Generalised Connections and Applications to Control Theory

Bavo Langerock

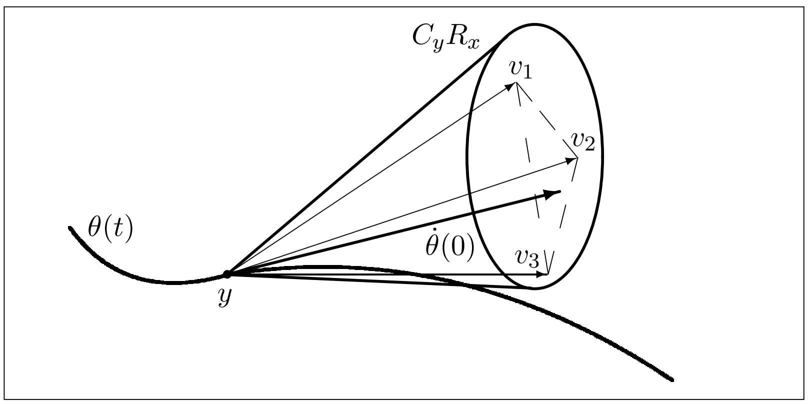

Proefschrift ingediend aan de Faculteit Wetenschappen
tot het behalen van de graad van
Doctor in de Wetenschappen: Wiskunde

Promotor: Prof. Dr. F. Cantrijn
Co-promotor: Prof. Dr. W. Sarlet

Universiteit Gent
Faculteit Wetenschappen
Vakgroep Wiskundige Natuurkunde en Sterrenkunde
Academiejaar 2002-2003

Voorwoord

Met genoegen wil ik dit voorwoord gebruiken om enkele mensen te bedanken. Ik denk aan mijn collega's en de gezamenlijke 'soeppauzes', die zeker zorgden voor een aangename onderbreking bij het eerder eenzame werk achter mijn bureau. Mijn ouders, broers en vrienden wil ik bedanken. Zij leverden geen wiskundige steun aan dit doctoraat, maar gaven die wel op ontelbare andere manieren. Misschien zijn ze zich hier niet van bewust, maar dit maakt het zeker niet minder waard. Antje-PK, mijn lief, zal zich hier waarschijnlijk in terugvinden. De vele toffe momenten die we samen doorbrachten, gaven me een hoop energie om alles tot een goed einde te brengen. Je bent een schat.

Frans Cantrijn, mijn promotor, wens ik zeker te bedanken. Meer nog, hij verdient zelfs een standbeeld. Hoeveel keer ben ik zijn bureau binnengestormd om vlug nog iets te vragen ? Wat was nu ook alweer een pre-symplectische 2-vorm ? Weet jij in welk boek ik iets kan vinden over . . . ? Hij heeft misschien door mijn toedoen enkele grijze haren gekregen, maar dat was geen belemmering om toch iedere keer bereidwillig, en zelfs met een noodzakelijk grapje tussendoor, mijn vragen te beantwoorden. Frans, bedankt.

Naast mijn promotor, heeft ook Jorge Cortés deze thesis nagelezen om allerlei fouten op te sporen. Ik ben ervan overtuigd, aan de hoeveelheid verbeteringen die in de marges te zien waren, dat er weinig speurwerk nodig was maar eerder veel doorzettingsvermogen. Zowel Jorge, Alberto Ibort, Manuel de León, Tom Mestdag als Willy Sarlet ben ik erkentelijk voor de vele boeiende gesprekken en hun interesse in dit werk. In het bijzonder vermeld ik Manuel, voor de vriendelijke ontvangst tijdens mijn verblijf in Madrid.

Gent, 29 april 2003 Bavo Langerock

Contents

Introduction

We start with an elementary classical problem of variational calculus, the so-called isoperimetric problem (see [46]). A variant of it is encountered in the legend about the founding of the city of Carthage, about 2850 years ago, and is described in the Aeneid of Vergil. This example is rewarding in that it touches the various topics encountered in this thesis.

Queen Dido had to flee from her brother Pygmalion who had already killed her husband. She landed, accompanied by some servants, at the African shore, where King Jarbas ruled. Dido begged Jarbas for help, and she persuaded him to give her as much land as she could enclose with the hide of a bull. Dido summoned her servants to cut the bull's hide into a single narrow strip. In order to obtain as much land as possible, Dido was confronted with the following geometric problem: *given a strip with endpoints on a fixed line L (the coastline) and with fixed length; what shape should the strip have in order that, together with the line L, it encloses a piece of land with the largest possible area*. The following picture sketches a possible situation.

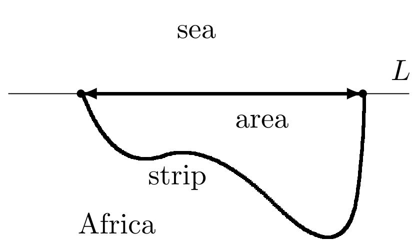

We shall study here the isoperimetric problem, formulated as follows. *Given a rectifiable strip with endpoints on a fixed line L and such that the domain enclosed by the strip and the line has a prescribed area S; what shape should the strip have in order that its length be minimal.* J. Steiner proved in 1841 that the solution to this problem is a half-circle. The isoperimetric problem is also called the dual to Dido's problem because one can prove that the solution of the isoperimetric problem is also the solution of Dido's problem (and vice versa).

When having a closer look at the formulation of the dual problem, one can see that the only shapes of the strip that are admissible, are those for which the enclosed domain has a given fixed area S. This condition is generally called a *constraint* and is an essential ingredient of the kind of variational problems that are investigated in this thesis. In order to formulate a more detailed characterisation of the constraint we are dealing with, we choose coordinates (x, y) in the plane of the strip such that we can describe a shape of the strip by a curve $c : [a, b] \to \mathbb{R}^2 : t \mapsto c(t) = (x(t), y(t))$. The initial point $c(a) = (x_a, y_a)$ and the endpoint $c(b) = (x_b, y_b)$ of the curve both coincide with the line L. We assume that the coordinate axes are chosen such that the origin $(0, 0)$ coincides with $c(a)$ and such that L coincides with the y-axis, i.e. $x_a = y_a = x_b = 0$. If Ω denotes the domain enclosed by c and L, then the area S of Ω equals:

$$S = \iint_\Omega dxdy.$$

By Stokes' theorem we know that the area S of Ω equals the line integral of the smooth vector field $-y\partial/\partial x + x\partial/\partial y$ along the oriented boundary of Ω:

$$S = \iint_\Omega dxdy = \frac{1}{2}\int_c xdy - ydx = \frac{1}{2}\int_a^b \big(x(t)\dot{y}(t) - y(t)\dot{x}(t)\big)dt;$$

The line segment on L between the initial and endpoint does not contribute to the line integral. If we define the following function of t:

$$z(t) = \frac{1}{2}\int_a^t \big(x(t')\dot{y}(t') - y(t')\dot{x}(t')\big)dt',$$

then, obviously, we have $z(b) = S$. This property encourages us to consider the curve $\hat{c}(t) = (x(t), y(t), z(t))$ in $\mathbb{R}^3$. The components of the velocity curve $\dot{\hat{c}}(t)$ satisfy the following differential equation:

$$\dot{z} = \frac{1}{2}(x\dot{y} - y\dot{x}). \tag{0.1}$$

On the other hand, assume that $\hat{c}(t) = (x(t), y(t), z(t))$ is an arbitrary curve in $\mathbb{R}^3$ with starting point $\hat{c}(a) = (0, 0, z_a)$ and endpoint $\hat{c}(b) = (0, y_b, z_b)$ and let $c(t) = (x(t), y(t))$ denote the projection of $\hat{c}(t)$ onto the (x, y)-plane. If the components of $\dot{\hat{c}}$ satisfy (0.1), then, after integration one obtains

$$z(b) - z(a) = \iint_\Omega dxdy = S,$$

where Ω equals the domain in the (x, y)-plane enclosed by c and the line segment on the y-axis bounded by the endpoints of c. More generally, if the initial point (x_a, y_a, z_a) of a curve $\hat{c}$ coincides with the origin $(0, 0, 0)$ and if the components of $\hat{c}$ satisfy (0.1), then $\hat{c}$ has the property that, for each $t_0 \in [a, b]$, the z-component of $\hat{c}(t_0)$ equals the area $S(t_0)$ of the domain in the (x, y)-plane enclosed by $c|_{[a,t_0]}$ and the line segment connecting the origin with $c(t_0)$.

