:

$$\omega = f(x^i(y^j)) \, J^{-1} dy^1 \wedge ... \wedge dy^n, \tag{5.7}$$

which shows that the definition (5.2) is indeed independent of the coordinate system adopted in $\mathcal{D}$.

A more fruitful way to exploit the coordinate independence property is to regard $\phi : \mathcal{D} \to \mathbb{R}^n$ not as a mere coordinate transformation but as an actual change of the domain of integration. In this case, the transformation formula is interpreted readily in terms of the pull-back of ω as:

$$\int_{\phi(\mathcal{A})} \omega = \int_{\mathcal{A}} \phi^*(\omega), \tag{5.8}$$

for every n-form ω defined over an open set containing $\phi(\mathcal{A})$.

3.5.2 Integration of forms on oriented manifolds

Let $\mathcal{M}$ be an oriented m-dimensional manifold and let $(\mathcal{U}, \psi)$ be a (consistently) oriented chart. The integral of an m-form ω over $\mathcal{U}$ is defined as:

$$\int_{\mathcal{U}} \omega := \int_{\psi(\mathcal{U})} \left(\psi^{-1}\right)^* \omega. \tag{5.9}$$

Notice that the right-hand side is a standard Riemann (or Lebesgue) integral of a function in $\mathbb{R}^m$, according to Equation (5.2). In other words, given an m-form defined over the domain of a chart in an m-dimensional manifold, we simply pull back this form to the codomain of the chart (an open subset of $\mathbb{R}^m$) and integrate. The result is independent of the chart used.

To integrate over a domain covered by more than one chart, we need to use the concept of *partition of unity*, whose detailed presentation we omit. Briefly and imprecisely stated, a partition of unity is a collection of real-valued non-negative smooth functions, one for each of the members of an open cover and vanishing outside it. Moreover, we assume that the cover is *locally finite*, in the sense that each point of the manifold belongs to only a finite number of members of the cover. Finally, the (finite) sum of all these functions at each point is equal to 1 (hence the name). It can be shown that partitions of unit exist provided the manifold is *paracompact*, namely, it admits a locally finite cover. Denoting by ϕ_i the functions making up the partition of unity, we define the integral as:

$$\int_{\mathcal{M}} \omega = \sum_i \int_{\mathcal{M}} \phi_i \omega, \tag{5.10}$$

where the integrand on the right-hand side is just the product of a function times a differential form. Each integral on the right-hand side is well defined by (5.9).

For the definition implied by Equation (5.10) to make sense, we must prove that the result is independent of the choice of charts and of the choice of partition of unity. This can be done quite straightforwardly by expressing each integral of the right-hand side of one choice in terms of the quantities of the second choice, and viceversa.

3.5.3 Stokes' theorem

The boundary of an n-dimensional manifold $\mathcal{M}$ will denoted by $\partial\mathcal{M}$, which we consider as a manifold of dimension $n-1$. For example, $\mathcal{M}$ may be an open ball in $\mathbb{R}^n$ and $\partial\mathcal{M}$ the bounding spherical surface. The boundary of an oriented manifold can be consistently oriented. Given an $(n-1)$-form ω on $\mathcal{M}$, it makes sense to calculate its integral over the (oriented) boundary $\partial\mathcal{D}$. Stokes' theorem asserts that:

$$\int_{\partial\mathcal{D}} \omega = \int_{\mathcal{D}} d\omega. \tag{5.11}$$

We omit the proof and limit ourselves to remark that this elegant formula encompasses all the integral theorems of ordinary vector calculus.

3.6 Distributions and the theorem of Frobenius

We have mentioned in Section **??** the notion of curvature of a connection as an indication of how the parallel transport of an entity along a curve depends on the curve itself. Before giving a precise definition of this concept, however, it may prove useful to introduce the more general concept of *involutivity* of a distribution. The reason for this digression is that a connection can always be regarded as a (horizontal) distribution. A k-dimensional distribution of an m-dimensional manifold $\mathcal{M}$ (with $m \geq k$) is defined as a smooth assignment of a k-dimensional subspace $\mathcal{D}_x$ of the tangent space $T_x\mathcal{M}$ to each point $x \in \mathcal{M}$. A fundamental question in the theory of distributions is whether or not there exist *integral embedded*

submanifolds, namely, embedded submanifolds of dimension k whose tangent space at each point x coincides with $\mathcal{D}_x$.

An *embedded submanifold* of dimension k is defined as a subset $\mathcal{S} \subset \mathcal{M}$ such for each point $s \in \mathcal{S}$ one can find a chart of $\mathcal{M}$ with coordinates x^i $(i = 1, ..., m)$ such that s belongs to this chart and such that the intersection of the set $\mathcal{S}$ with the chart coincides with the set obtained by keeping the last $m - k$ coordinates constant. This idea becomes clear if one thinks of the particular case of $\mathbb{R}^2$ as embedded in $\mathbb{R}^3$ with coordinates x, y, z. The equation of the embedded submanifold $\mathbb{R}^2$ can be given as $z = 0$.

In some sense, the question of existence of integral submanifolds can be regarded as a generalization to many dimensions of the question of integrability of systems of ordinary differential equation, which would correspond to the case $k = 1$, namely to the case in which the subspaces of the distribution are mere lines. While in the particular case $k = 1$ we are assured, by the fundamental theorem of ODEs, of the (local) existence of integral curves, the answer in the general case $k > 1$ is usually negative. A k-dimensional distribution is said to be *completely integrable* if for each point of the manifold $\mathcal{M}$ there exists a chart x^i $(i = 1, ..., m)$ such that each set obtained by keeping the last $n - k$ coordinates thereat constant is an integral submanifold (of dimension k). Assume that a completely integrable distribution has been given. Then, according to our definition, the first k natural vectors of the local coordinate system just described belong to the distribution and constitute a basis of $\mathcal{D}_x$ at each point x in the chart. Any vector fields $\mathbf{v}_\alpha$ $(\alpha = 1, ..., k)$ with this property (of constituting a basis of the distribution) is said to *span* the distribution. Within the chart, any vector fields $\mathbf{v}_\alpha$ $(\alpha = 1, ..., k)$ that span the distribution must be expressible, therefore, as

$$\mathbf{v}_\alpha = v_\alpha{}^\beta \, \frac{\partial}{\partial x^\beta}, \tag{6.1}$$

where the summation convention applies for Greek indices within the range $1, ..., k$. We now calculate the Lie bracket of any pair of the spanning vectors as

$$[\mathbf{v}_\alpha, \mathbf{v}_\beta] = v_\alpha{}^\rho \, \frac{\partial v_\beta{}^\sigma}{\partial x^\rho} \, \frac{\partial}{\partial x^\sigma} - v_\beta{}^\sigma \, \frac{\partial v_\alpha{}^\rho}{\partial x^\sigma} \, \frac{\partial}{\partial x^\rho}. \tag{6.2}$$

Notice that, in calculating the Lie brackets, we have used the fact that the components of the vectors $\mathbf{v}_\alpha$ vanish on the natural base vectors $\partial/\partial x^i$ with $i > k$. Moreover, since the given vectors are linearly independent, the matrix with entries $v_\alpha{}^\beta$ is nonsingular. Inverting, therefore, Equation (6.1), we can express the natural base vectors $\partial/\partial x^\alpha$ $(\alpha = 1, ..., k)$ in terms of the vectors $\mathbf{v}_\beta$, with the result that the Lie brackets are themselves linear combinations of these vectors, namely, there exist scalars $C^\gamma_{\alpha\beta}$ such that

$$[\mathbf{v}_\alpha, \mathbf{v}_\beta] = C^\gamma_{\alpha\beta} \, \mathbf{v}\gamma. \tag{6.3}$$

A distribution with this property (namely, that the Lie bracket of any two vector fields in the distribution is also in the distribution) is said to be *involutive*. We have proven, therefore, that every completely integrable distribution is involutive. The converse of this result (that is, that every involutive distribution is completely integrable) is also true, and is the content of the theorem of Frobenius, whose prove we omit.

3.7 Connections

3.7.1 Introduction

All the fibres of a fibre bundle are, by definition, diffeomorphic to each other. In the absence of additional structure, however, there is no canonical way to single out a particular diffeomorphism between fibres.

In the case of a product bundle, for example, such a special choice is indeed available because of the existence of the second projection map onto the typical fibre. In this extreme case, we may say that we are in the presence of a *canonical distant parallelism* in the fibre bundle. An equivalent way to describe this situation is by saying that we have a canonical family of non-intersecting smooth cross sections such that each point in the fibre bundle belongs to one, and only one, of them. In a general fibre bundle we can only afford this luxury non-canonically and locally. A *connection* on a fibre bundle is, roughly speaking, an additional structure defined on the bundle that permits to establish intrinsic fibre diffeomorphisms for fibres lying along curves in the base manifold. In other words, a connection can be described as a curve-dependent parallelism. Given a connection, it may so happen that the induced fibre parallelisms turn out to be curve-independent. A quantitative measure of this property or the lack thereof is provided by the vanishing, or otherwise, of the *curvature* of the connection.

3.7.2 Ehresmann connection

Consider the tangent bundle $T\mathcal{C}$ of the total space $\mathcal{C}$ of an arbitrary fibre bundle $(\mathcal{C}, \pi, \mathcal{B}, \mathcal{F}, \mathcal{G})$, and denote by $\tau_C : T\mathcal{C} \to \mathcal{C}$ its natural projection. If the dimension of the base manifold $\mathcal{B}$ and the typical fibre $\mathcal{F}$ are, respectively, m and n, the dimension of $\mathcal{C}$ is $m+n$, and the typical fibre of $(T\mathcal{C}, \tau_C)$ is $\mathbb{R}^{m+n}$, with structure group $GL(m+n;\mathbb{R})$. At each point $c \in \mathcal{C}$ the tangent space $T_c\mathcal{C}$ has a canonically defined *vertical subspace* V_c, which can be identified with the tangent space $T_c\mathcal{C}_{\pi(c)}$ to the fibre of $\mathcal{C}$ at c. The dimension of V_c is n. A vector in $T_c\mathcal{C}$ belongs to the vertical subspace V_c (or: is *vertical*) if, and only if, its projection by π_* is the zero vector of $T_{\pi(c)}\mathcal{B}$. If a vector in $T_c\mathcal{C}$ is not vertical, there is no canonical way to assign to it a vertical component. It is this deficiency, and only this deficiency, that the Ehresmann connection remedies. Formally, an Ehresmann connection consists of a smooth *horizontal distribution* in $\mathcal{C}$. This is a smooth assignment to each point $c \in \mathcal{C}$ of an (m-dimensional) subspace $H_c \subset T_c\mathcal{C}$ (called the *horizontal subspace at* c), such that:

$$T_c\mathcal{C} = H_c + V_c. \tag{7.1}$$

In this equation, $+$ denotes the direct sum of vector spaces. Each tangent vector $\mathbf{u} \in T_c\mathcal{C}$ is, accordingly, uniquely decomposable as the sum of a horizontal part $h(\mathbf{u})$ and a vertical part $v(\mathbf{u})$. A vector is *horizontal*, if its vertical part vanishes. The only vector that is simultaneously horizontal and vertical is the zero vector. Since H_c and $T_{\pi(c)}\mathcal{B}$ have the same dimension (m), the restriction $\pi_*|_{H_c} : H_c \to T_{\pi(c)}\mathcal{B}$, is a vector-space isomorphism. We denote its inverse by Γ_c. Thus, given a vector $\mathbf{v}$ tangent to the base manifold at a point $b \in \mathcal{B}$, there is a unique horizontal vector: $\Gamma_c\mathbf{v}$ at $c \in \pi^{-1}(\{b\})$ such that $\pi_*(\Gamma_c\mathbf{v}) = \mathbf{v}$. This unique vector is called the *horizontal lift* of $\mathbf{v}$ to c. In particular: $\Gamma_c(\pi_*(\mathbf{u})) = \Gamma_c(\pi_*(h(\mathbf{u}))) = h(\mathbf{u})$. These ideas are schematically illustrated in Figure 3.4.

3.7.3 Parallel transport along a curve

Let:

$$\gamma : (-\epsilon, \epsilon) \longrightarrow \mathcal{B} \tag{7.2}$$

be a smooth curve in the base manifold $\mathcal{B}$ of the fibre bundle $(\mathcal{C}, \pi)$, and let $c \in \mathcal{C}_{\gamma(0)}$ be a point in the fibre at $\gamma(0)$. A *horizontal lift* of γ through c is defined as a curve:

$$\hat{\gamma} : (-\epsilon, \epsilon) \longrightarrow \mathcal{C}, \tag{7.3}$$

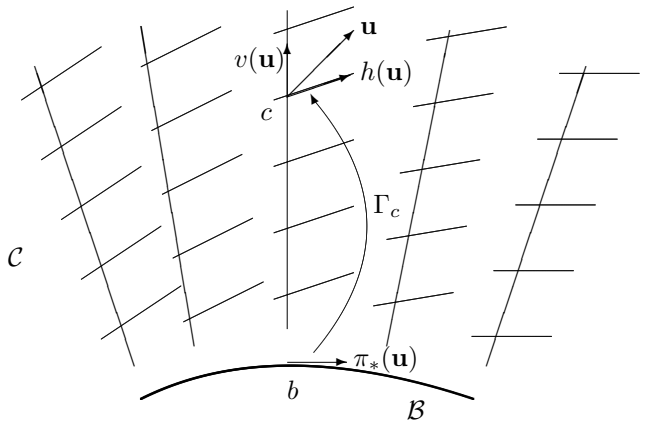

Figure 3.4: Ehresmann connection

such that:

$$\hat{\gamma}(0) = c, \tag{7.4}$$

$$\pi\left(\hat{\gamma}(t)\right) = \gamma(t), \qquad \forall t \in (-\epsilon, \epsilon), \tag{7.5}$$

and

$$\hat{\gamma}'(t) \in H_{\hat{\gamma}(t)} \qquad \forall t \in (-\epsilon, \epsilon), \tag{7.6}$$

where a prime denotes the derivative with respect to the curve parameter t. A horizontal lift is thus a curve that projects onto the original curve and, moreover, has a horizontal tangent throughout.