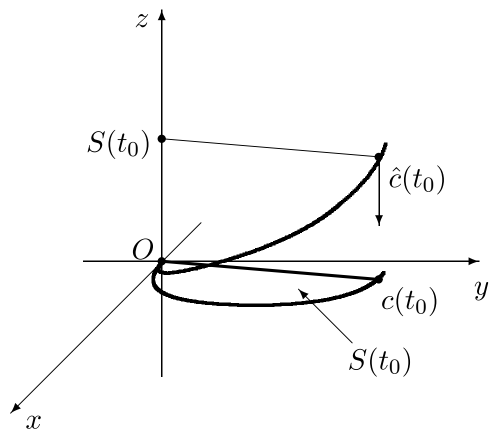

This allows us to reformulate the dual problem as follows. *Given a point* $(0, y_b, S)$ *in* $\mathbb{R}^3$*; what curve* $\hat{c}$*, connecting the origin with* $(0, y_b, S)$*, satisfies equation (0.1) and minimises the length of the projected curve in the* (x, y)*-plane.* (Note that, in comparison with Dido's problem, we assumed here that the endpoint $(0, y_b)$ of the projected curve is a fixed point on L. We return to this assumption when discussing the solution of Dido's problem.) We now briefly discuss, one by one, the various elements encountered in the above formulation.

Control theory and anchored bundles

The restriction (0.1) imposed on the class of curves $\hat{c}(t)$ is an example of what is called a *nonholonomic constraint*. It determines a restriction on the velocity vector of the curve: in the present case, the z-component of the velocity is completely determined by $x(t)$, $y(t)$, $\dot{x}(t)$ and $\dot{y}(t)$. We now rewrite this condition in order to introduce the concept of a control. We introduce two new variables v and w and we consider the following function:

$$\gamma : \mathbb{R}^3 \times \mathbb{R}^2 \to \mathbb{R}^3 : (x, y, z, v, w) \mapsto (v, w, \tfrac{1}{2}(xw - yv)).$$

A curve $\hat{c}$ in $\mathbb{R}^3$ satisfies condition (0.1) if and only if there exists functions $v(t)$ and $w(t)$ such that

$$\gamma\Big(\hat{c}(t), \big(v(t), w(t)\big)\Big) \equiv \frac{d\hat{c}}{dt}(t). \tag{0.2}$$

This correspondence easily follows from the definition of γ and by taking $v(t) = \dot{x}(t)$, $w(t) = \dot{y}(t)$ and $\frac{1}{2}(x(t)w(t) - y(t)v(t)) = \dot{z}(t)$.

We take the structure of (0.2) as an example in order to consider the following more general situation. Let γ be a smooth function, defined on $\mathbb{R}^n \times \mathbb{R}^k$, with values in $\mathbb{R}^n$. We call $\mathbb{R}^n$ the configuration space and $\mathbb{R}^k$ the control domain. Coordinates on the configuration space are denoted by $(q^1, \ldots, q^n)$ and on the control domain by $(u^1, \ldots, u^k)$. Consider a curve $u : [a, b] \to \mathbb{R}^k$ in the control domain and a curve $q : [a, b] \to \mathbb{R}^n$ in the configuration space. We say that $u(t)$ *controls* $q(t)$ if:

$$\dot{q}^i(t) = \gamma^i\big(q^1(t), \ldots, q^n(t), u^1(t), \ldots, u^k(t)\big), \quad \text{for } i = 1, \ldots, n.$$

This terminology is justified by noting that $q(t)$ is a solution to the following set of time dependent differential equations: $\dot{q} = \gamma(q, u(t))$. From uniqueness of solutions of differential equations it follows that $q(t)$ is completely determined by $u(t)$ and the initial condition $q(a)$. A pair $(q(t), u(t))$ is called *admissible* if $u(t)$ controls $q(t)$.

Besides the relation to Dido's problem, the "control systems" introduced above are important for the study of physical systems in technological applications. The curve $u(t)$ can be regarded as an external (human) input in a physical system, whose evolution is modeled by $q(t)$. The function γ expresses how the human input $u(t)$ affects the configuration of the system. The curve u is called the *control*. Let us illustrate these ideas in the case of a mechanical problem.

Consider a mechanical system with n degrees of freedom on which k forces act, represented by $F_i : \mathbb{R}^{2n} \to \mathbb{R}^n$ and which are controlled by k independent smooth functions $u^i(t)$ such that the total force acting on the system is given by $u^i F_i$. For the sake of simplicity we assume in addition that the kinetic energy metric of the system equals the standard euclidian metric on $\mathbb{R}^n$, i.e. all masses are equal to unity. We define a control system as follows. The configuration space of the control system is taken to be the phase space $\mathbb{R}^{2n}$ of the physical system and the control domain is identified with $\mathbb{R}^k$. Define the function γ as, with $(q, v) = (q^1, \ldots, q^n, v^1, \ldots, v^n) \in \mathbb{R}^{2n}$:

$$\gamma((q, v), u) = (v^1, \ldots, v^n, u^i F_i^1(q, v), \ldots, u^i F_i^n(q, v)).$$

It is now easily seen that a curve $(q(t), v(t))$ in $\mathbb{R}^{2n}$ is controlled by $u(t) = (u^1(t), \dots, u^k(t))$ iff

$$\begin{aligned} \dot{q}^i(t) &= v^i(t) \text{ and} \\ \dot{v}^i(t) &= u^j(t) F^i_j(q^1(t), \dots, q^n(t), v^1(t), \dots, v^n(t)) \end{aligned}$$

hold, for $i = 1, \dots, n$. In particular this implies that Newton's second law is satisfied:

$$\ddot{q}^i(t) = u^i(t) F_i(q(t), \dot{q}(t)).$$

Let us keep in mind the example of a hovercraft. Roughly speaking, the input force is a variable magnitude and variable direction force (i.e. two inputs) influenced by the turbine on the boat. The control $(u^1(t), u^2(t))$ is a measure for the "steering" of the pilot, i.e. u^1 and u^2 parameterise the magnitude and direction of the input force (for a detailed treatment, see [32]).

The above mentioned example, is only one of the many possible applications of control theory. We will see that the differential geometric structure, in which a control system is defined, is that of an *anchored bundle*. Roughly speaking, an anchored bundle consists of a bundle over the configuration manifold (here $\mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^n$) and an "anchor map", which is a mapping from the bundle to the tangent bundle of the configuration manifold (here $\gamma : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^n$). In Chapter I we will study some aspects of anchored bundles and introduce the notion of an admissible curve in this general setting. In particular, we will be interested in a geometric description of the set of points that can be reached by admissible curves having a fixed initial point. Moreover, we also consider the notion of *admissible loops* and prove that these curves generate a subgroup of the first fundamental group of the configuration space.

Optimal control theory

Let us return to the treatment of the dual of Dido's problem and assume that $(\hat{c}(t), u(t))$ is admissible, with $u(t) = (v(t), w(t)) \in \mathbb{R}^2$. We know that being admissible is equivalent to the condition (0.1) for the curve $\hat{c}$. As is clearly expressed in the formulation of the problem, we have to find the curve minimising the length of the projected curve $c(t)$ in the (x, y)-plane.