Consider the "cylindrical" subbundle $\gamma^*\mathcal{C}$ obtained by pulling back the bundle $\mathcal{C}$ to the curve γ or, less technically, by restricting the base manifold to the curve γ. The tangent vector field of γ has a unique horizontal lift at each point of this bundle. In other words, the curve generates a (horizontal) vector field throughout this restricted bundle. By the fundamental theorem of the theory of ODEs, it follows that, at least for small enough ϵ, there is a unique horizontal lift of γ through any given point in the fibre at $\gamma(0)$, namely, the corresponding integral curve of the horizontal vector field. We conclude, therefore, that the horizontal lift of a curve through a point in a fibre bundle exists and is locally unique. As the horizontal curve issuing from c cuts the various fibres lying on γ, the point c is said to undergo a *parallel transport* relative to the given connection and the given curve. Thus, given a point $c \in \mathcal{C}$ and a curve γ through $\pi(c) \in \mathcal{B}$, we obtain a unique parallel transport of c along γ by solving a system of ODEs (so as to travel always horizontally). These concepts are illustrated schematically in Figure 3.5

3.7.4 Connections in principal bundles

A connection in a principal bundle $(\mathcal{P}, \pi, \mathcal{B}, \mathcal{G}, \mathcal{G})$ is an Ehresmann connection which is compatible with the right action R_g of $\mathcal{G}$ on $\mathcal{P}$, namely:

$$(R_g)_*(H_p) = H_{R_g p} \qquad \forall g \in \mathcal{G} \quad p \in \mathcal{P}. \tag{7.7}$$

This condition can be stated verbally as follows: the horizontal distribution is invariant under the group action.

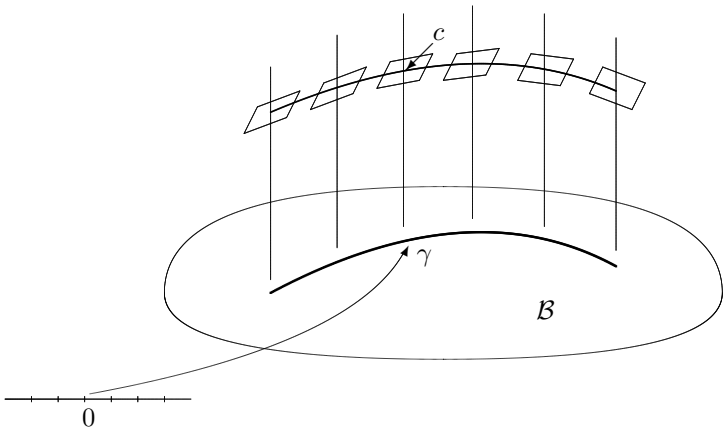

Figure 3.5: Parallel transport along a curve

Recall that the group $\mathcal{G}$ acts freely (to the right) on $\mathcal{P}$. Consequently, the fundamental vector field $\mathbf{v_g}$ associated with any non-zero vector $\mathbf{g}$ in the Lie algebra $\mathfrak{g}$ of $\mathcal{G}$ does not vanish anywhere on $\mathcal{P}$. Moreover, since the action of $\mathcal{G}$ maps fibres into themselves, the fundamental vector fields are all vertical. The correspondence between vectors in the Lie algebra and tangent vectors to the fibre at any point is clearly linear and one-to-one. Since the dimension of $\mathcal{G}$ is equal to the dimension of each fibre, we conclude that the map $\mathfrak{g} \to V_p$ given by: $\mathbf{g} \mapsto \mathbf{v_g}(p)$ is a linear isomorphism between the Lie algebra and each of the vertical subspaces of the principal bundle.

Let $\mathbf{v} \in T\mathcal{P}$ be any tangent vector to the fibre bundle. A connection Γ assigns to it a unique vertical part and, as we have just seen, the action of the group assigns to this vertical part an element of the Lie algebra $\mathfrak{g}$. This means that we have a well-defined linear map:

$$\omega : T\mathcal{P} \longrightarrow \mathfrak{g}, \tag{7.8}$$

associated with a given connection in a principal bundle. This map can be regarded as a *Lie-algebra valued 1-form*. It is called the *connection form* associated with Γ.

3.7.5 Curvature

Suppose that we draw through a point b of the base manifold $\mathcal{B}$ a small closed curve γ. If we now choose a point p in the fibre on b, we have learned that there exists a unique horizontal lift $\tilde{\gamma}$, namely, a horizontal curve containing p and projecting on γ. Is this curve closed? To clarify the meaning of this question and its possible answer, recall that a connection on a principal fibre bundle is a special case of a distribution, which we have called horizontal (the dimension of the horizontal distribution equals the dimension of the base manifold and is thus strictly smaller than the dimension of the fibre bundle, assuming that the typical fibre is of dimension greater than zero). Clearly, if the horizontal distribution is involutive, any horizontal lift of a small curve in the base manifold will lie entirely on an integral surface and, therefore, will be closed. This observation suggests that a measure of the lack of closure of the horizontal lift of closed curves is the fact that the Lie bracket between horizontal vector fields has a vertical component. We want to see now how to extract this information from the connection itself. More particularly, since a connection is specified by its connection form ω, we want to extract this information from ω alone.

Consider two horizontal vector fields $\mathbf{u}$ and $\mathbf{v}$. Let us evaluate the 2-form[3] $d\omega$ on this pair. Using Equation (3.26) we obtain:

$$\langle d\omega \mid \mathbf{u} \wedge \mathbf{v} \rangle = \mathbf{u}\left(\langle \omega \mid \mathbf{v} \rangle\right) - \mathbf{v}\left(\langle \omega \mid \mathbf{u} \rangle\right) - \langle \omega \mid [\mathbf{u}, \mathbf{v}] \rangle, \tag{7.9}$$

which, in view of the fact that $\mathbf{u}$ and $\mathbf{v}$ are assumed to be horizontal, yields:

$$\langle d\omega \mid \mathbf{u} \wedge \mathbf{v} \rangle = -\langle \omega \mid [\mathbf{u}, \mathbf{v}] \rangle. \tag{7.10}$$

The right-hand side of this equation will vanish if, and only if, the Lie bracket is horizontal. This means that we have found a way to extract the right information from ω by just taking its exterior derivative and applying it to two horizontal vector fields. Notice, however, that $d\omega$ can be applied to arbitrary pairs of vector fields, not necessarily horizontal. To formalize this point, given a connection, we define the *exterior covariant derivative* $D\alpha$ of an r-form α as the $(r+1)$-form given by:

$$\langle D\alpha \mid \mathbf{U}_1 \wedge ... \wedge \mathbf{U}_{r+1} \rangle = \langle d\alpha \mid h(\mathbf{U}_1) \wedge ... \wedge h(\mathbf{U}_{r+1}) \rangle, \tag{7.11}$$

where $h(.)$ denotes the horizontal component of a vector. Accordingly, we define the *curvature* 2*-form* Ω of a connection ω on a principal fibre bundle as:

$$\Omega := D\omega. \tag{7.12}$$

3.7.6 Cartan's structural equation

Our definition of curvature, by using both the connection 1-form and the horizontal projection map, is a hybrid that mixes both (equivalent) definitions of a connection. It is possible, on the other hand, to obtain an elegant formula that involves just the connection 1-form. This formula, known as *Cartan's structural equation*, reads:

$$\Omega = d\omega + \frac{1}{2}\,[\omega, \omega], \tag{7.13}$$

or, more precisely, for any two vectors $\mathbf{U}$ and $\mathbf{V}$ at a point[4] of the frame bundle:

$$\langle \Omega \mid \mathbf{U} \wedge \mathbf{V} \rangle = \langle d\omega \mid \mathbf{U} \wedge \mathbf{V} \rangle + \frac{1}{2}\,[\omega(\mathbf{U}), \omega(\mathbf{V})]. \tag{7.14}$$

The proof of this formula, whose details we omit, is based on a careful examination of three cases: (i) $\mathbf{U}$ and $\mathbf{V}$ are horizontal, whereby the formula is obvious; (ii) $\mathbf{U}$ is horizontal and $\mathbf{V}$ is vertical, in which case one can extend them, respectively, to a horizontal and a fundamental (vertical) vector field; (iii) $\mathbf{U}$ and $\mathbf{V}$ are both vertical, in which case they can both be extended to fundamental fields.

3.7.7 Bianchi identities

Unlike the ordinary exterior derivative d, the operator D (of exterior covariant differentiation) is not necessarily nilpotent, namely, in general $D^2 \neq 0$. Therefore, there is no reason to expect that $D\Omega$, which is equal to $D(D\omega)$, will vanish identically. But in fact it does. To see that this is the case, notice that, by definition of D, we need only verify the vanishing of $D\Omega$ on an arbitrary triple of *horizontal* vectors. It can be shown that:

$$D\Omega = 0. \tag{7.15}$$

In terms of components, we obtain differential identities to be satisfied by any curvature form. They are known as the *Bianchi identities*.

[3]Notice that this is a Lie-algebra valued differential form.

[4]Notice that this formula is valid point-wise, since the Lie bracket on the right-hand side is evaluated in the Lie algebra, not in the manifold.

3.7.8 Linear connections

A connection on the bundle of linear frames $F\mathcal{B}$ is called a *linear connection* on $\mathcal{B}$. Among principal bundles, the bundle of linear frames occupies a special position for various reasons. In the first place, the bundle of linear frames is canonically defined for any given base manifold $\mathcal{B}$. Moreover, the associated bundles include all the tensor bundles, thus allowing for a unified treatment of all such entities. Another way to express this peculiar feature of the bundle of linear frames is that, whereas the quantities parallel-transported along curves in a general principal bundle are of a nature not necessarily related to the base manifold, in the case of the bundle of linear frames the quantities transported are precisely the very frames used to express the components of vectors and forms defined on the base manifold. An elegant manifestation of this property is the existence of a canonical 1-form that ties everything together. A direct consequence of the existence of this 1-form is the emergence of the idea of the *torsion* of a connection. We start the treatment of linear connections by lingering for a while on the definition of the canonical 1-form.

3.7.9 The canonical 1-form

Given a tangent vector $\mathbf{v} \in T_x\mathcal{B}$ at a point x in the base manifold, and a point $p \in F_x\mathcal{B}$ in the fibre over x, and recalling that p consists of a frame (or basis) $\{\mathbf{e}_1, ..., \mathbf{e}_m\}$ of $T_x\mathcal{B}$, we can determine uniquely the m components of $\mathbf{v}$ in this frame, namely:

$$\mathbf{v} = v^a \mathbf{e}_a. \tag{7.16}$$

In other words, at each point $p \in F\mathcal{B}$, we have a well-defined non-singular linear map[5]:

$$u(p) : T_{\pi(p)}\mathcal{B} \longrightarrow \mathbb{R}^m. \tag{7.17}$$

The *canonical 1-form* θ on $F\mathcal{B}$ is defined as:

$$\theta(\mathbf{V}) := u(p) \circ \pi_*(\mathbf{V}) \qquad \forall\ \mathbf{V} \in T_p(F\mathcal{B}). \tag{7.18}$$

Note that this is an $\mathbb{R}^m$-valued form. The canonical 1-form of the frame bundle is a particular case of a more general construct known as a *soldering form*.

It may prove instructive to exhibit the canonical form in components. Let $x^1, ..., x^m$ be a local coordinate system on $\mathcal{U} \subset \mathcal{B}$. Every frame $\{\mathbf{e}_1, ..., \mathbf{e}_m\}$ at $x \in \mathcal{U}$ can be expressed uniquely by means of a non-singular matrix with entries x^i_j as:

$$\mathbf{e}_a = x^i_a \frac{\partial}{\partial x^i}. \tag{7.19}$$

This means that the $m + m^2$ functions $\{x^i, x^i_a\}$ constitute a coordinate system for the linear frame bundle $\pi^{-1}(\mathcal{U})$. We call it the coordinate system *induced* by x^i. The projection map $\pi : F\mathcal{B} \to \mathcal{B}$ has the coordinate representation:

$$(x^i, x^i_a) \mapsto \pi(x^i, x^i_a) = (x^i), \tag{7.20}$$

with some notational abuse.

Consider now the tangent bundle $TF(\mathcal{B})$ with projection $\tau : TF(\mathcal{B}) \to F(\mathcal{B})$. The coordinate system $\{x^i, x^i_a\}$ induces naturally a coordinate system in $TF(\mathcal{B})$. A vector $\mathbf{X} \in TF(\mathcal{B})$ is expressed in these coordinates as follows:

$$\mathbf{X} \mapsto \left(x^i, x^i_a, X^i \frac{\partial}{\partial x^i} + X^i_a \frac{\partial}{\partial x^i_a} \right) = \left(x^i, x^i_a, X^i, X^i_a \right). \tag{7.21}$$

[5]This map is, in fact, an alternative definition of a linear frame at a point of a manifold $\mathcal{B}$.