The length $\ell(c)$ of the projected curve c with respect to the standard metric is defined as:

$$\ell(c) = \int_a^b \sqrt{\dot{x}^2(t) + \dot{y}^2(t)}dt.$$

By making use of the fact that $v(t) = \dot{x}(t)$ en $w(t) = \dot{y}(t)$, the integrand of the above integral equals $\sqrt{v^2(t) + w^2(t)}$. When taking the latter expression as a function L on the control domain, i.e. $L(u) = \sqrt{v^2 + w^2}$ with $u = (v, w) \in \mathbb{R}^2$ arbitrary, we can define a *functional* $\mathcal{J}$ on the class of admissible curves:

$$\mathcal{J}(\hat{c}, u) = \int_a^b L(u(t))dt = \ell(c).$$

Using these definitions, we have that, the admissible curve solving the dual of Dido's problem, is precisely the admissible curve minimising this functional.

The theory dealing with similar problems on an arbitrary control system, is called *optimal control theory* and this is the topic of Chapter III. Recall the notations introduced in the previous section. Assume that L denotes a smooth function on $\mathbb{R}^n \times \mathbb{R}^k$ which is called the *cost function*, and let $(q, u) : [a, b] \to \mathbb{R}^n \times \mathbb{R}^k$ denote an admissible curve with $q(a) = q_a$ and $q(b) = q_b$. Then, in turn, the cost function determines a functional $\mathcal{J}$ on the class of admissible curves, as follows:

$$\mathcal{J}(q, u) = \int_a^b L(q(t), u(t))dt.$$

An admissible curve (q, u) and the associated control u is called *optimal* if for every admissible curve (q', u'), defined on $[a, b]$ with $q'(a) = q_a$ and $q'(b) = q_b$, the inequality $\mathcal{J}(q, u) \leq \mathcal{J}(q', u')$ holds.

In [47] necessary conditions for an admissible curve to be optimal are formulated and these conditions are stated in the *maximum principle* (for the sake of simplicity we leave out all technical details).

Theorem (The maximum principle). *Assume that the admissible curve $(q, u) : [a, b] \to \mathbb{R}^n \times \mathbb{R}^k$ is optimal with respect to a given cost function $L(q, u)$. Then, there exists a pair $(\lambda, p(t))$, with $p : [a, b] \to \mathbb{R}^n$ and $\lambda \in \mathbb{R}$, such that the following conditions are satisfied:*

1. *$(q(t), p(t))$ is a solution to the following Hamiltonian system,*

$$\begin{aligned} \dot{p}_i &= -\frac{\partial H}{\partial q^i}(q,u,p), \\ \dot{q}^i &= \frac{\partial H}{\partial p_i}(q,u,p) = \gamma^i(q,u), \end{aligned}$$

   where the Hamiltonian H is given by $\mathbb{R}^n \times \mathbb{R}^k \times \mathbb{R}^n \to \mathbb{R} : (q,u,p) \mapsto p_i\gamma^i(q,u) + \lambda L(q,u)$;

2. *$H(q(t), u(t), p(t)) \geq H(q(t), u', p(t))$ with $u' \in \mathbb{R}^k$ arbitrary and for all $t \in [a,b]$ (or: $H(q(t), u', p(t))$ reaches a maximal value at $u' = u(t)$);*

3. *the pair $(p(t), \lambda)$ is not trivial, i.e. $p(t) \neq 0$ or $\lambda \neq 0$;*

4. *$\lambda = 0$ or $\lambda = -1$.*

An important observation is that the maximum principle also allows the constant λ to be zero. In the past, this subtlety was often wrongly interpreted (see e.g. [50, 51]). If $\lambda = 0$ then the Hamiltonian H equals $p_i\gamma^i(q,u)$ and, therefore, in that case the conditions of the maximum principle become independent of the cost function ! This extraordinary situation motivated some people to call the corresponding admissible curves *abnormal* extremals, precisely because the above conditions are necessary for admissible curves to be optimal and that one would expect that these conditions naturally depend on the cost function. R. Montgomery proved in 1994 that there exist abnormal extremals that are minimising [45].

In Chapter III, we study optimal control theory from a differential geometric point of view. We will prove a differential geometric version of the maximum principle in the case of optimal control theory for time dependent and time independent (autonomous) control systems. We also consider the case where the endpoints of an optimal control are allowed to vary on a given submanifold. We introduce the notion of a *variation of an admissible curve* and we define a *cone of variations* in the tangent space of the configuration manifold at the endpoint of an admissible curve. The geometric picture behind this cone, is that it will contain all possible directions near this endpoint that are also reachable by an admissible curve. The notion of abnormal extremal is thoroughly investigated and it will turn out that these abnormal extremals are precisely those admissible curves that do not allow variations in all possible directions, i.e. in this case, the cone of variations is not equal to the full tangent space at the endpoint. Amongst others, we will prove necessary and sufficient conditions for an admissible curve to be abnormal. At the end of Chapter III, we mention several possible applications of the maximum principle in differential geometry, such as Hamiltonian and Lagrangian mechanics on (affine) Lie algebroids, simple mechanical control systems, etc.

Generalised connections

In the treatment of the maximum principle for control systems on anchored bundles, an important role is assigned to the notion of a generalised connection. In Chapter II we develop the theory of generalised connections over a bundle map and show, among others, that it allows us to define a *transport operator* along admissible curves. In particular, the variational cone of an admissible curve can be defined by transporting certain tangent vectors along the admissible curve under consideration.

The notion of a generalised connection, presented in Chapter II, was inspired on recent work by R.L. Fernandes [11, 14, 15]. He introduced a generalisation of the notion of connection, some of the essential elements of which can be found in earlier work by, among others, I. Vaisman [56], Y.C. Wong [59] and F. Kamber and P. Tondeur [21] on respectively, contravariant connections, pseudo-connections and partial connections. The relevance of creating such a model in which all the above mentioned different notions of connection fit, lies within the fact that it brings within the reach of connection theory certain differential geometric structures which have not been considered previously from such a point of view. As is already mentioned, one of these possible fields of applications is geometric control theory.

Sub-Riemannian geometry versus nonholonomic mechanics

Another field of applications of optimal control theory can be found in sub-Riemannian geometry. Let us first describe very briefly what is meant by a sub-Riemannian geometry using the definitions above. Let $\gamma : \mathbb{R}^k \times \mathbb{R}^n \to \mathbb{R}^n$ denote a control system. Consider the map $\gamma_q : u' \mapsto \gamma(q, u')$, defined for an arbitrary fixed point $q \in \mathbb{R}^n$. Assume that γ_q is *injective* and *linear* for all $q \in \mathbb{R}^n$, i.e. $\gamma^i(q, u) = \sum_{j=1}^{k} \gamma^i_j(q) u^j$ for $i = 1, \ldots, n$. If $g_{ij}(q)$, for $i, j = 1, \ldots, n$ denotes a Riemannian metric on $\mathbb{R}^n$, then we can define a

Riemannian bundle metric on the trivial bundle $\mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^n$ as follows: $h_{mn} = g_{ij}\gamma^i_m\gamma^j_n$ for $m, n = 1, \ldots, k$. Since γ_q is assumed to be injective for all $q \in \mathbb{R}^n$, this bundle metric is well defined. The control system and the bundle metric determine what is called a *sub-Riemannian structure*. Using this metric one can define the length of an admissible curve $(q(t), u(t))$ as usual by

$$\int_a^b \sqrt{g_{ij}(q(t))\dot{q}^i(t)\dot{q}^j(t)}dt.$$

Using the definition of an admissible curve, it is easily seen that the following equality holds:

$$g_{ij}(q(t))\dot{q}^i(t)\dot{q}^j(t) = h_{mn}(q(t))u^m(t)u^n(t).$$

One of the main topics is to find those admissible curves that minimise length among all admissible curves connecting two given points or, equivalently, since $\dot{q}(t) = \gamma(q(t), u(t))$, those admissible curves that are optimal with respect to the cost function $L(q, u) = \sqrt{h_{mn}(q)u^m u^n}$. It is easily seen that the dual of Dido's problem fits into the definition of a sub-Riemannian geometry.