The derivative of the projection π is a map: $\pi_* TF(\mathcal{B}) \to T\mathcal{B}$. Its coordinate representation is:

$$\left(x^i, x^i_a, X^i, X^i_a\right) \mapsto \left(x^i, X^i\right). \tag{7.22}$$

The map u defined in Equation (7.17) is given in coordinates by:

$$[u(x^i, x^i_a)](x^j, w^j) = x_i^{-a} w^i \quad (a = 1, ..., m), \tag{7.23}$$

where we have denoted by x_i^{-a} the entries of the inverse of the matrix with entries x^i_a. Combining (7.22) and (7.23), we obtain from (7.18) the following coordinate representation of the $(\mathbb{R}^m)$-valued canonical form θ:

$$\theta^a = x_i^{-a}\, dx^i \qquad (a = 1, ..., m). \tag{7.24}$$

3.7.10 The Christoffel symbols

The canonical form θ exists independently of any connection. Let us now introduce a connection on $F(\mathcal{B})$, that is, a linear connection on $\mathcal{B}$. If we regard a connection as a horizontal distribution, there must exist non-singular linear maps $\Gamma(x, p)$ from each $T_x\mathcal{B}$ to each of the tangent spaces $T_pF(\mathcal{B})$ (with $\pi(p) = x$) defining the distribution. Noticing that the same distribution may correspond to an infinite number of such maps, we pin down a particular one by imposing the extra condition that they must be also horizontal lifts. In other words, we demand that:

$$\pi_* \circ \Gamma(x, p) = id_{T_x\mathcal{B}}. \tag{7.25}$$

The implication of this condition is that, when written in components, we must have:

$$\Gamma(x, p)\left(v^i \frac{\partial}{\partial x^i}\right) = v^i \frac{\partial}{\partial x^i} - \hat{\Gamma}^j_{ia}(x, p)\; v^i \frac{\partial}{\partial x^j_a}, \tag{7.26}$$

where $\hat{\Gamma}^j_{ia}(x, p)$ are smooth functions of x and p. The minus sign is introduced for convenience.

These functions, however, cannot be arbitrary, since they must also satisfy the compatibility condition (7.7). It is not difficult to verify that this is the case if, and only if, the functions Γ^j_{ik} defined by:

$$\Gamma^j_{ik} := \hat{\Gamma}^j_{ia}(x, p)\; x_k^{-a}(p), \tag{7.27}$$

are independent of p along each fibre.

We conclude that a linear connection is completely defined (on a given coordinate patch) by means of m^3 smooth functions. These functions are known as the *Christoffel symbols* of the connection.

3.7.11 Parallel transport and the covariant derivative

Now that we are in possession of explicit coordinate expressions for the horizontal distribution, we can write explicitly the system of ODEs that effects the horizontal lift of a curve in $\mathcal{B}$. A solution of this system is a one-parameter family of frames being parallel-transported along the curve. Let the curve γ in the base manifold be given by:

$$x^i = x^i(t). \tag{7.28}$$

On this curve, the connection symbols are available as functions of t, by composition. The non-trivial part of the system of equations is given by:

$$\frac{dx_a^i(t)}{dt} = -\Gamma_{jk}^i(t)\ x_a^k(t)\ \frac{dx^j(t)}{dt}. \tag{7.29}$$

The local solution of this system with given initial condition (say, $x_a^i(0) = \bar{x}_a^i$) is the desired curve in $F(\mathcal{B})$, representing the parallel transport of the initial frame along the given curve.

Let now $\bar{\mathbf{v}}$ be a vector in $T_{x^i(0)}\mathcal{B}$, that is, a vector at the initial "time" $t = 0$. We say that the curve $\mathbf{v} = \mathbf{v}(t)$ in $T\mathcal{B}$ is the parallel transport of $\bar{\mathbf{v}}$ if it projects on γ, with $\mathbf{v}(0) = \bar{\mathbf{v}}$, and if the components of $\mathbf{v}(t)$ in a parallel-transported frame along γ are constant[6]. For this definition to make sense, we must make sure that the constancy of the components is independent of the particular initial frame chosen. This, however, is a direct consequence of the fact that our linear connection is, by definition, consistent with the right action of the group.

To obtain the system of ODEs corresponding to the parallel transport of $\bar{\mathbf{v}}$ along γ, we enforce the constancy conditions:

$$v^i(t)\ x_i^{-a}(t) = C^a, \tag{7.30}$$

where each C^a $(a = 1, ..., m)$ is a constant and where v^i denotes components in the coordinate basis. Differentiating this equation with respect to t and invoking Equation (7.29), we obtain:

$$\frac{dv^i}{dt}\ +\ \Gamma_{jk}^i\ \frac{dx^j}{dt}\ v^k\ =\ 0. \tag{7.31}$$

A vector field along γ satisfying this equation is said to be *covariantly constant* or *parallel along* γ. For a given vector field $\mathbf{w}$ on $\mathcal{B}$, the expression on the left-hand side makes sense in a point-wise manner whenever a vector $\mathbf{u}$ is defined at a point (whereby the curve γ can be seen as a representative at $t = 0$). The expression:

$$\nabla_{\mathbf{u}}\mathbf{w} := \left(\frac{dw^i}{dt}\ +\ \Gamma_{jk}^i\ u^j\ w^k\right)\ \frac{\partial}{\partial x^i}, \tag{7.32}$$

is called the *covariant derivative* of $\mathbf{v}$ in the direction of $\mathbf{u}$. From the treatment above, it can be seen that the covariant derivative is precisely the limit:

$$\nabla_{\mathbf{u}}\mathbf{w} = \lim_{t\to 0}\ \frac{\rho_{t,0}\mathbf{w} - \mathbf{w}(0)}{t}, \tag{7.33}$$

where $\rho(a, b)$ denotes the parallel transport along γ from $t = b$ to $t = a$.

3.7.12 Geodesics

A (parametrized) curve γ in a manifold $\mathcal{B}$ with a linear connection Γ is a *geodesic* if its tangent vector field γ' is covariantly constant along γ. In coordinates, a curve $x^i = x^i(t)$ is a geodesic if, in accordance with Equation (7.31), it satisfies the following system of second-order ODEs:

$$\frac{d^2x^i}{dt^2}\ +\ \Gamma_{jk}^i\ \frac{dx^j}{dt}\ \frac{dx^k}{dt}\ =\ 0. \tag{7.34}$$

[6]The same criterion for parallel transport that we are using for the tangent bundle can be used for any associated bundle.

3.7.13 Curvature and torsion

To obtain an explicit equation for the curvature form Ω, we should start by elucidating the connection form ω on the basis of the connection symbols Γ. Given a vector $\mathbf{X} \in T_pF\mathcal{B}$, we know that its horizontal component $h(\mathbf{X})$ is given by:

$$h(\mathbf{X}) = \Gamma(\pi(p), p) \circ \pi_* (\mathbf{X}). \tag{7.35}$$

Its vertical component must, therefore, be given by:

$$v(\mathbf{X}) = \mathbf{X} - h(\mathbf{X}) = \mathbf{X} - \Gamma(\pi(p), p) \circ \pi_* (\mathbf{X}). \tag{7.36}$$

Recall that the connection form ω assigns to $\mathbf{X}$ the vector in $\mathfrak{g}$ such that $v(\mathbf{X})$ belongs to its fundamental vector field. Let the coordinates of p be (x^i, x^i_a). The right action of $GL(m; \mathbb{R}$ is given by:

$$(R_g(p))^i_a = x^i_b \; g^b_a, \tag{7.37}$$

where we have shown only the action on the fibre component and where g^b_a is the matrix corresponding to $g \in GL(m; \mathbb{R})$. Consequently, if $g(t)$ is a one-parameter subgroup represented by the vector:

$$\hat{g}^a_b = \left. \frac{dg^a_b(t)}{dt} \right|_{t=0}, \tag{7.38}$$

the value of the corresponding fundamental vector field at p is:

$$\tilde{g}^i_a = x^i_b \; \hat{g}^b_a. \tag{7.39}$$

The coordinate expression of Equation (7.36) is:

$$(v(\mathbf{X}))^i_a = X^i_a - h(\mathbf{X}) = \mathbf{X} - \Gamma(\pi(p), p) \circ \pi_* (\mathbf{X}). \tag{7.40}$$

Let the main part of the vector $\mathbf{X}$ be given by:

$$\mathbf{X} = v^i \frac{\partial}{\partial x^i} + X^i_a \frac{\partial}{\partial x^i_a}. \tag{7.41}$$

Then, Equation (7.40) delivers:

$$(v(\mathbf{X}))^i_a = X^i_a + \Gamma^j_{ik} \; v^i \; x^k_a. \tag{7.42}$$

According to Equation (7.39), the corresponding element of the Lie algebra is:

$$\hat{g}^b_a = \left(X^j_a + \Gamma^j_{ik} \; v^i \; x^k_a \right) \; x^{-b}_j. \tag{7.43}$$

Accordingly, the Lie-algebra valued connection form ω is given by:

$$\omega^b_a = \Gamma^j_{ik} \; x^k_a \; x^{-b}_j \; dx^i \; + \; x^{-b}_j \; dx^j_a. \tag{7.44}$$

The exterior derivative is given by:

$$\begin{aligned} d\omega^b_a &= \frac{\partial \Gamma^j_{ik}}{\partial x^m} \; x^k_a \; x^{-b}_j \; dx^m \wedge dx^i \\ &+ \Gamma^j_{ik} x^{-b}_j \; dx^k_a \wedge dx^i - \Gamma^j_{ik} x^k_a x^{-b}_s x^{-c}_j dx^s_c \wedge dx^i - x^{-c}_j x^{-b}_s \; dx^s_c \wedge dx^j_a. \end{aligned} \tag{7.45}$$

A vector such as (7.41) has the following horizontal component:

$$h(\mathbf{X}) = v^i \frac{\partial}{\partial x^i} - \Gamma^j_{ik} x^k_a v^i \frac{\partial}{\partial x^j_a}. \tag{7.46}$$

With a similar notation, the horizontal component of another vector $\mathbf{Y}$ is given by:

$$h(\mathbf{Y}) = w^i \frac{\partial}{\partial x^i} - \Gamma^j_{ik} x^k_a w^i \frac{\partial}{\partial x^j_a}. \tag{7.47}$$

Consider now the following evaluations:

$$\langle dx^j \wedge dx^i \mid h(\mathbf{X}) \wedge h(\mathbf{Y}) \rangle = v^j w^i - v^i w^j, \tag{7.48}$$

$$\langle dx^k_a \wedge dx^i \mid h(\mathbf{X}) \wedge h(\mathbf{Y}) \rangle = -\Gamma^k_{rs} x^s_a \, (v^r w^i - v^i w^r), \tag{7.49}$$

and

$$\langle dx^j_c \wedge dx^s_a \mid h(\mathbf{X}) \wedge h(\mathbf{Y}) \rangle = -\Gamma^j_{rn} x^n_c \, \Gamma^s_{ik} x^k_a \, (v^r w^i - v^i w^r). \tag{7.50}$$

Putting all these results together, we obtain:

$$\langle \Omega \mid \mathbf{X} \wedge \mathbf{Y} \rangle = \langle \omega \mid h(\mathbf{X}) \wedge h(\mathbf{Y}) \rangle = x^k_a x^{-b}_j R^j_{kri} v^r w^i, \tag{7.51}$$

where

$$R^j_{kri} = \frac{\Gamma^j_{ik}}{\partial x^r} - \frac{\Gamma^j_{rk}}{\partial x^i} + \Gamma^j_{rh} \Gamma^h_{ik} - \Gamma^j_{ih} \Gamma^h_{rk} \tag{7.52}$$

is called the *curvature tensor* of the linear connection.

In analogy with the concept of curvature form, we define the *torsion form* of a connection as:

$$\Theta = D\theta. \tag{7.53}$$

Notice that the coupling with the connection is in the fact that the operator D is the exterior *covariant* derivative, which involves the horizontal projection. To understand the meaning of the torsion, consider a case in which the curvature vanishes. This means that there exists a *distant* (or curve independent) parallelism in the manifold $\mathcal{B}$. Thus, fixing a basis of the tangent space at any one point $x_0 \in \mathcal{B}$, a field of bases is uniquely determined at all other points. The question that the torsion tensor addresses is the following: does there exist a coordinate system such that these bases coincide at each point with its natural basis? An interesting example can be constructed in $\mathbb{R}^3$ as follows. Starting from the standard coordinate system, move up the x^3 axis and, while so doing, apply a linearly increasing rotation to the horizontal planes, imitating the action of a corkscrew. Thus, we obtain a system of (orthonormal) bases which are perfectly Cartesian plane by horizontal plane, but twisted with respect to each other as we ascend. These frames can be used to define a distant parallelism (two vectors are parallel if they have the same components in the local frame). It is not difficult to show (or to see intuitively) that there is no coordinate system that has these as natural bases (use, for example, Frobenius' theorem). This example explains the terminology of "torsion".

To obtain the coordinate expression of the torsion form, we start by calculating the exterior derivative of Equation (7.24) as:

$$d\theta^a = dx^{-a}_i \ \wedge \ dx^i = -x^{-a}_j x^{-b}_i dx^j_b \ \wedge \ dx^i. \tag{7.54}$$

Using Equation (7.49), we obtain:

$$\langle D\theta \mid \mathbf{X} \wedge \mathbf{Y} \rangle = \langle d\theta \mid h(\mathbf{X}) \wedge h(\mathbf{Y}) \rangle = x^{-a}_j \, T^j_{ri} \, v^r w^i, \tag{7.55}$$

where

$$T^j_{ri} := \Gamma^j_{ri} - \Gamma^j_{ir} \tag{7.56}$$

are the components of the *torsion tensor* of the connection.