In Chapter IV we discuss the general setting of sub-Riemannian geometry and its relation to optimal control theory. Further, we shall demonstrate the applicability of the theorems on abnormal curves from Chapter III and we present an alternative approach to, respectively, sub-Riemannian geometry and nonholonomic mechanics (i.e. mechanical systems subjected to nonholonomic constraints) by means of the generalised connections introduced in Chapter II. We show that the geometrical framework in which nonholonomic mechanics is modeled also fixes a sub-Riemannian structure. Necessary and sufficient conditions are derived for "geodesics" (curves satisfying the necessary conditions of the maximum principle for $\lambda = -1$) to be solutions of the equations of motion of a free nonholonomic particle. In Chapter V we further discuss an alternative reduction procedure of nonholonomic mechanics with symmetry using the tools developed in Chapter II on generalised connections. The importance of this new reduction procedure lies within the fact that the conditions on the nonholonomic constraint distribution, necessary for applying the theory, are less restrictive in comparison with other reduction procedures. We consider an example of a nonholonomic mechanical system with symmetry, the Snakeboard, to fix these ideas.

The solution of Dido's problem

To close this introduction, we point out how the maximum principle can be used to solve the dual of Dido's problem. Using the results from Chapter IV we can easily show that no abnormal extremals can occur for this particular problem. For $\lambda = -1$, the Hamiltonian system in the maximum principle can be reformulated such that the solution $(x(t), y(t), z(t))$ has to satisfy the following system of differential equations:

$$\begin{aligned} \dot{p}_x &= -\tfrac{1}{2}\dot{y}p_z, & \dot{x} &= p_x - \tfrac{1}{2}yp_z, \\ \dot{p}_y &= \tfrac{1}{2}\dot{x}p_z, & \dot{y} &= p_y + \tfrac{1}{2}xp_z, \\ \dot{p}_z &= 0, & \dot{z} &= \tfrac{1}{2}(x\dot{y} - y\dot{x}), \end{aligned}$$

for some $p(t) = (p_x(t), p_y(t), p_z(t))$. From this it follows that $\ddot{x} = -\dot{y}\omega$ and that $\ddot{y} = \dot{x}\omega$, with $p_z = \omega$ constant. The general solution for $(x(t), y(t))$ through the origin is of the form:

$$\begin{aligned} x(t) &= A\sin(\omega t) + B\big(\cos(\omega t) - 1\big), \\ y(t) &= A(1 - \cos(\omega t)) + B\sin(\omega t), \end{aligned}$$

for $t \in [0,1]$ and with A and B constants. Let R be the positive constant defined by $R^2 = A^2 + B^2$. There exists an $\alpha \in [0, 2\pi[$ such that $A = R\cos\alpha$ and $B = R\sin\alpha$ and we obtain:

$$\begin{aligned} x(t) &= R\big(\sin(\omega t + \alpha) - \sin(\alpha)\big), \\ y(t) &= R\big(\cos(\alpha) - \cos(\omega t + \alpha)\big), \end{aligned}$$

for $t \in [0,1]$. The shape of the above curve is the circle segment with centre at the point $(-R\sin(\alpha), R\cos(\alpha))$ and with radius R. The constants R, α, ω follow from the condition that the circle passes through a fixed endpoint (x_1, y_1) with $x_1 = 0$ and that $z(1) = S$, where $z(t)$ is given by:

$$z(t) = \frac{R^2}{2}\big(\omega t + \sin(\omega t)\big).$$

In Chapter III we prove that, in case the endpoint is a variable point on L, then the constants R, α, ω are determined by the conditions $p_y(1) = 0$, $x(1) = 0$ and $z(1) = S$. In particular, this implies that $\alpha = 0$ and $\omega = \pi$, i.e. the solution curve is a half circle, with centre at a point of L. The condition $z(1) = S$ then precisely expresses that $S = \frac{1}{2}\pi R^2$.

Notations and conventions

We only consider real, Hausdorff, second countable smooth manifolds, and by smooth we will always mean C^∞. The set of (real valued) smooth functions on a manifold M will be denoted by $C^\infty(M)$, the set of smooth vector fields by $\mathfrak{X}(M)$ and the set of smooth one-forms by $\mathfrak{X}^*(M)$. Given a fibre bundle $\pi : E \to M$, then the set of all smooth sections defined on an open neighbourhood of a point $x \in M$ will be denoted by $\Gamma_x(\pi)$, and we further put $\Gamma(\pi) = \cup_{x\in M}\Gamma_x(\pi)$ (sometimes we also write $\Gamma(E)$). Note, in particular, that any global section of π, if it exists, belongs to $\Gamma_x(\pi)$ for all x. The fibre of π over a point $x \in M$ will be indicated by E_x. Given a smooth map $f : N_1 \to N_2$ between two manifolds, we will denote the tangent map of f by $Tf : TN_1 \to TN_2$.

Let V be a real vector space, and W a subspace, then the annihilator space of W is given by

$$W^0 = \{\beta \in V^* \mid \langle \beta, w\rangle = 0 \;\forall w \in W\}.$$

If E is a vector bundle over a manifold M and F any vector subbundle, then the annihilator bundle F^0 of F is the subbundle of the dual bundle E^* of E over M whose fibre over a point $x \in M$ is the annihilator space of the subspace F_x of E_x. The domain of a curve will usually be taken to be a closed (compact) interval in $\mathbb{R}$. If a group G' is a subgroup of a group G, then we write $G' < G$.

Parts of the work presented in this thesis have been published in the following papers [5, 27, 28, 29, 30, 31].

I

Anchored bundles

In this chapter we describe the basic structure on which our study of generalised connections relies, namely that of an anchored bundle. Let M denote an arbitrary n-dimensional manifold with tangent bundle $\tau_M : TM \to M$. The conceptual idea of an anchored bundle is that one considers a bundle over M which is related to TM, in such a way that, for further developments, the bundle can be taken as an alternative to the tangent bundle of M. The notion of an anchored bundle already appears, for instance, in the work of P. Popescu [48], who also uses the denomination “relative tangent space”.

Definition 0.1. An anchored bundle on M is a pair (ν, ρ) where, $\nu : N \to M$ denotes a fibre bundle over M, and $\rho : N \to TM$ is a bundle map, fibred over the identity on M. We call ρ the *anchor map* of the anchored bundle.

The following diagram is commutative:

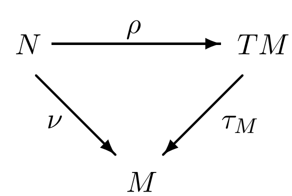

We say that an anchored bundle (ν, ρ) is *linear*, if ν is a vector bundle and ρ is a linear bundle morphism.

Consider two anchored bundles (ν', ρ') and (ν, ρ) with base manifolds respectively M' and M. An *anchored bundle morphism* $(f, \overline{f})$ from (ν', ρ') to (ν, ρ) consists of a smooth mapping $\overline{f} : M' \to M$ and a bundle morphism $f : N' \to N$ fibred over $\overline{f}$, in such a way that the following equality holds:

$$T\overline{f} \circ \rho' = \rho \circ f.$$

We say that $(f, \overline{f})$ is an anchored bundle isomorphism if f is a bundle isomorphism (see e.g. [49]) and if, in addition, $(f^{-1}, \overline{f}^{-1})$ is also an anchored bundle morphism. In this case we can write $\rho' = T(\overline{f})^{-1} \circ \rho \circ f$ and, conversely, $\rho = T\overline{f} \circ \rho' \circ f^{-1}$. Next, if $\overline{f}$ is an injective immersion, then we say that (ν', ρ') is an anchored subbundle of (ν, ρ). Note that, although we made no assumption on the restriction of f to the fibres of ν', the anchor ρ' is completely determined by $\rho' = T(\overline{f})^{-1} \circ \rho \circ f$, which is well defined since $\overline{f}$ is an immersion. Assume that both anchored bundles are linear. Then, we say that f is a *linear homomorphism* if $f : N' \to N$ is a linear bundle map. The following commutative diagram represents an anchored bundle morphism:

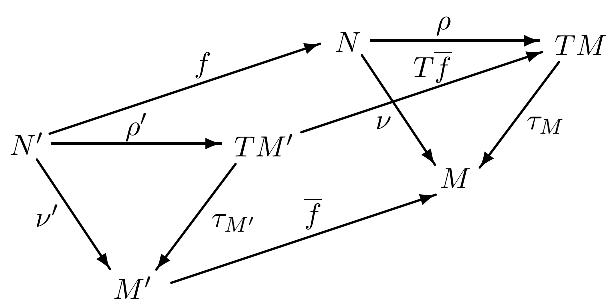

For brevity, we will often refer to the bundle morphism $f : N' \to N$ as an anchored bundle morphism, with the base mapping $\overline{f}$ then being understood.

1 The foliation on anchored bundles

In this section we need some elements from the theory of integrability of distributions, developed by H.J. Sussmann [52] (see also [37]). We first briefly recall the basic definitions and some key results on distributions, before applying them to anchored bundles. We also use this section to fix some notations regarding composite flows and concatenations of integral curves of vector fields.

Consider a manifold M and assume that F is a *differentiable distribution* on M, i.e. F is a subset of TM such that, for any point $x \in M$, the fibre $F_x = F \cap T_x M$ is a linear subspace of $T_x M$ and such that F_x is spanned by a finite number, say k, of (local) independent vector fields on M evaluated at x and such that these vector fields satisfy $X_i(y) \in F_y$ for all y in the domain of X_i $(i = 1, \ldots, k)$. The number k is called the *rank* of the distribution F at x: it is the dimension of F_x. Note that, in the above definition, a distribution need not have constant rank in general. If F has constant rank, we say that F is a *regular distribution* .

A distribution is said to be *completely integrable* if there exists, for every $x \in M$, an immersed connected submanifold $i : L \hookrightarrow M$ containing x and such that $T_yL = F_y$, for each $y \in L$. A submanifold L satisfying the above conditions is called a leaf of the distribution if it is maximal in the sense that, given any other immersed submanifold L' verifying the above conditions and containing L, then $L' = L$. It can be proven that these leaves are unique and form a partition of M which is called the foliation induced by the completely integrable distribution. Note that, by definition, the distribution F has constant rank on the points of each leaf.

Let $\mathcal{F}$ be a family of vector fields on M, each defined on an open subset of M. We say that $\mathcal{F}$ is *everywhere defined* if, given any $x \in M$, there exists an element X of $\mathcal{F}$ containing x in its domain. An everywhere defined family of vector fields $\mathcal{F}$ generates a distribution F in the following way. Put

$$F_x = \operatorname{span}\{X(x) \mid X \in \mathcal{F}, x \in \operatorname{dom} X\},$$

then it is readily seen that F is a differentiable distribution. On the other hand, it easily follows from the definition that any differentiable distribution is generated by an everywhere defined family of vector fields. H.J. Sussmann [52] has shown that one can always construct the smallest completely integrable distribution $\widetilde{F}$ containing the distribution F. In order to describe this construction, we first need the notion of a composite flow.

Assume that we have fixed an ordered ℓ-tuple $\mathcal{X} = (X_\ell, \ldots, X_1)$ of (not necessarily different) vector fields on M, and let us represent the flow of X_i by $\{\phi^i_t\}$. With a view on later constructions, we have chosen the ordering of the vector fields in $\mathcal{X}$ with indices decreasing from left to right.

The *composite flow* of $\mathcal{X}$ is the map

$$\Phi : V \subset \mathbb{R}^\ell \times M \to M : ((t_\ell, \ldots, t_1), x) \mapsto \phi^\ell_{t_\ell} \circ \cdots \circ \phi^1_{t_1}(x),$$

defined on some open subset V of $\mathbb{R}^\ell \times M$. For brevity we shall write $\Phi((t_\ell, \ldots, t_1), x) = \Phi_T(x)$, where $T = (t_\ell, \ldots, t_1)$. We shall sometimes refer to T as the *composite flow parameter*. For each fixed T, Φ_T determines a diffeomorphism from an open subset of M (which may be empty) to another open subset of M. It can also be proven that, if we fix a point $x \in M$, then the map $T' \mapsto \Phi_{T'}(x)$ is smooth and is defined on an open neighbourhood of T. For further details about the domain of composite flows, we refer the reader to [37].

Assume that we are given two composite flows: Φ of $\mathcal{X} = (X_\ell, \dots, X_1)$ and Ψ of $\mathcal{Y} = (Y_{\ell'}, \dots, Y_1)$. The *composition of* Φ and Ψ is the composite flow, denoted by $\Psi \star \Phi$, of the $(\ell' + \ell)$-tuple $(Y_{\ell'}, \dots, Y_1, X_\ell, \dots, X_1)$. One can verify that the composition of composite flows is associative. Using these notations, it is easily seen that, for instance, Φ equals $\phi^\ell \star \cdots \star \phi^1$. If T is a composite flow parameter for Φ and T' for Ψ, then we put $T' \star T = (T', T) \in \mathbb{R}^{\ell' + \ell}$, which is a composite flow parameter for $\Psi \star \Phi$.

The composite flow Φ of $\mathcal{X} = (X_\ell, \dots, X_1)$ is said to be *generated* by an everywhere defined family of vector fields $\mathcal{F}$ if $\mathcal{X}$ is an ordered ℓ-tuple of elements of $\mathcal{F}$. Using all composite flows generated by $\mathcal{F}$, we can define an equivalence relation on the points of M, denoted by $\cdot \overset{\mathcal{F}}{\longleftrightarrow} \cdot$.

Definition 1.1. Assume that $x, y \in M$. Then $x \overset{\mathcal{F}}{\longleftrightarrow} y$ if there exists a composite flow Φ generated by $\mathcal{F}$ and a composite flow parameter T such that $\Phi_T(x) = y$.

In the sequel, when no confusion can arise, we will also simply write $x \leftrightarrow y$, dropping the explicit reference to $\mathcal{F}$. It is easily seen that the relation $\leftrightarrow$ is transitive (see the above definition of the composition of composite flows) and reflexive (take $T = (0, \dots, 0)$). If Φ is a composite flow of $\mathcal{X} = (X_\ell, \dots, X_1)$ and $\Phi_T(x) = y$ for some $T = (t_\ell, \dots, t_1)$, then the composite flow $\widetilde{\Phi}$ of $\widetilde{\mathcal{X}} = (X_1, \dots, X_\ell)$ and the composite flow parameter $\widetilde{T} = (-t_1, \dots, -t_\ell)$ satisfy $\widetilde{\Phi}_{\widetilde{T}}(y) = x$. Since $\widetilde{\Phi}$ is also generated by $\mathcal{F}$, this makes the relation symmetric. In the following we always assume that the distribution F is the one generated by a given family $\mathcal{F}$. The following theorem is due to H.J. Sussmann and can be found in [37, 52].