3.8 Riemannian manifolds

3.8.1 Inner-product spaces

We have come a long way without the need to speak about metric concepts, such as the length of a vector or the angle between two vectors. That even the concept of power of a force can be introduced without any metric background may have seemed somewhat surprising, particularly to those accustomed to hear about "the magnitude of the force multiplied by the magnitude of the velocity and by the cosine of the angle they form". It is very often the case in applications to particular fields (Mechanics, Theoretical Physics, Chemistry, Engineering, and so on) that there is much more structure to go around than really needed to formulate the basic concepts. For the particular application at hand, there is nothing wrong in taking advantage of this extra structure. Quite to the contrary: the extra structure may be the carrier of implicit assumptions that permit, consciously or not, the formulation of the physical laws. The most dramatic example is perhaps the adherence to Euclidean Geometry as the backbone of Newtonian Physics. On the other hand, the elucidation of the minimal (or nearly so) structure necessary for the formulation of a fundamental notion, has proven time and again to be the beginning of an enlightenment that can lead to further developments and, no less importantly, to a better insight into the old results.

We have seen how the concept of the space dual to a given vector space arises naturally from the consideration of linear functions on the original vector space. On the other hand, we have learned that, intimately related as they are, there is no natural isomorphism between these two spaces. In other words, there is no natural way to associate a covector to a given vector, and viceversa. In Newtonian Mechanics, however, the assumption of a Euclidean metric, whereby the theorem of Pythagoras holds globally, provides such identification. In Lagrangian Mechanics, it is the kinetic energy of the system that can be shown to provide such extra structure, at least locally. In General Relativity, this extra structure (but of a somewhat different nature) becomes the main physical quantity to be found by solving Einstein's equations. In all these cases, the identification of vectors with covectors is achieved by means of the introduction of a new operation called an *inner product* (or a *dot product* or, less felicitously, a *scalar product*).

A vector space V is said to be an *inner-product space* if it is endowed with an operation (called an inner product):

$$\begin{aligned} \cdot : V \times V &\longrightarrow \mathbb{R} \\ (\mathbf{u}, \mathbf{v}) &\mapsto \mathbf{u} \cdot \mathbf{v}, \end{aligned} \tag{8.1}$$

satisfying the following properties[7]:

(1) Commutativity:

$$\mathbf{u} \cdot \mathbf{v} = \mathbf{v} \cdot \mathbf{u}, \quad \forall\, \mathbf{u}, \mathbf{v} \in V; \tag{8.2}$$

(2) Bi-linearity:[8]

$$(\alpha \mathbf{u}_1 + \beta \mathbf{u}_2) \cdot \mathbf{v} = \alpha(\mathbf{u}_1 \cdot \mathbf{v}) + \beta(\mathbf{u}_2 \cdot \mathbf{v}), \quad \forall\, \alpha, \beta \in \mathbb{R},\ \mathbf{u}_1, \mathbf{u}_2, \mathbf{v} \in V; \tag{8.3}$$

[7] It is to be noted that in the case of a complex vector space, such as in Quantum Mechanics applications, these properties need to be altered somewhat.

[8] The term bi-linearity refers to the fact that the inner product is linear in each of its two arguments. Nevertheless, given that we have already assumed commutativity, we need only to show linearity with respect to one of the arguments.

(3) Positive definiteness:[9]

$$\mathbf{v} \neq \mathbf{0} \implies \mathbf{v} \cdot \mathbf{v} > 0. \tag{8.4}$$

One can show that $\mathbf{0} \cdot \mathbf{v} = 0$, for all $\mathbf{v}$. The *magnitude* or *length* of a vector $\mathbf{v}$ is defined as the non-negative number $\sqrt{\mathbf{v} \cdot \mathbf{v}}$. Two vectors $\mathbf{u}, \mathbf{v} \in V$ are called *orthogonal* (or *perpendicular*) to each other if $\mathbf{u} \cdot \mathbf{v} = 0$.

We want now to show how the existence of an inner product induces an isomorphism between a space and its dual (always in the finite-dimensional case). Let $\mathbf{v} \in V$ be a fixed element of V. By the linearity of the inner product, the product $\mathbf{v} \cdot \mathbf{u}$ is linear in the second argument. Accordingly, we define the covector $\omega_v \in V^*$ corresponding to the vector $\mathbf{v} \in V$, by:

$$\langle \omega_v, \mathbf{u} \rangle \ := \ \mathbf{v} \cdot \mathbf{u}, \quad \forall \, \mathbf{u} \in V. \tag{8.5}$$

It is not difficult to prove that this linear map from V to V^* is one-to-one and that, therefore, it constitutes an isomorphism between V and V^*. We conclude that in an inner product space there is no need to distinguish notation-wise between vectors and covectors.

We call *reciprocal basis* the basis of V that corresponds to the dual basis in the isomorphism induced by the inner product. We already know that the dual basis operates on vectors in the following way:

$$\langle \mathbf{e}^i, \mathbf{v} \rangle = v^i, \quad \forall \mathbf{v} \in V, \tag{8.6}$$

where v^i is the i-th component of $\mathbf{v} \in V$ in the basis $\{\mathbf{e}_j\}$ $(j = 1, ...n)$. The reciprocal basis, therefore, consists of *vectors* $\{\mathbf{e}^j\}$ $(j = 1, ...n)$ such that:

$$\mathbf{e}^i \cdot \mathbf{v} = v^i, \quad \forall \mathbf{v} \in V. \tag{8.7}$$

Let the components of the reciprocal base vectors be expressed as:

$$\mathbf{e}^i = g^{ij} \mathbf{e}_j. \tag{8.8}$$

In other words, we denote by g^{ij} the j-th component of the i-th member of the reciprocal basis we are seeking. It follows from (8.7) that:

$$\mathbf{e}^i \cdot \mathbf{v} = (g^{ij} \mathbf{e}_j) \cdot (v^k \mathbf{e}_k) = g^{ij} \ (\mathbf{e}_j \cdot \mathbf{e}_k) \ v^k = v^i, \quad \forall v^k \in \mathbb{R}. \tag{8.9}$$

Looking at the very last equality, it follows that

$$g^{ij} \ (\mathbf{e}_j \cdot \mathbf{e}_k) \ = \ \delta^i_k. \tag{8.10}$$

Indeed, regarded as a matrix equation, (8.9) establishes that the matrix with entries $[g^{ij} \ (\mathbf{e}_j \cdot \mathbf{e}_k)]$ (summation convention understood), when multiplied by an arbitrary column-vector, leaves it unchanged. It follows that this matrix must be the identity. This is only possible if the matrix with entries:

$$g_{ij} = \mathbf{e}_i \cdot \mathbf{e}_j, \tag{8.11}$$

is the inverse of the matrix with entries g^{ij}. So, the procedure to find the reciprocal basis is the following: (i) Construct the (symmetric) square matrix with entries $g_{ij} = \mathbf{e}_i \cdot \mathbf{e}_j$; (ii) Invert this matrix to obtain the matrix with entries g^{ij}; (iii) Define $\mathbf{e}^i = g^{ij} \mathbf{e}_j$. Note that the *metric matrix* $\{g_{ij}\}$ is always invertible, as it follows from the linear independence of the basis.

[9] In Relativity this property is removed.

A basis of an inner-product space is called *orthonormal* if all its members are of unit length and mutually orthogonal. The reciprocal of an orthonormal basis coincides with the original basis.

Having identified an inner-product space with its dual, and having brought back the dual basis to the original space under the guise of the reciprocal basis, we have at our disposal contravariant and covariant components of vectors. Recall that before the introduction of an inner product, the choice of a basis in V condemned vectors to have contravariant components only, while the components of covectors were covariant.

Starting from $\mathbf{v} = v^i \mathbf{e}_i = v_i \mathbf{e}^i$ and using Equations (8.10, 8.11), the following formulas can be derived:

$$v^i = g^{ij} v_j, \tag{8.12}$$

$$v_i = g_{ij} v^j, \tag{8.13}$$

$$\mathbf{e}_i = g_{ij} \mathbf{e}^j, \tag{8.14}$$

$$v^i = \mathbf{v} \cdot \mathbf{e}^i, \tag{8.15}$$

$$v_i = \mathbf{v} \cdot \mathbf{e}_i, \tag{8.16}$$

$$\mathbf{e}^i \cdot \mathbf{e}^j = g^{ij}, \tag{8.17}$$

$$\mathbf{e}^i \cdot \mathbf{e}_j = \delta^i_j. \tag{8.18}$$

A linear map $Q : U \to V$ between inner-product spaces is called *orthogonal* if $QQ^T = id_V$ and $Q^T Q = id_U$, where id stands for the identity map in the subscript space. The components of an orthogonal linear map in orthonormal bases of both spaces comprise an *orthogonal matrix*. A linear map T between inner-product spaces preserves the inner product if, and only if, it is an orthogonal map. By preservation of inner product we mean that: $T(\mathbf{u}) \cdot \mathbf{T}(\mathbf{v}) = \mathbf{u} \cdot \mathbf{v}, \ \ \forall \mathbf{u}, \mathbf{v} \in \mathbf{U}$.

3.8.2 Riemannian manifolds

If each tangent space $T_x\mathcal{M}$ of the manifold $\mathcal{M}$ is endowed with an inner product, and if this inner product depends smoothly on $x \in \mathcal{M}$, we say that $\mathcal{M}$ is a *Riemannian manifold*. To clarify the concept of smoothness, let $\{\mathcal{U}, \phi\}$ be a chart in $\mathcal{M}$ with coordinates $x^1, ..., x^n$. This chart induces the (smooth) basis field $\frac{\partial}{\partial x^1}, ..., \frac{\partial}{\partial x^n}$. We define the *contravariant components of the metric tensor* $\mathbf{g}$ associated with the given inner product (indicated by $\cdot$) as:

$$g_{ij} := \left(\frac{\partial}{\partial x^i}\right) \cdot \left(\frac{\partial}{\partial x^j}\right). \tag{8.19}$$

Smoothness means that these components are smooth functions of the coordinates within the patch. The *metric tensor* itself is given by:

$$\mathbf{g} = g_{ij} \, dx^i \otimes dx^j. \tag{8.20}$$

We have learned how an inner product defines an isomorphism between a vector space and its dual. When translated to Riemannian manifolds, this result means that the tangent and cotangent bundles are naturally isomorphic (via the point-wise isomorphisms of the tangent and cotangent spaces induced by the inner product).

A non-trivial physical example is found in Lagrangian Mechanics, where the kinetic energy (assumed to be a positive-definite quadratic form in the generalized velocities) is used to view the configuration space $\mathcal{Q}$ as a Riemannian manifold.

3.8.3 Riemannian connections

The theory of Riemannian manifolds is very rich in results. Classical differential geometry was almost exclusively devoted to their study and, more particularly, to the study of two-dimensional surfaces embedded in $\mathbb{R}^3$, where the Riemannian structure is derived from the Euclidean structure of the surrounding space.

A *Riemannian connection* is a linear connection on a Riemannian manifold. The most important basic result for Riemannian connections is contained in the following theorem:

Theorem 3.8.1 *On a Riemannian manifold there exists a unique linear connection (called the Levi-Civita connection) with vanishing torsion and such that the covariant derivative of the metric vanishes identically.*

We omit the proof. The curvature tensor associated with this special connection is called the *Riemann-Christoffel* curvature tensor. A Riemannian manifold is said to be *locally flat* if, for each point, a coordinate chart can be found such that the metric tensor components everywhere in the chart reduce to the identity matrix. It can be shown that local flatness is equivalent to the identical vanishing of the Riemann-Christoffel curvature tensor.

It can be shown that the Christoffel symbols of the Levi-Civita connection are given by:

$$\Gamma^k_{ij} = \frac{1}{2} g^{kh} \left(\frac{\partial g_{ih}}{\partial x^j} + \frac{\partial g_{jh}}{\partial x^i} - \frac{\partial g_{ij}}{\partial x^h} \right). \tag{8.21}$$

3.8.4 Geodesics in a Riemannian manifold

A geodesic in a Riemannian manifold satisfies Equation (7.34), namely:

$$\frac{d^2 x^i}{dt^2} \; + \; \Gamma^i_{jk} \, \frac{dx^j}{dt} \, \frac{dx^k}{dt} \; = \; 0, \tag{8.22}$$

where $\Gamma^i_{jk} = \Gamma^i_{kj}$ are the Christoffel symbols of the unique Levi-Civita connection associated with the metric $\mathbf{g}$. We note that, since the Levi-Civita connection is compatible with the metric, upon parallel transport the length of vectors is conserved. In particular, for a geodesic,

$$g_{ij} \frac{dx^i}{dt} \, \frac{dx^j}{dt} \; = \; \text{constant}. \tag{8.23}$$

This means that the parameter of a geodesic curve is necessarily proportional to its running length and, therefore, it can be identified with it.

Consider now the variational problem associated with the extreme values of the length functional:

$$L[\gamma] = \int_a^b \sqrt{g_{ij} \frac{dx^i}{dt} \, \frac{dx^j}{dt}}, \tag{8.24}$$

for a curve γ with equation $x^i = x^i(t)$. The Euler-Lagrange equation associated with this variational problem is precisely the geodesic equation (8.22), as can be verified directly. In other words, a geodesic in a Riemannian manifold is an extremal of the length functional.