Theorem 1.1. *The smallest completely integrable distribution $\widetilde{F}$ containing F is the distribution generated by the everywhere defined family $\widetilde{\mathcal{F}}$ containing all vector fields of the form $\Phi_T^* Y$, where $Y \in \mathcal{F}$ and Φ is a composite flow generated by $\mathcal{F}$. Moreover, the equivalence relations associated with $\mathcal{F}$ and $\mathcal{F}'$ are equal and the leaves of the distribution $\widetilde{F}$ are the equivalence classes of $\overset{\mathcal{F}}{\longleftrightarrow}$.*

Consider the distribution $\widetilde{F}$ and let $[X, Y]$ denote the Lie bracket of two vector fields in $\mathcal{F}$. It is easily seen that $[X, Y]$ is a vector field in $\widetilde{F}$. Indeed, let $\{\phi_t\}$ be the flow of X and observe that $\phi_t^* Y$ is in $\widetilde{\mathcal{F}}$. Then, for each $x \in M$, the curve $t \mapsto (\phi_t^* Y)(x)$ is entirely contained in the linear space $\widetilde{F}_x$, and so is its tangent vector:

$$\left.\frac{d}{dt}\right|_0 \big(\phi_t^* Y\big)(x) = [X, Y](x).$$

This reasoning can be easily extended to any finite number of iterated Lie brackets of vector fields in $\mathcal{F}$. In fact, this observation is rather important since it leads to an alternative proof of the Chow's Theorem (see [52]), which says that if the distribution generated by the family $\mathcal{F}$ and the family of all iterated Lie brackets of vector fields in $\mathcal{F}$, equals TM, then any two points in a connected component of M can be connected by a curve whose tangent vector is everywhere contained in F.

Assume that $\mathcal{X} = (X_\ell, \ldots, X_1)$ is an arbitrary ordered ℓ-tuple of vector fields, with composite flow Φ. Fix a value $(t_\ell, \ldots, t_1)$ of the composite flow parameter T. A *concatenation of integral curves* of $\mathcal{X}$ trough $x \in M$ is a piecewise smooth curve $\gamma : [a, a + |t_1| + \cdots + |t_\ell|] \to M$ defined as follows,

$$\gamma(t) = \begin{cases} \phi^1_{\mathrm{sgn}(t_1)(t-a)}(x) & \text{for} \quad t \in [a, a_1], \\ \phi^2_{\mathrm{sgn}(t_2)(t-a_1)}(\phi^1_{t_1}(x)) & \text{for} \quad t \in\,]a_1, a_2], \\ \vdots & \\ \phi^\ell_{\mathrm{sgn}(t_\ell)(t-a_{\ell-1})}(\cdots(\phi^1_{t_1}(x))\cdots) & \text{for} \quad t \in\,]a_{\ell-1}, a_\ell], \end{cases}$$

where $a_i = a + \sum_{j=1}^{i} |t_j|$, $\mathrm{sgn}(t_i) = t_i/|t_i|$ for $t_i \neq 0$ and $\mathrm{sgn}(0) = 0$. Note that if $t \in]a_{i-1}, a_i[$ then $\dot\gamma(t) = \mathrm{sgn}(t_i) X_i(\gamma(t))$ and, hence, the restriction of γ to $]a_{i-1}, a_i[$ is an integral curve of X_i if $t_i > 0$ (or of $-X_i$ if $t_i < 0$). Note that $\gamma(a_\ell) = \Phi_T(x)$, i.e. the endpoint of γ coincides with the image of x under the composite flow Φ_T. It is easily seen that in the specific case where $\mathcal{X}$ is generated by a given everywhere defined family $\mathcal{F}$ of vector fields, the concatenation of integral curves of $\mathcal{X}$ through $x \in M$ is entirely contained in the leaf L_x through x of the associated completely integrable distribution $\widetilde{F}$.

Let us now proceed towards the construction of an everywhere defined family of vector fields on M, given an anchored bundle (ν, ρ) on M. Consider an arbitrary (local) section σ of ν, i.e. $\sigma : M \to N$ is a smooth map with $(\nu \circ \sigma)(x) = x$. We can associate to the section σ of ν, the vector field $\rho \circ \sigma$ on M. Let $\mathcal{D}$ denote the set of all vector fields of the form $\rho \circ \sigma$ for some section σ of ν. Clearly, $\mathcal{D}$ is everywhere defined and using the notations introduced above, the manifold M is equipped with a distribution D generated by $\mathcal{D}$ (with $D = \mathrm{im}\,\rho$ if (ν, ρ) is linear) and we can consider the smallest completely integrable distribution $\widetilde{D}$ containing D. The leaf of the foliation determined by $\widetilde{D}$ through x is denoted by L_x.

Consider the immersion $i : L_x \hookrightarrow M$, and let $\nu' : N' = L_x \times_M N \to L_x$ denote the pull-back bundle of ν under i, i.e. $(y, s) \in N'$ if $i(y) = \nu(s)$. Since

i is an immersion, we can define an anchor map $\rho' : N' \to TL_x$ as follows: $T_y i(\rho'(y,s)) = \rho(s)$, given any $(y,s) \in N'$. The projection $\pi_2 : N' \to N$ of N' onto the second factor, determines an anchored bundle morphism, fibred over the immersion i, i.e. (ν', ρ') is an anchored subbundle of (ν, ρ). We shall call (ν', ρ') *the pull-back anchored bundle under* i.

Before passing to the next section, we first give two examples of anchored bundles and the distribution they induce. The first example is taken from [45], where it was used in the context of sub-Riemannian geometry to construct length-minimising strictly abnormal extremals (see the Introduction and Chapter IV). The other example is taken from [37] and provides a non-trivial completely integrable distribution on $\mathbb{R}^2$.

Example 1.2. Take $M = \mathbb{R}^3$ and $\nu : N = M \times \mathbb{R}^2 \to M$ a trivial bundle over M (we use cylindrical coordinates (r, θ, z) on M). Consider the following two vector fields on M: $X_1 = \partial/\partial r$ and $X_2 = \partial/\partial\theta - p(r)\partial/\partial z$, where $p(r)$ is a function on $\mathbb{R}$ with a single non degenerate maximum at $r = 1$:

$$\left.\frac{d}{dr}p(r)\right|_{r=1} = 0 \quad \text{and} \quad \left.\frac{d^2}{dr^2}p(r)\right|_{r=1} < 0.$$

Such a function always exists (take, for instance, $p(r) = \frac{1}{2}r^2 - \frac{1}{4}r^4$). Let $\rho : N \to TM$ denote the map defined by $\rho(x, u^1, u^2) = u^1 X^1(x) + u^2 X^2(x)$, with $x = (r, \theta, z) \in M$. Note that the vector fields in the family $\mathcal{D}$ associated with this anchored bundle are of the form $x \mapsto \rho(\sigma(x)) = \sigma^1(x) X^1(x) + \sigma^2(x) X^2(x)$ with $\sigma(x) = (x, \sigma^1(x), \sigma^2(x))$ a section of ν (where $\sigma^1, \sigma^2 \in C^\infty(M)$). It is easily seen that (ν, ρ) is a linear anchored bundle. The flows of X_1, X_2 are denoted by $\{\phi_t\}$, $\{\psi_t\}$, respectively. In particular, we have $\phi_t(r, \theta, z) = (t+r, \theta, z)$, $\psi_t(r, \theta, z) = (r, \theta+t, z-p(r)t)$. The foliation induced by $D = \operatorname{im}\rho$ is trivial. Indeed, all iterated Lie brackets of the two vector fields X^1 and X^2 span the total tangent space at each point of M, implying that $\widetilde{D} = TM$ and M itself is the only leaf.