Chapter 4

Physical illustrations

4.1 Mechanics in the configuration space

4.1.1 Virtual displacements and velocity vectors

Let us return to the example of the double pendulum considered in Section 2.1. Given a configuration $q \in \mathcal{Q}$, we consider a small perturbation to arrive at another, neighbouring, configuration, always moving over the surface of the torus $\mathcal{Q}$ (since the system cannot escape the trap of its own configuration space). Intuitively, what we have is a small piece of a curve in $\mathcal{Q}$, which we can identify with *tangent vector*.

To make this notion more precise, imagine that we have an initially un-stretched thin elastic ruler on which equally spaced markers have been drawn, including a zero mark. If we now stretch or contract this ruler, bend it and then apply it to the surface of the torus at some point q, in such a way that the zero mark falls on q, we obtain an intuitive representation of a *parametrized curve* γ on the configuration manifold. Let us now repeat this procedure ad infinitum with all possible amounts of bending and stretching, always applying the deformed ruler with its zero mark at the same point q. Among all the possible curves obtained in this way, there will be a sub-collection that shares the same tangent and the same stretch with γ. We call this whole collection (technically known as an *equivalence class* of parametrized curves) a *tangent vector to the configuration manifold* at q. Notice that, although when we draw this tangent vector $\mathbf{v}$ in the conventional way as an arrow, it seems to contradict the fact that we are supposed to stay on the surface, the definition as an equivalence class of curves (or, less precisely, a small piece of a curve) removes this apparent contradiction. Any of the curves in the equivalence (e.g., the curve γ of departure) can be used as the *representative* of the vector. The vector can also be regarded as a derivation with respect to the curve parameter (the equally spaced markers).

The collection of all tangent vectors at a point $q \in \mathcal{Q}$ is called the *tangent space of* $\mathcal{Q}$ *at* q and is denoted by $T_q\mathcal{Q}$. In the case of the torus, the interpretation of $T_q\mathcal{Q}$ is the tangent plane to the torus at q, as shown in Figure 4.1. The tangent space at a point q of the configuration space is the carrier of all the possible virtual displacements away from the configuration represented by q. A physically appealing way to look at virtual displacements is as *virtual velocities* multiplied by a small time increment.

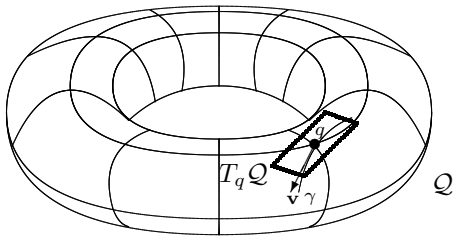

Figure 4.1: A tangent space

4.1.2 Force fields

A *force* α_q at a point $q \in \mathcal{Q}$ is a linear functional on the space of virtual displacements at q. This definition corresponds exactly to Lagrange's idea of a *generalized force* as an entity that produces *virtual work* on virtual displacements. Geometrically, a force is nothing but a covector at q, that is, $\alpha_q \in T_q^*\mathcal{Q}$. A *force field* α is, accordingly, a one-form on $\mathcal{Q}$, namely, a section of $T^*\mathcal{Q}$.

A force field α is said to be *conservative* or to *derive from a potential* if α is exact. That is, a force field is conservative if there exists a scalar function $V : \mathcal{Q} \to \mathbb{R}$ such that

$$\alpha = -dV. \tag{1.1}$$

4.1.3 The Lagrangian density

In Lagrange's view of Mechanics, a mechanical system is characterized by a single scalar function defined on the configuration space of the system. This function $\mathcal{L} : \mathcal{T}Q \to \mathbb{R}$ is called the *Lagrange density function* of the system. It encompasses both the force fields acting on the system and its inertia properties. In many cases of practical application, the Lagrangian density function can be decomposed into two parts, each one corresponding to one of the properties just mentioned. Notice that any function $V : \mathcal{Q} \to \mathbb{R}$ can be trivially extended to a function $V : T\mathcal{Q} \to \mathbb{R}$, which we denote by the same symbol. If we encompass the force fields in a single potential function V, the part associated with the inertia properties of the system, denoted by T, is known as the *kinetic energy density*. The Lagrangian density is defined as the difference:

$$\mathcal{L} = T - V. \tag{1.2}$$

For many mechanical systems of interest, the kinetic energy density $T : T\mathcal{Q} \to R$ is a point-wise positive-definite quadratic form in the velocities, namely (in a coordinate patch):

$$T = \frac{1}{2} m_{ij} \frac{dq^i}{dt} \frac{dq^j}{dt}, \tag{1.3}$$

where $m_{ij} = m_{ij}(q)$ are components of a positive-definite symmetric matrix, known as the *mass matrix* of the system. Clearly, in this frequent particular case, the configuration space acquires the structure of a Riemannian manifold whose metric is induced by the kinetic energy density. From now on, we will refer exclusively to these *decomposable systems*.

4.1.4 Lagrange's postulate and the equations of motion

The objective of Mechanics is to predict, out of given *initial conditions* $q_0 \in \mathcal{Q}$ and $\dot{q}_0 \in T_{q_0}$ at time $t = 0$, the system trajectory for times $t > 0$ within a certain time interval. We now formulate a fundamental postulate, which we call *Lagrange's postulate*, as follows:

Postulate 4.1.1 *For a decomposable mechanical system, in the absence of any force field, the system follows a geodesic of the Levi-Civita connection induced by the kinetic-energy density.*

Given any possible system trajectory, that is, given any curve γ in $\mathcal{Q}$ we define its *generalized acceleration field* $\mathbf{a}$ as the covariant derivative of its velocity in the direction of the velocity, namely:

$$\mathbf{a} = \nabla_{\mathbf{v}}\mathbf{v}. \tag{1.4}$$

The above postulate, therefore, can be rephrased as: In the absence of any force field, a decomposable system follows a trajectory of vanishing generalized acceleration.

Consider now the case of the presence of a non-vanishing force field α. According to the Newtonian viewpoint,[1] we postulate that the trajectory of the system is governed by the second-order ODE system:

$$\tau_{\mathbf{a}} = \alpha, \tag{1.5}$$

where we have denoted by $\tau_{\mathbf{a}}$ the covector corresponding to the vector $\mathbf{a}$ in the isomorphism between $T_q\mathcal{Q}$ and $T_q^*\mathcal{Q}$ induced by the Riemannian metric. In components:

$$(\tau_{\mathbf{a}})_i = m_{ij}a^j. \tag{1.6}$$

Box 4.1.2 Double power
It is worthwhile noticing that the Riemannian structure induced by the kinetic energy density needs to be exploited not just once but rather twice in the writing of the equations of motion. The first use of this structure is necessary for the very definition of the acceleration vector. This would have been feasible with only the specification of a linear connection (not necessarily Riemannian) on $\mathcal{Q}$. Even in the absence of a connection, it is always possible to lift any given trajectory γ to the tangent bundle $T\mathcal{Q}$[a] and, thereafter, to lift once again this lifted trajectory to the iterated bundle $T(T\mathcal{Q})$. Thus, a primitive notion of acceleration would give us an element of the iterated tangent bundle. A linear connection, on the other hand, implies a connection in the associated bundle $T\mathcal{Q}$, which permits us to calculate the vertical part of this primitive acceleration. This vertical part is tangent to a fibre of $T\mathcal{Q}$. But in a vector space there exists a natural isomorphism between the tangent spaces and the space itself. In this way, it can be said that the primitive acceleration, originally belonging to the iterated bundle, induces via a linear connection a vector tangent to $\mathcal{Q}$ itself. We have called this vector the acceleration. The second use of the Riemannian structure is essential to establish the equation of motion. Indeed, the force is a one-form whereas the acceleration is a vector. The existence of a metric permits us to relate one with the other.

[a] By attaching to each point of γ its tangent vector.

[1] For a mote thorough treatment of Newtonian Mechanics in the geometrical setting, see Segev R and Ailon A (1986), Newtonian Mechanics of Robots, it Journal of the Franklin Institute **322**/3, 173-183.

Recalling that a geodesic in a Riemannian manifold is also an extremal of the length functional, we may adopt the Lagrangian viewpoint and establish that the system follows a trajectory corresponding to an extreme value of the *Lagrangian functional* defined as:

$$L[\gamma] = \int_a^b \mathcal{L} dt. \tag{1.7}$$

In the case of conservative force fields, both formulations give identical results, as can be verified directly by writing the Euler-Lagrange equations of this variational problem.

4.2 Hamiltonian mechanics

In Lagrangian Mechanics, the fundamental geometric arena is precisely the tangent bundle $T\mathcal{Q}$. Indeed, the *Lagrangian density* $\mathcal{L}$ of a mechanical system is given by a function $\mathcal{L} : T\mathcal{Q} \to \mathbb{R}$, assigning to each configuration and each velocity (at this configuration) a real number.

A covector Q at q is a linear function that assigns to each tangent vector (virtual displacement δq) at q a real number $\delta W = \langle Q, \delta q \rangle$, whose meaning is the *virtual work* of the *generalized force* Q on the virtual displacement δq (or the *power* of the generalized force on the corresponding velocity). The terminology and the notation are due to Lagrange. The interesting feature of the geometric approach is that, once the basic geometric entity has been physically identified as a manifold, its tangent and cotangent bundles are automatically the carriers of physical meaning. In Hamiltonian Mechanics, covectors at $q \in \mathcal{Q}$ can be regarded as *generalized momenta* of the system. Thus, the cotangent bundle $T^*\mathcal{Q}$ is identified with the *phase space* of the system, namely, the repository of all configurations and momenta. The *Hamiltonian function* of a mechanical system is a function $\mathcal{H} : T^*\mathcal{Q} \to \mathbb{R}$. The cotangent bundle of a manifold is endowed with a canonical *symplectic structure*, which we now discuss briefly.

4.2.1 Symplectic vector spaces

A tensor T of type $(0, r)$ on V is a multilinear function acting on r vector arguments, $(\mathbf{v}_1, ..., \mathbf{v}_r)$. Fixing one argument, say $\mathbf{v}_1$, we obtain a tensor $T_{\mathbf{v}_1}$ of type $(0, r-1)$. In particular, a tensor T of type $(0, 2)$ assigns to each vector $\mathbf{u} \in V$ the covector $T_{\mathbf{u}}$ defined by:

$$T_{\mathbf{u}}(\mathbf{v}) = T(\mathbf{u}, \mathbf{v}) \quad \forall \mathbf{v} \in V. \tag{2.1}$$

The tensor T of type $(0, 2)$ is *non-degenerate* if $T_{\mathbf{u}} = \mathbf{0}$ implies that $\mathbf{u} = \mathbf{0}$. Since in a given basis the components of the covector $T_{\mathbf{u}}$ are $T_{ij}u^i$, we conclude that a necessary and sufficient condition for T to be nondegenerate is that the matrix with entries $[T_{ij}]$ must have a non-vanishing determinant, a condition that is independent of the basis chosen.

A *symplectic vector space* is a vector space in which a non-degenerate 2-covector ω has been singled out. The standard example is provided by a vector space of even dimension $2m$. Choosing a basis $\{\mathbf{e_1}, ..., \mathbf{e_m}, \mathbf{f}_1, ..., \mathbf{f_m}\}$, the 2-covector

$$\omega_{\mathbf{ef}} = \sum_{i=1}^{m} \mathbf{e}^i \wedge \mathbf{f}^i \tag{2.2}$$

is nondegenerate. It can be shown that every symplectic vector space is necessarily even-dimensional and that there exists a basis for which ω has the form (2.2).

An important property of a symplectic vector space is that, due to the nondegeneracy of the 2-covector ω, there exists a natural correspondence between vectors and covectors.

4.2.2 Symplectic manifolds

Recall that an r-form ω on a manifold $\mathcal{M}$ is a smooth r-covector field, namely, a smooth assignment of an r-covector ω_p at each point $p \in \mathcal{M}$. Equivalently, ω is a (smooth) section of the bundle $\Lambda_r(\mathcal{M})$. A *symplectic form* on $\mathcal{M}$ is a nondegenerate closed 2-form ω. A *symplectic manifold* $(\mathcal{M}, \omega)$ is a manifold in which a symplectic form ω has been singled out. According to our discussion above, a symplectic manifold is necessarily even-dimensional.

Given an m-dimensional manifold $\mathcal{Q}$ (for example, the configuration space of a mechanical system), the tangent and cotangent bundles are manifolds of even dimension $2m$. It is a remarkable fact that the cotangent bundle T^*Q of any manifold is automatically endowed with a *canonical symplectic form*. By 'canonical' we mean that this form is defined intrinsically (i.e., independently of any coordinate chart). It is not surprising, therefore, that this canonical structure results in a corresponding physical interpretation. For a mechanical system, the cotangent bundle represents the phase space (of positions and momenta) and the canonical form plays a fundamental role in Hamiltonian mechanics.

A generic point $s \in T^*\mathcal{Q}$ has the form $s = (q, p)$, where $q = \pi(s) \in \mathcal{Q}$ and $p \in T^*_q\mathcal{Q}$. Put differently, a point in the cotangent bundle consists of a point q in the base manifold and a 1-covector p at q. Let $\mathbf{V}$ be a tangent vector to $T^*\mathcal{Q}$ at the point $s = (q, p) \in T^*\mathcal{Q}$, namely, $\mathbf{V} \in T(T^*\mathcal{Q})$. Since the projection $\pi : T^*Q \to Q$ is a differentiable map, its differential $\pi_* : T(T^*\mathcal{Q}) \to T\mathcal{Q}$ is well-defined. In particular, $\pi_*(\mathbf{V}_s) \in T_q\mathcal{Q}$. But the tangent bundle $T(T^*\mathcal{Q})$, as a tangent bundle, has its own projection $\hat{\tau} : T(T^*\mathcal{Q}) \to T^*\mathcal{Q}$. In particular, $\hat{\tau}(\mathbf{V}_s) = s = (q, p)$. Since this is a covector at $q \in \mathcal{Q}$, it makes sense to evaluate it on the tangent vector $\pi_*(\mathbf{V}_s) \in T_q\mathcal{Q}$.

Recall that a 1-form on $T^*\mathcal{Q}$ is a smooth assignment of a covector θ_s at each point $s = (q, p) \in T^*\mathcal{Q}$. We define the *canonical 1-form* θ on $T^*\mathcal{Q}$ by the formula:

$$\theta(\mathbf{V}_s) = \langle \hat{\tau}(\mathbf{V}_s), \pi_*(\mathbf{V}_s) \rangle. \tag{2.3}$$

The *canonical symplectic form* ω on $T^*\mathcal{Q}$ is defined as:

$$\omega = -d\theta. \tag{2.4}$$

Thus, ω is exact and, therefore, closed. Moreover, it is nondegenerate. It is, in fact, not difficult to obtain a coordinate expression of the canonical symplectic form. We have seen that a chart $(q^1, ..., q^m)$ in $\mathcal{Q}$ induces a chart in $T^*\mathcal{Q}$. Indeed, any 1-form p on $\mathcal{Q}$ has the coordinate expression $p = p_i dq^i$, where the summation convention is in force. The induced chart in T^*Q uses as coordinates the $2m$ numbers $(q^1, ..., q^m, p_1, ..., p_m)$. The canonical 1-form θ is given by $\theta = p_i dq^i$. It follows that the canonical symplectic form is expressed as: $\omega = -dp_i \wedge dq^i = dq^i \wedge dp_i$.

4.2.3 Hamiltonian systems

A *Hamiltonian system* consists of a symplectic manifold $(\mathcal{M}, \omega)$ and a smooth real-valued function $\mathcal{H} : \mathcal{M} \to \mathbb{R}$ called the system *Hamiltonian*. In Classical Mechanics, the symplectic manifold is identified with the phase space $\mathcal{M} = T^*\mathcal{Q}$ of the underlying configuration manifold $\mathcal{Q}$.

A key concept in Hamiltonian systems is that of *Hamiltonian vector field*. Since the Hamiltonian $\mathcal{H}$ is differentiable, its differential $d\mathcal{H}$ is a well-defined 1-form on $\mathcal{M}$. In a symplectic manifold, on the other hand, to each 1-form we can assign uniquely a vector field, by exploiting the point-wise nondegeneracy of the symplectic form. We thus obtain the associated Hamiltonian vector field $\mathbf{V}_H$. More explicity, at each point $s \in \mathcal{M}$ we have:

$$\langle d\mathcal{H}, \mathbf{U} \rangle = \omega(\mathbf{V}_H, \mathbf{U}) \qquad \forall \mathbf{U} \in T_s\mathcal{M}. \tag{2.5}$$

A curve γ in $\mathcal{M}$ is a *trajectory* of the Hamiltonian system if it satisfies *Hamilton's equations*, namely, if it is an integral curve of the Hamiltonian vector field, viz.:

$$\frac{d\gamma}{dt} = \mathbf{V}_H(\gamma(t)). \tag{2.6}$$

In the natural coordinates of a cotangent bundle, the curve γ consists of the $2m$ functions $q^i = q^i(t)$ and $p_i = p_i(t)$, with $i = 1, ..., m$. The Hamiltonian vector field has the components $\partial H/\partial p_i$ and $-\partial H/\partial q^i$. We thus recover the standard form of Hamilton's equations:

$$\frac{dq^i}{dt} = \frac{\partial H}{\partial p_i}, \tag{2.7}$$

and

$$\frac{dp_i}{dt} = -\frac{\partial H}{\partial q^i}, \tag{2.8}$$

Notice that the construction (2.5) applies to any smooth real-valued function defined on $\mathcal{M}$, not just the Hamiltonian. Namely, to any such function $\mathcal{G}$ we can uniquely assign a vector field $\mathbf{V}_G$. We can thus define an operation between any two scalar fields $\mathcal{G}$ and $\mathcal{K}$, called the *Poisson bracket* $\{\mathcal{G}, \mathcal{K}\}$, by any of the equivalent prescriptions:

$$\{\mathcal{G}, \mathcal{K}\} = \mathbf{V}_K(\mathcal{G}) = \langle d\mathcal{G}, \mathbf{V}_K \rangle = \omega(\mathbf{V}_G, \mathbf{V}_K). \tag{2.9}$$

The derivative of a scalar function $\mathcal{G}$ along a trajectory γ of the Hamiltonian system $(\mathcal{M}, \mathcal{H})$ is obtained as:

$$\frac{d\mathcal{G}}{dt} = \frac{d\gamma}{dt}(\mathcal{G}) = \langle d\mathcal{G}, \frac{d\gamma}{dt} \rangle = \langle d\mathcal{G}, \mathbf{V}_H \rangle = \{\mathcal{G}, \mathcal{H}\}. \tag{2.10}$$

Thus, the Poisson bracket of a function $\mathcal{G}$ (representing some physical property of the system) with the Hamiltonian function describes the time evolution of $\mathcal{G}$. The vanishing of this Poisson bracket indicates, therefore, a conserved quantity.

4.3 Fluxes in Continuum Physics

One of the basic notions of Continuum Physics is that of an *extensive property*, a term that describes a property that may be assigned to *subsets* of a given universe, such as the mass of various parts of a material body, the electrical charge enclosed in a certain region of space, and so on. Mathematically speaking, therefore, an extensive property is expressed as a real-valued *set function* p, whose argument ranges over subsets $\mathcal{R}$ of a universe $\mathcal{U}$. It is usually assumed, on physical grounds, that the function p is additive, namely,

$$p(\mathcal{R}_1 \cup \mathcal{R}_2) = p(\mathcal{R}_1) + p(\mathcal{R}_2) \qquad \text{whenever } \mathcal{R}_1 \cap \mathcal{R}_2 = \emptyset. \tag{3.1}$$

With proper regularity assumptions, additivity means that, from the mathematical standpoint, p is a *measure* in $\mathcal{U}$.

In the appropriate space-time context, the *balance* of an extensive property expresses a relation between the rate of change of the property in a given region and the causes responsible for that change. Of particular importance is the idea of *flux* of the property through the boundary of a region, which is an expression of the rate of change of the property as a result of interaction with other regions. It is a common assumption that the flux between regions takes place through, and only through, common boundaries. In principle, the flux is a set function on the boundaries of regions. In most physical theories, however, this complicated dependence can be greatly simplified by means of the so-called *Cauchy postulates* and Cauchy's theorem.

4.3.1 Extensive-property densities

We will identify the universe $\mathcal{U}$ as an m-dimensional differentiable manifold. Under appropriate continuity assumptions, a set function such as the extensive property p is characterized by a *density*. Physically, this means that the property at hand cannot be concentrated on subsets of dimension lower than m. More specifically, we assume that the density ρ of the extensive property p is a smooth m-form on $\mathcal{U}$ such that

$$p(\mathcal{R}) = \int_{\mathcal{R}} \rho, \tag{3.2}$$

for any subset $\mathcal{R} \subset \mathcal{U}$ for which the integral is defined. Clearly, the additivity condition (3.1) is satisfied automatically.

We introduce the time variable t as if space-time were just a product manifold $\mathbb{R} \times \mathcal{U}$. In fact, this trivialization is observer-dependent, but it will serve for our present purposes. The density ρ of the extensive property p should, accordingly, be conceived as a function $\rho = \rho(t, x)$, where $x \in \mathcal{U}$. Notice that, since for fixed x and variable t, ρ belongs to the same vector space $\Lambda^m\left(T_x^*\mathcal{U}\right)$, it makes sense to take the partial derivative with respect to t to obtain the new m-form

$$\beta = \frac{\partial \rho}{\partial t}, \tag{3.3}$$

defined on $\mathcal{U}$. For a fixed (i.e., time-independent) region $\mathcal{R}$, we may write

$$\frac{dp(\mathcal{R})}{dt} = \int_{\mathcal{R}} \beta. \tag{3.4}$$

In other words, the integral of the m-form β over a fixed region measures the rate of change of the content of the property p inside that region.

4.3.2 Balance laws, flux densities and sources

In the classical setting of Continuum Mechanics it is assumed that the change of the content of a smooth extensive property p within a fixed region $\mathcal{R}$ can be attributed to just two causes: (1) the rate at which the property is produced (or destroyed) within $\mathcal{R}$ by the presence of sources and sinks, and (2) the rate at which the property enters or leaves $\mathcal{R}$ through its boundaries, namely the *flux* of p. For the sake of definiteness, in this section we adopt the convention that the production rate is positive for sources (rather than sinks) and that the flux is positive when there is a an outflow (rather than an inflow) of the property. The *balance equation* for the extensive property p states that *the rate of change of p in a fixed region $\mathcal{R}$ equals the difference between the production rate and the flux.* A good physical example is

the balance of internal energy in a rigid body due to volumetric heat sources and heat flux through the boundaries.

Since we have assumed continuity for p as a set function, we will do the same for both the production and the flux. As a result, we postulate the existence of an m-form s, called the *source density* such that the production rate in a region $\mathcal{R}$ is given by the integral

$$\int_{\mathcal{R}} s. \tag{3.5}$$

Just as ρ itself, the m-form s is defined over all of $\mathcal{U}$ and is independent of $\mathcal{R}$. Thus, from the physical point of view, we are assuming that the phenomena at hand can be described locally. This assumption excludes interesting phenomena, such as internal actions at a distance or surface-tension effects.

As far as the flux term is concerned, we also assume that it is a continuous function of subsets of the boundary ∂R. We postulate the existence, for each region $\mathcal{R}$, of a smooth $(m-1)$-form $\tau_{\mathcal{R}}$, called the *flux density*, such that the flux of p is given by

$$\int_{\partial\mathcal{R}} \tau_{\mathcal{R}}. \tag{3.6}$$

Thus, the classical balance law of the property p assumes the form

$$\int_{\mathcal{R}} \beta = \int_{\mathcal{R}} s - \int_{\partial\mathcal{R}} \tau_{\mathcal{R}}. \tag{3.7}$$

An equation of balance is said to be a *conservation law* if both s and $\tau_{\mathcal{R}}$ vanish identically.

4.3.3 Flux forms and Cauchy's formula

We note that (beyond the obvious fact that β and s are m-forms, whereas $\tau_{\mathcal{R}}$ is an $(m-1)$-form), there is an essential complication peculiar to the flux densities $\tau_{\mathcal{R}}$. Indeed, in order to specify the flux for the various regions of interest, it seems that one has to specify the form $\tau_{\mathcal{R}}$ for each and every region $\mathcal{R}$. In other words, while the rate of change of the property and the production term are specified by forms whose domain (for each time t) is the entire space $\mathcal{U}$, the flux term must be specified by means of a set function, whose domain is the collection of all regions. We refer to the set function $\mathcal{R} \mapsto \tau_{\mathcal{R}}$ as a *system of flux densities*. Consider, for example, a point $x \in \mathcal{U}$ belonging simultaneously to the boundaries of two different regions. Clearly, we do not expect that the flux density will be the same for both. The example of sun-tanning should be sufficiently convincing in this regard. Consider, however, the following particular case. Let the natural *inclusion* map

$$\iota : \partial\mathcal{R} \longrightarrow \mathcal{U}, \tag{3.8}$$

be defined by

$$\iota(x) = x \quad \forall x \in \partial\mathcal{R}. \tag{3.9}$$

Notice that this formula makes sense, since $\partial\mathcal{R} \subset \mathcal{U}$. Moreover, the map ι is smooth. It can, therefore, be used to pull back forms of any order on $\mathcal{U}$ to forms of the same order on $\partial\mathcal{R}$. In particular, we can define

$$\int_{\partial\mathcal{R}} \phi = \int_{\partial\mathcal{R}} \iota^*(\phi), \tag{3.10}$$

for any form ϕ on $\mathcal{U}$. Let us now assume the existence of a globally defined $(m-1)$-*flux form* Φ on $\mathcal{U}$ and let us define the associated system of flux densities by means of the formula

$$\tau_{\mathcal{R}} = \iota^*_{\partial\mathcal{R}}(\Phi), \tag{3.11}$$

where we use the subscript $\partial\mathcal{R}$ to emphasize the fact that each region requires its own inclusion map. Equation (3.11) is known as *Cauchy's formula*. Clearly, this is a very special system of flux densities (just as a conservative force field is a special vector field derivable from a single scalar field). Nevertheless, it is one of the fundamental results of classical Continuum Mechanics that, under rather general assumptions (known as *Cauchy's postulates*), every system of flux densities can be shown to derive from a unique flux form using Cauchy's formula (3.11). We will omit the general proof of this fact, known as Cauchy's theorem.

In less technical terms, Cauchy's formula is the direct result of assuming that the flux is given by a *single* 2-form defined over the three-dimensional domain of the body. The fact that one and the same form is to be used for a given location, and integrated over the given boundary, is trivially seen to imply (and generalize) the linear dependence of the flux on the normal to the boundary, as described in the standard treatments.