Example 1.3. Let $M = \mathbb{R}^2$ and let $N = M \times \mathbb{R}^2$, with $\rho(x, y, u^1, u^2) = u^1 X(x, y) + u^2 Y(x, y)$, where $X = \partial/\partial x$ and $Y = y\partial/\partial y$. The distribution F on M defined by $F = \operatorname{im}\rho$ satisfies $F = \widetilde{F}$, since $[X, Y] = 0$, i.e. F is completely integrable. The two 2-dimensional submanifolds $\{y < 0\}$, $\{y > 0\}$ and the 1-dimensional submanifold $\{y = 0\}$ are the leaves of the foliation on M. We will use this example to show that Lemma 2.1 in the following section is non-trivial. We shall construct a curve, which is tangent to F, i.e. has tangent vector everywhere contained in F, but, such that the curve itself is not entirely contained in a single leaf. Indeed, consider $\widetilde{c} : \mathbb{R} \to$

$M : t \mapsto (t, t^3)$. It is readily seen that $\dot{\widetilde{c}}(t) = X(\widetilde{c}(t)) + 3t^{-1}Y(\widetilde{c}(t)) \in F_{\widetilde{c}(t)}$ for $t \neq 0$ and $\dot{\widetilde{c}}(0) = X(0,0) \in F_x$. However $\widetilde{c}$ passes through the three leaves of F.

2 ρ-admissible curves

We introduce here the notion of a *ρ-admissible curve*. By a smooth curve in a manifold M we will always mean a C^∞ map $c : I \to M$, where $I \subseteq \mathbb{R}$ may be either an open or a closed (compact) interval. In the latter case, the denominations "path"or "arc"are also frequently used in the literature but, for simplicity, we will make no distinction in terminology between both cases. For a curve defined on a closed interval, say $[a, b]$, it is tacitly assumed that it admits a smooth extension to an open interval containing $[a, b]$. Fix an anchored bundle (ν, ρ) on M.

Definition 2.1. Let $c : [a, b] \to N$ denote a smooth curve in N, and let $\widetilde{c} = \nu \circ c$ denote the projected curve in M, called the *base curve of* c. Then c is called a *smooth ρ-admissible curve* if $\rho \circ c = \dot{\widetilde{c}}$.

Local coordinates on M will be denoted by (q^i) and corresponding bundle adapted coordinates on N by (q^i, u^a), with $i = 1, \ldots, n$ and $a = 1, \ldots, k$, where k is the dimension of the typical fibre of N. If we write the anchor map ρ locally as

$$\rho(q^i, u^a) = \gamma^j(q^i, u^a)\frac{\partial}{\partial q^j}, \tag{2.1}$$

then a smooth ρ-admissible curve $c(t) = (q^i(t), u^a(t))$ locally satisfies the equation $\gamma^j(q^i(t), u^a(t)) = \dot{q}^j(t)$. In order to introduce a suitable concept of "control curves" (see Chapter III) or of "leafwise holonomy" in the framework of ρ-lifts, it turns out that the class of ρ-admissible curves in N should be further extended to curves admitting (a finite number of) discontinuities in the form of certain 'jumps' in the fibres of N, such that the corresponding base curve is piecewise smooth. For a more precise definition of a "piecewise" ρ-admissible curve we first consider the composition of smooth ρ-admissible curves.

The *composition* of a finite number of, say ℓ, smooth ρ-admissible curves $c_i : [a_{i-1}, a_i] \to N$ for $i = 1, \ldots, \ell$, whose base curves satisfy the boundary conditions $\widetilde{c}_i(a_i) = \widetilde{c}_{i+1}(a_i)$ for $i = 1, \ldots, \ell - 1$, is the map $c_\ell \cdot \ldots \cdot c_1$:

$[a_0, a_\ell] \to N$ defined by

$$c_\ell \cdot \ldots \cdot c_1(t) = \begin{cases} c_1(t) & t \in [a_0, a_1], \\ c_2(t) & t \in]a_1, a_2], \\ \vdots & \\ c_\ell(t) & t \in]a_{\ell-1}, a_\ell]. \end{cases} \tag{2.2}$$

Note that the base curve of $c_\ell \cdot \ldots \cdot c_1$ is a piecewise smooth curve. However, in general $c_\ell \cdot \ldots \cdot c_1$ is discontinuous at $t = a_i$, $i = 1, \ldots, \ell - 1$. The composition $c = c_\ell \cdot \ldots \cdot c_1$ is called *a piecewise ρ-admissible curve*, or simply a *ρ-admissible curve*. We now proceed towards the following important result, which tells us that the base curve of a ρ-admissible curve is always entirely contained in a leaf of the foliation on M, induced by the everywhere defined family of vector fields $\mathcal{D}$ on M (see the previous section).

Lemma 2.1. *The base curve $\widetilde{c}$ of a ρ-admissible curve $c : [a, b] \to N$ is entirely contained in the leaf L_x, with $x = \widetilde{c}(a)$.*

Proof. It is sufficient to prove this result for c smooth. For any point $x \in M$, consider a coordinate neighbourhood U centred at x with coordinates $(q^1, \ldots, q^n)$ adapted to the foliation induced by $\mathcal{D}$, such that: (1) if $q^{p+1}(y) = \cdots = q^n(y) = 0$, then $y \in L_x$, and (2) the coordinate functions $q^1, \ldots, q^p$ determine coordinates on the leaf L_x (this is always possible since L_x is an immersed submanifold). Upon restricting U to a smaller subset, if necessary, we may always assume, in addition, that the fibre bundle ν is trivial over U and we denote the adapted bundle coordinates on N by (q^i, u^a) for $i = 1, \ldots, n$ and $a = 1, \ldots, k$. In the following we only consider such coordinate charts. Recall the definition of the pull-back anchored bundle (ν', ρ') under $i : L_y \hookrightarrow M$. Note that $(q^1, \ldots, q^p, u^1, \ldots, u^k)$ is a bundle adapted coordinate chart on N'.

Fix a suitable coordinate chart on M (in the sense specified above) containing the point $x = \widetilde{c}(a)$ and assume that c is written in the adapted bundle coordinates as $(\widetilde{c}^i(t), u^a(t))$. Let $\widetilde{d}$ denote the solution curve in L_x of the following system of differential equations:

$$\dot{\widetilde{d}}^i(t) = \rho'^i(\widetilde{d}^1(t), \ldots, \widetilde{d}^p(t), u^1(t), \ldots, u^k(t)), \; i = 1, \ldots, p,$$

with initial condition $\widetilde{d}(a) = x$. From standard arguments we know that $\widetilde{d}$ is defined on an interval containing $[a, a + \epsilon[$ for some $\epsilon > 0$.

Consider the curve $\widetilde{d}' = i \circ \widetilde{d} : [a, a+\epsilon[\to M$ in M. Then we have, by uniqueness of solutions of differential equations, that $\widetilde{d}' = \widetilde{c}|_{[a,a+\epsilon[}$, since the curves $\widetilde{d}'(t)$ and $\widetilde{c}(t)$ both satisfy the system of differential equations $\dot{q}^i = \rho^i(q, u(t))$, $i = 1, \ldots, n$. Indeed, for $\widetilde{c}$ this is trivial and for $\widetilde{d}'$ we have

$$\dot{\widetilde{d}}'(t) = Ti(\rho'(\widetilde{d}(t), c(t))) = \rho(\widetilde{d}^i(t), 0, u^a(t)).$$

Therefore, we conclude that $\widetilde{c}|_{[a,a+\epsilon[}$ is contained in the leaf L_x, since by making use of the coordinate system, we have $\widetilde{c}^i(t) = 0$ for $t \in [a, a+\epsilon[$ and $i = p+1, \ldots, n$. Taking the limit from the left at $t = a+\epsilon$, we obtain that $\widetilde{c}^i(a+\epsilon) = 0$ for $i = p+1, \ldots, n$, i.e. $\widetilde{c}(a+\epsilon) \in L_x$. We can repeat the above reasoning for the curve $c|_{[a+\epsilon,b]}$, i.e. starting from the point $\widetilde{c}(a+\epsilon)$ instead of the point x. We thus obtain that $\widetilde{c}(t) \in L_x$ for all $t \in [a, a+\epsilon+\epsilon']$ for some $\epsilon' > 0$. Continuing this way, we eventually obtain that the entire curve $\widetilde{c}$ is contained in L_x, which concludes the proof. □

It can be seen that the curve $\widetilde{c}$ constructed in Example 1.3 does not contradict the previous lemma, although $\widetilde{c}$ is a curve tangent to the distribution $\operatorname{im}\rho$. Indeed, $\widetilde{c}$ can not be written as the base curve of a ρ-admissible curve since it has a singularity at $t = 0$.