4.3.4 Differential expression of the balance law

Assuming the existence of a flux form Φ, the general balance law (3.7) can be written as

$$\int_{\mathcal{R}} \beta = \int_{\mathcal{R}} s - \int_{\partial\mathcal{R}} \iota^*_{\partial\mathcal{R}}(\Phi). \tag{3.12}$$

Using Stokes' theorem (Equation (5.11)), we can rewrite the last term as

$$\int_{\partial\mathcal{R}} \iota^*_{\partial\mathcal{R}}(\Phi) = \int_{\mathcal{R}} d\Phi, \tag{3.13}$$

where the dependence on $\partial\mathcal{R}$ has evaporated. Using this result, we write (3.12) as

$$\int_{\mathcal{R}} \beta = \int_{\mathcal{R}} s - \int_{\mathcal{R}} d\Phi. \tag{3.14}$$

Since this balance law should be valid for arbitrary $\mathcal{R}$, and since the forms β, s and Φ are defined globally and independently of the region of integration, we obtain

$$\beta = s - d\Phi. \tag{3.15}$$

This equation is known as the *differential expression of the general balance law*.

4.4 Microstructure

The idea of endowing bodies with a microstructure represented by affine deformations of micromedia, or grains, embedded in a matrix goes back to the pioneering work of the Cosserat brothers. If each of the grains is permitted to undergo just affine deformations (namely, deformations with a constant gradient),

it is clear that the extra kinematics can be described in terms of a linear mapping of any basis attached to the grain. Knowing how one basis deforms is enough to determine how all other bases at the same point deform. The choice of basis remaining arbitrary, we are naturally led to the conclusion that the appropriate geometric counterpart of a body $\mathcal{B}$ with linear microstructure (a *general Cosserat body*[2]) is the bundle of linear frames $F\mathcal{B}$.

4.5 Kinematics of a Cosserat body

A *Cosserat body* is, by definition, the frame bundle $F\mathcal{B}$ of an ordinary body $\mathcal{B}$, usually called the *macromedium*, the *matrix* or the *underlying body*. We recall that the frame bundle of a differentiable manifold consists roughly of the collection of all frames at each point of the manifold. Thus, in the case of a three-dimensional manifold $\mathcal{B}$, we obtain the collection of all triads of linearly independent vectors forming all possible bases of the tangent spaces of $\mathcal{B}$. Since each triad is attached to a particular point of $\mathcal{B}$, we have a *projection* map:

$$\pi : F\mathcal{B} \longrightarrow \mathcal{B}, \tag{5.1}$$

which assigns to each triad the point at which it is attached. In the terminology of differential geometry, the macromedium $\mathcal{B}$ is known as the *base manifold* of the frame bundle. Given a point $\mathbf{X} \in \mathcal{B}$, the inverse image $\pi^{-1}(\mathbf{X})$ is called the *fibre* at $\mathbf{X}$. It consists of all the possible bases of the tangent space $T_X\mathcal{B}$. In the physical picture, the fibre is the carrier of the information about the events taking place at the"grain" level. Since, as conceived by the Cosserats, any particular basis (rather that the whole collection thereof) should carry that very information, we will take this fact into consideration when defining the concept of configuration and deformation of a Cosserat body.

Assume now that a coordinate chart with coordinates X^I $(I = 1, 2, 3)$ is specified on an open set $\mathcal{U}$ of the base manifold $\mathcal{B}$. The natural basis of this chart:

$$\mathbf{E}_I = \frac{\partial}{\partial X^I}, \quad I = 1, 2, 3, \tag{5.2}$$

determines, at each point of $\mathcal{U} \subset \mathcal{B}$, a basis of the tangent space. In other words, the coordinate chart induces a smooth local section of the frame bundle. Any frame $\mathbf{H}_I$ $(I = 1, 2, 3)$ within the domain $\mathcal{U}$ can be expressed in terms of components in the coordinate-induced frame by means of a matrix, viz.:

$$\mathbf{H}_I = H_I^J \mathbf{E}_J. \tag{5.3}$$

We can say, therefore, that, as far as the domain of the chart is concerned, every element of $F\mathcal{B}$ can be represented uniquely by the twelve numbers (X^I, H_J^K). In fact one can prove that the frame bundle $F\mathcal{B}$ is itself a differentiable manifold of dimension 12 and that the numbers just described constitute admissible coordinates of this manifold. If we should consider a different coordinate system, Y^I say, on an open set $\mathcal{V} \subset \mathcal{B}$ such that $\mathcal{U} \cap \mathcal{V} \neq \emptyset$, the natural bases of both systems can be related point-wise by an arbitrary non-singular 3×3-matrix, that is, by an arbitrary member of the general linear group $GL(3, \mathbb{R})$. This means that fibre-wise the coordinate transformations are governed by this group, which is, therefore, called the *structural group* of the bundle $F\mathcal{B}$. On the other hand, for a fixed basis at a point, all the elements in the fibre, according to Equation (5.3), are precisely spanned by the collection

[2]The terminology *Cosserat medium* is often used in the literature to designate the particular case in which the grains can undergo rotations only. For this reason, we use here the longer and more descriptive title. An alternative terminology distinguishes between micropolar and micromorphic media.

of non-singular 3×3-matrices. This special situation, whereby the nature of the fibres and the nature of the structural group are identical, is described in differential geometry by saying that the bundle of frames of a manifold is a *principal bundle*.

We have demonstrated the manifold character of the frame bundle by means of fibre-wise coordinates which consist of components of the frames in terms of the natural basis of a coordinate system of the base manifold. It should be clear, however, that we could as well have singled out at each point any basis of the tangent space of the base manifold, not necessarily a coordinate basis, and expressed the fibre-wise coordinates in terms of the matrix of components of the frames in that particular basis.

In a principal bundle we have at our disposal a special operation called the *right action* of the structural group on the principal bundle. We will describe this operation for the particular case at hand. Let $\mathbf{M}$ belong to our structural group. $\mathbf{M}$ is, therefore, a nonsingular matrix with entries $\{M^I{}_J\}$. We want to define the right action R_M of $\mathbf{M}$ as it applies to each element of the principal bundle $F\mathcal{B}$ to produce another element of $\mathcal{B}$. We will do this as follows: let $(X^I, H^I{}_J)$ be the components of an element of $F\mathcal{B}$ in some coordinate system. Then the image of this element by the right action of $\mathbf{M}$ is given, by definition, as the element of $F\mathcal{B}$ with components $(X^I, H^I{}_J M^J{}_K)$ in the same coordinate system. It is not difficult to prove that this definition, although expressed in a particular chart, is in fact independent of the chart chosen. Note that a frame at a point is always mapped to another frame at the same point, so that the right action just defined is fibre preserving.

We now seek an appropriate definition of a configuration of a Cosserat body. To this end, we start by noting that the physical space (which we have identified with $\mathbb{R}^3$) is itself a differentiable manifold and, therefore, it has a naturally defined frame bundle $F\mathbb{R}^3$ with projection π_R. We want to define a configuration of a Cosserat body as a map K between these two principal bundles, namely:

$$K : F\mathcal{B} \longrightarrow F\mathbb{R}^3. \tag{5.4}$$

But it is clear that an arbitrary map will not do, so this concept needs further clarification. When we map a principal bundle into another, there are three elements at play. Firstly, there are the two base manifolds, which in our case are $\mathcal{B}$ and $\mathbb{R}^3$. Secondly, there are the fibres at each point of these manifolds. And finally, there are the two structural groups. We will assume that the configuration K incorporates an ordinary configuration of the base manifold (the macromedium) $\mathcal{B}$, that is, an embedding:

$$\kappa : \mathcal{B} \longrightarrow \mathbb{R}^3. \tag{5.5}$$

This map is, as we know, smooth and has a smooth inverse defined on the image $\kappa(\mathcal{B})$. Secondly, we want that fibres don't get mixed up: a frame at a point of $\mathbf{X} \in \mathcal{B}$ must be mapped to a frame at the image point $\kappa(\mathbf{X})$. In the physical picture, we want each point in the matrix to carry its own "grain" in the process of deformation. Mathematically, this means that the map K must satisfy the equation:

$$\pi_R \circ K = \kappa \circ \pi. \tag{5.6}$$

This restriction is nicely represented in the following commutative diagram:

$$\begin{array}{ccc} F\mathcal{B} & \xrightarrow{\;K\;} & F\mathbb{R}^3 \\ \pi \big\downarrow & & \big\downarrow \pi_R \\ \mathcal{B} & \xrightarrow{\;\kappa\;} & \mathbb{R}^3 \end{array} \tag{5.7}$$

But we are not done yet, and this is because in a principal bundle we also have to take into consideration the structural groups and provide an appropriate map between them. Since in our particular case the two structural groups are identical, namely $GL(3,\mathbb{R})$, we will agree that the map between them is just the identity map. Finally, we will require that the right action of the structural group commute with the map between fibres. This can be represented by the following commutative diagram:

$$\begin{array}{ccc} F\mathcal{B} & \xrightarrow{K} & F\mathbb{R}^3 \\ {\scriptstyle R_M}\downarrow & & \downarrow{\scriptstyle R_M} \\ F\mathcal{B} & \xrightarrow{K} & F\mathbb{R}^3 \end{array} \tag{5.8}$$

Physically, this means that the deformation of a grain is an intrinsic quantity independent of the particular triad that one chooses to represent that grain. This is precisely the consistency condition that reconciles the original Cosserat picture (one frame representing the grain at a point) with the principal bundle picture (the collection of all frames at a point representing the same grain).

In the terminology of differential geometry, with all the above restrictions, the map between $F\mathcal{B}$ and its image $K(F\mathcal{B}) \subset F\mathbb{R}^3$ is called a *principal-bundle isomorphism*. In terms of components in a given coordinate system in the body and in space, a configuration of a Cosserat body is defined by twelve smooth functions:

$$x^i = \kappa^i(X^J), \tag{5.9}$$

and

$$K^i{}_I = K^i{}_I(X^J). \tag{5.10}$$

We see that in a Cosserat body there exist two independent mechanisms, as it were, of dragging vectors by means of a deformation (Figure 4.2): The first mechanism is the ordinary dragging of vectors by means of the deformation gradient of the macromedium, represented by the matrix with entries $F^i_I = x^i_{,I}$. The second mechanism is the one associated with the deformation of the "micromedium" or grain, and is represented by the matrix with entries $K^i{}_I$. To recover an ordinary medium (without microstructure) these two mechanisms are identified with each other.

We have been freely talking about configurations and deformations almost interchangeably. Indeed, since principal-bundle morphisms can be inverted and composed, we can clearly adopt a reference configuration of a Cosserat body and define the notion of a deformation in the same way as we have done for ordinary bodies.

4.6 Dislocations

4.6.1 An intuitive picture

Let an atomic lattice be given by, say, all points with integer coordinates in $\mathbb{R}^2$. To each atom we can associate two vectors (in this instance unit and orthogonal) determined by joining it to its immediate neighbours to the right and above, respectively. If the lattice is deformed regularly, these vectors will deform accordingly, changing in length and angle, but always remaining linearly independent at each

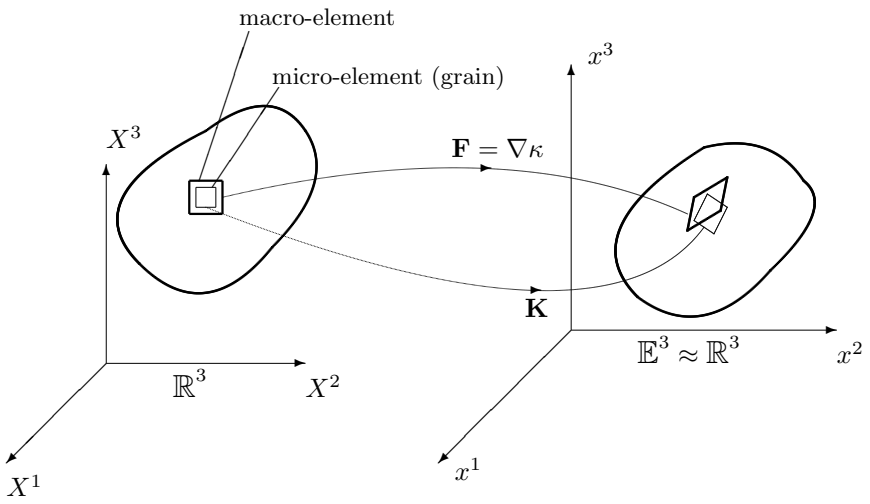

Figure 4.2: The two kinematically independent dragging mechanisms

atom. In the (not precisely defined) continuous limit, we can imagine that each point of $\mathbb{R}^2$ has been endowed with a basis or frame, the collection of which is called a *moving frame* (or *repère mobile*)[3].

Returning to the discrete picture, if there is a dislocation (for example, a half-line of atoms is missing, as shown on the right-hand side of Figure 4.3), the local bases will be altered differently from the case of a mere deformation. The engineering way to recognize this is the so-called *Burgers' circuit*, which consists of a four-sided path made of the same number of atomic spacings in each direction. The failure of such a path to close is interpreted as the presence of a local dislocation in the lattice. We want to show that in the putative continuous limit this failure is represented by the non-vanishing of a Lie bracket. What we have in the continuous case as the only remnant of the discrete picture is a smoothly distributed collection of bases, which we have called a moving frame, and which can be seen as two vector fields $\mathbf{E}_\alpha$ ($\alpha = 1, 2$) over $\mathbb{R}^2$.

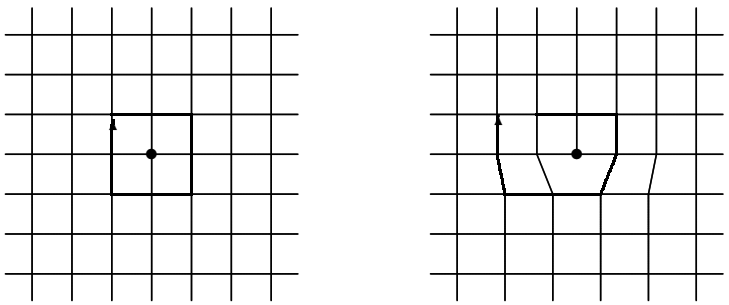

Figure 4.3: Dislocation in a crystal lattice

From the theory of ordinary differential equations, we know that each vector field gives rise, at least locally, to a well-defined family of parametrized integral curves, where the parameter is determined up

[3]This idea was introduced mathematically by Cartan and, in a physical context, by the brothers Cosserat.

to an additive constant. More specifically, these curves are obtained as the solutions $\mathbf{r} = \mathbf{r}(s^\alpha)$ of the systems of equations:

$$\frac{d\mathbf{r}(s^\alpha)}{ds^\alpha} = \mathbf{E}_\alpha[\mathbf{r}(s^\alpha)], \qquad (\alpha = 1, 2; \quad \text{no sum on } \alpha), \tag{6.1}$$

where $\mathbf{r}$ represents the natural position vector in $\mathbb{R}^2$. The parameter s^α (one for each of the two families of curves) can be pinned down in the following way. Select a point p_0 as origin and draw the (unique) integral curve γ_1 of the first family passing through this origin. Adopting the value $s^1 = 0$ for the parameter at the origin, the value of s^1 becomes uniquely defined for all the remaining points of the curve. Each of the curves of the second family must intersect this curve of the first family. We adopt, therefore, for each of the curves of the second family the value $s^2 = 0$ at the corresponding point of intersection with that reference curve (of the first family). In this way we obtain (at least locally) a new coordinate system s^1, s^2 in $\mathbb{R}^2$. By construction, the second natural base vector of this coordinate system is $\mathbf{E}_2$. But there is no guarantee that the first natural base vector will coincide with $\mathbf{E}_1$, except at the curve γ_1 through the adopted origin. In fact, if we repeat the previous construction in reverse, i.e. with the same origin but adopting the curve γ_2 of the *second* family as a reference, we obtain in general a different system of coordinates, which is well adapted to the basis vectors $\mathbf{E}_1$, but not necessarily to $\mathbf{E}_2$.

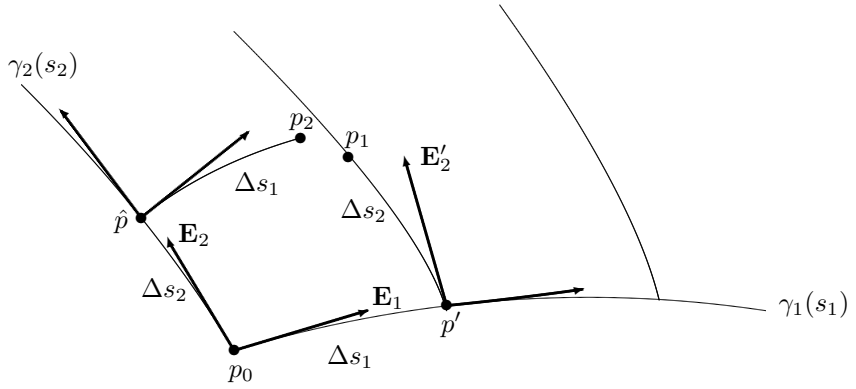

Figure 4.4: The continuous case

Assume now that, starting at the adopted origin, we move an amount of Δs^1 along γ_1 to arrive at a point p' and thereafter we climb an amount of Δs^2 along the encountered curve of the second family through p'. We arrive at some point p_1. Incidentally, this is the point with coordinates $(\Delta s^1, \Delta s^2)$ in the coordinate system obtained by the first construction. If, however, starting at the same origin we move by Δs^2 along the curve γ_2 to a point $\hat{p}$ and then move by Δs^1 along the encountered curve of the first family, we will arrive at a point p_2 (whose coordinates are $(\Delta s^1, \Delta s^2)$ in the *second* construction) which is, in general, different from p_1. Thus, we have detected the failure of a four-sided circuit to close! The discrete picture has, therefore, its continuous counterpart in the non-commutativity of the flows along the two families of curves.

Let us calculate a first-order approximation to the difference between p_2 and p_1. For this purpose, let us evaluate, to the first order, the base vector $\mathbf{E}_2$ at the auxiliary point p'. The result is:

$$\mathbf{E}_2' = \mathbf{E}_2(p_0) + \frac{\partial \mathbf{E}_2}{\partial x^i}\frac{dx^i}{ds^1}\,\Delta s^1, \tag{6.2}$$

where derivatives are calculated at p_0. The position vector of p_1, always to first-order approximation, is

obtained, therefore, as:

$$\mathbf{r}_1 = \Delta s^1 \mathbf{E}_1(p_0) + \Delta s^2 \left(\mathbf{E}_2(p_0) + \frac{\partial \mathbf{E}_2}{\partial x^i} \frac{dx^i}{ds^1} \, \Delta s^1, \right). \tag{6.3}$$

In a completely analogous fashion, we calculate the position vector of p_2 as:

$$\mathbf{r}_2 = \Delta s^2 \mathbf{E}_2(p_0) + \Delta s^1 \left(\mathbf{E}_1(p_0) + \frac{\partial \mathbf{E}_1}{\partial x^i} \frac{dx^i}{ds^2} \, \Delta s^2, \right). \tag{6.4}$$

By virtue of (6.1), however, we have:

$$\frac{dx^i}{ds^\alpha} = E^i_\alpha, \tag{6.5}$$

where E^i_α is the i-th component in the natural basis of $\mathbb{R}^2$ of the base vector $\mathbf{E}_\alpha$. From the previous three equations we obtain:

$$\mathbf{r}_2 - \mathbf{r}_1 = \left(\frac{\partial \mathbf{E}_1}{\partial x^i} E^i_2 \; - \frac{\partial \mathbf{E}_2}{\partial x^i} E^i_1 \right) \Delta s^1 \Delta s^2 = [\mathbf{E}_1, \mathbf{E}_2] \, \Delta s^1 \Delta s^2. \tag{6.6}$$

We thus confirm that the closure of the infinitesimal circuits generated by two vectors fields is tantamount to the vanishing of their Lie bracket. This vanishing, in turn, is equivalent to the commutativity of the flows generated by these vector fields. For this reason, the Lie bracket is also called the *commutator* of the two vector fields. In physical terms, we may say that the vanishing of the Lie brackets between the vector fields representing the limit of a lattice is an indication of the absence of dislocations.

Since in this example we have introduced the notion of a moving frame, that is, a smooth field of bases $\mathbf{E}_i$ $(i = 1, ..., n)$ over an n-dimensional manifold, it makes sense to compute all the possible Lie brackets between the base vectors and to express them in terms of components in the local basis. Since a Lie bracket of two vector fields is itself a vector field, there must exist unique scalar fields c^k_{ij} such that:

$$[\mathbf{E}_i, \mathbf{E}_j] = c^k_{ij} \mathbf{E}_k \qquad (i, j, k = 1, ..., n). \tag{6.7}$$

These scalars are known as the *structure constants* of the moving frame. The structure constants vanish identically if, and only if, the frames can be seen locally as the natural base vectors of a coordinate system.

4.6.2 Distant parallelism

An equivalent way to analyze the presence of dislocations within the previous picture consists of regarding the moving frame associated with an underlying (perfect or defective) lattice as establishing a *distant parallelism* in the body manifold $\mathcal{B}$, a notion that we have already encountered. Two tangent vectors at different points are said to be parallel if they have the same respective components in the respective local bases. Clearly, a distant parallelism establishes a fixed isomorphism between any pair of fibres of the frame bundle $F\mathcal{B}$ and, consequently, a linear connection.

To determine the Christoffel symbols of this parallelism in some coordinate system x^i, we start by writing the moving frame in terms of components as:

$$\mathbf{E}_\alpha = E^i_\alpha \, \frac{\partial}{\partial x^i}. \tag{6.8}$$

According to our criterion for parallelism, it is obvious that these base vectors (having trivially constant components upon themselves) are parallel (along any curve). Accordingly, their covariant derivatives must vanish identically along the coordinate lines. Using Equation (7.32), we obtain:

$$\nabla_{\frac{\partial}{\partial \mathbf{x}^h}} \mathbf{E}_\alpha = \left(\frac{dE^i_\alpha}{dx^h} \; + \; \Gamma^i_{jk} \; \delta^j_h \; E^k_\alpha \right) \; \frac{\partial}{\partial x^i} = 0, \tag{6.9}$$

for each coordinate line x^h. Consequently,

$$\Gamma^i_{jk} = -\frac{dE^i_\alpha}{dx^j} \; E^{-\alpha}_{\;\;k}. \tag{6.10}$$

The corresponding Riemann-Christoffel curvature tensor vanishes identically. The components of the torsion tensor are proportional to the components of the Lie brackets of corresponding pairs of vectors of the frames. Thus, the vanishing of the torsion is indicative of the absence of dislocations.

More generally, material parallelisms can be introduced in a uniform body via material isomorphisms, as described in Section 2.5. The corresponding *material connections*, whether or not curvature-free, open the door for the detection of material inhomogeneity. Locally, homogeneity is tantamount to the existence of a torsion-free material connection.[4]

4.6.3 Bravais planes and differential forms

From the previous treatment it might appear that the theory of continuous distributions of dislocations is irretrievably tied to the specification of a basis (or a collection of bases) at each point of the continuum. It comes, therefore, as a surprise that in fact a single differential one-form is sufficient to display the physical idea of a dislocation. A covector ω on an n-dimensional vector space $\mathbf{V}$ induces a family of hyperplanes. A hyperplane can be regarded either as a subspace of dimension $n-1$ in the affine space associated with $\mathbf{V}$ or, equivalently, as the collection of vectors of $\mathbf{v} \in \mathbf{V}$ such that the evaluation $\omega(\mathbf{v})$ is equal to a fixed constant k. If, for example, we restrict k to be an integer, then the evaluation of the covector ω on a vector $\mathbf{v}$ can be roughly regarded as the "number of hyperplanes" cut by the arrow representing the vector. Thus, a covector can be seen as a system of parallel planes, with a given "density", of a Bravais lattice.

In the case of a manifold $\mathcal{B}$, a one form ω is a covector field or, more precisely, a cross section of the cotangent bundle $T^*\mathcal{B}$. According to our previous description, therefore, a one-form represents at each point a density of parallel layers. We note that no additional structure, metric or otherwise, is required for this interpretation. We call the one-form ω the *local layering form*. Assume now that the local layering form ω is closed, namely, $d\omega = 0$ identically at the nieghbourhood of a point. Then, a (perhaps smaller) neighbourhood exists where the form is also exact. In other owrds, there exists locally a scalar function $f : \mathcal{B} \to \mathbb{R}$ where $\omega = df$. This means that, at least locally, the hyperplanes constituting the pointwise layering are, in fact, tangent to the hypersurfaces $f =$ constant and also that the local densities at neighbouring points are mutually compatible. In this case we say that the given layering form determines a *locally coherent system*. If, on the other hand, ω is not closed, we define its exterior derivative $\delta = d\omega$ as the *dislocation density 2-form*. A non-vanishing δ can indeed be interpreted as a local incoherence in the single family of Bravais planes.

[4]See Noll W (1967), Materially uniform bodies with inhomogeneities, *Archive for Rational Mechanics and Analysis* **27**, 1-32; Wang C-C (1967), On the geometric structure of simple bodies, *Archive for Rational Mechanics and Analysis* **27**, 33-94; Bloom F (1979), *Modern differential geometric techniques in the theory of continuous distributions of dislocations*, Springer; Epstein M and Elżanowski M (2007), *Material inhomogeneities and their evolution*, Springer.

4.6.4 Singular dislocations and de Rham currents

If a collection of n independent one-forms is stipulated on a manifold $\mathcal{B}$, it is possible to reproduce all the results of the continuous theory of dislocations (including the role of the torsion tensor of a distant parallelism). It would appear, therefore, that, apart from the fact that a dislocation can be associated to a single family of Bravais planes, not much has been gained beyond the elegance of the dual formulation. On the other hand, the power of this dual formulation resides in the ability of the theory of differential forms to sustain a generalization that can encompass singular dislocations (as irregular as point defects) as well as the smooth case. Such a generalization is provided by the concept of *currents* introduced by de Rham in a now classic work.[5] De Rham's work is a generalization to differential forms of the notion of *distribution* in analysis. A distribution is a linear functional in the space of C^∞ functions with compact support in $\mathbb{R}^n$. It comprises such well-known entities as the Dirac delta 'function' and its 'derivatives'. In an imprecise way it can be said that a current is a differential form whose coefficients in a coordinate system are distributions. A unified treatment of dislocations from this point of view can be found in a recent paper,[6]

Acknowledgment This work has been partially supported by the Natural Sciences and Engineering Research Council of Canada.

[5]de Rham G (1984), *Differentiable manifolds*, Springer. This is a translation of the French original *Varietés différentiables* published in 1955 by Hermann, Paris.

[6]Epstein M and Segev R (2012), Geometric aspects of singular dislocations, *Mathematics and Mechanics of Solids*.