Consider two anchored bundles (ν', ρ') and (ν, ρ), and an anchored bundle morphism f between them, i.e. $f : N' \to N$ fibred over $\overline{f} : M' \to M$. Let c' denote a ρ'-admissible curve. Consider the curve $c = f \circ c'$ in N, and let $\widetilde{c}$, resp. $\widetilde{c}'$, denote the base curve of c, resp. c'. Then, we have that c is ρ-admissible, since

$$\rho \circ c = \rho \circ f \circ c' = T\overline{f} \circ \rho' \circ c' = T\overline{f} \circ \dot{\widetilde{c}}' = \dot{\widetilde{c}}.$$

Let $c : [a, b] \to N$ denote a ρ-admissible curve. If $x = \widetilde{c}(a)$ and $y = \widetilde{c}(b)$, then we say that c *takes* x *to* y, and we write $x \xrightarrow{c} y$ (or, shortly $x \to y$ if there is no need to mention the ρ-admissible curve explicitly). The relation $\to$ on M is transitive, and is preserved by an anchored bundle morphism $(f, \overline{f})$, i.e. if $x' \to y'$, for $x', y' \in M'$, then $\overline{f}(x') \to \overline{f}(y')$. The set of points $y \in M$ such that $x \to y$ for some fixed x is denoted by R_x and is called *the set of reachable points from* x. Above we have proven that the base curve of a ρ-admissible curve is contained in a leaf L_x of the foliation on M, i.e. $R_x \subset L_x$. It is interesting to ask the question if every point in L_x can be reached from x following a ρ-admissible curve. In general this is not the case. However, if we consider the composition of ρ- and $(-\rho)$-admissible curves, then every point in L_x can indeed be reached from x.

Note: here and in the following, a minus sign in front of an element of a vector bundle or of a mapping with values in a vector bundle obviously applies to the fibre component.

Definition 2.2. Given an anchored bundle (ν, ρ). The *inverse anchored bundle* is defined by $(\nu, -\rho)$, where $-\rho : N \to TM : s \mapsto -\rho(s)$.

An anchored bundle (ν, ρ) is related to its inverse in the following way. Assume that $c : [a, b] \to N$ is a ρ-admissible curve taking x to y, i.e. $x \xrightarrow{c} y$. Then the curve $c^* : [a, b] \to N : t \mapsto c((b - t) + a)$ is $(-\rho)$-admissible and takes y to x. We shall call this curve *the* $(-\rho)$*-admissible curve associated with* c, or simply the *reverse of* c. Note that, using these notations, $(c^*)^* = c$. If we write the relation on M induced by the inverse anchored bundle as $\to_*$, we have the following equivalence:

$$x \xrightarrow{c} y \text{ iff } y \xrightarrow{c^*}_* x.$$

Note that the family of vector fields on M associated to the inverse anchored bundle structure equals $-\mathcal{D} = \{-\rho \circ \sigma \mid \sigma \in \Gamma(\nu)\}$ and, therefore, produces the same distribution D and the same foliation as $\mathcal{D}$. The set of reachable points from x induced by the inverse anchored bundle $(\nu, -\rho)$ is denoted by R_x^{-1}, i.e. $y \in R_x^{-1}$ if $x \to_* y$. We now consider the composition of ρ- and $(-\rho)$-admissible curves. Thus, assume that we have ℓ curves $c_i : [a_{i-1}, a_i] \to N$ for $i = 1, \dots, \ell$ such that $\widetilde{c}_{i-1}(a_{i-1}) = \widetilde{c}_i(a_{i-1})$ and such that c_i is either ρ-admissible or $(-\rho)$-admissible. The composition $c = c_\ell \cdot \ldots \cdot c_1$ of the curves c_i, defined as in Equation 2.2, is called a $\pm\rho$*-admissible curve.*

The projection $\widetilde{c}$ of a $\pm\rho$-admissible curve c onto M is a piecewise smooth curve which is called *the base curve of the* $\pm\rho$*-admissible curve.* If $\widetilde{c}(a_0) = x$ and $\widetilde{c}(a_\ell) = y$ we say that the $\pm\rho$-admissible curve takes x to y. Note that, in this case, the $\pm\rho$-admissible curve c^* defined by $c^* = (c_1)^* \cdot \ldots \cdot (c_\ell)^*$ takes y to x.

We thus obtain an alternative characterisation of the leaves of the foliation of M generated by the anchored bundle structure (ν, ρ).

Theorem 2.2. *We have that* $x \leftrightarrow y$*, or* $y \in L_x$*, iff there exists a* $\pm\rho$*-admissible curve taking* x *to* y*.*

Proof. The 'if'-part of the proof follows straightforwardly from Lemma 2.1. The 'only if'-part is proven by the following argument. Assume that $y \in L_x$ and consider a composite flow Φ of $\mathcal{X} = (X_\ell, \dots, X_1)$, with $X_i = \rho \circ \sigma_i$

and $\sigma_i \in \Gamma(\nu)$ (Φ is generated by $\mathcal{D}$) such that $\Phi_T(x) = y$. Consider the following curves:

$$c_i : [a_{i-1}, a_i] \to N : t \mapsto \sigma_i \circ \gamma|_{[a_{i-1},a_i]}, i = 1, \ldots, \ell,$$

where γ is the concatenation of integral curves associated with $\mathcal{X}$ and T through x (where we have used the notations from the preceding section). It is easily seen that c_i is ρ-admissible if $\mathrm{sgn}(t_i) > 0$, and $(-\rho)$-admissible if $\mathrm{sgn}(t_i) < 0$. If we put $c = c_\ell \cdot \ldots \cdot c_1$, then c takes x to y and is $\pm\rho$-admissible. □

The proof of the following theorem is a straightforward consequence of Theorem 2.2. Note that any anchored bundle morphism f between (ν', ρ') and (ν, ρ), which is fibred over $\overline{f} : M' \to M$, is also a morphism of the corresponding inverse anchored bundles, i.e. $f : (\nu', -\rho') \to (\nu, -\rho)$. This implies that, if $x' \to_* y'$ then $\overline{f}(x') \to_* \overline{f}(y')$, for $x', y' \in M'$.

Theorem 2.3. *Let f denote a morphism between (ν', ρ') and (ν, ρ), fibred over $\overline{f} : M' \to M$. Then $\overline{f}(L_{x'}) \subset L_{\overline{f}(x')}$. If (ν', ρ') is the pull-back bundle along $i : L_x \hookrightarrow M$ and $f = \pi_2$, then $i(L_x) = L_{i(x)}$.*

It is interesting to consider the special case of *linear* anchored bundles.

Theorem 2.4. *Let (ν, ρ) denote a linear anchored bundle on M and take any $x, y \in M$. Then $y \in L_x$ or $x \leftrightarrow y$ iff there exists a ρ-admissible curve that takes x to y, i.e. we have $R_x = L_x$.*

Proof. This theorem follows from the fact that, given a linear anchored bundle, then $x \to y$ iff $y \to x$. Indeed, assume that $c : [a, b] \to N$ is a ρ-admissible curve taking x to y. Then the curve $c^{-1} : [a, b] \to N : t \mapsto -c((b - t) + a)$ is also ρ-admissible and takes y to x. Note that $c^{-1} = -c^*$. The curve c^{-1} is called *the inverse of* c. In particular, the base curve of a $\pm\rho$-admissible curve is the base curve of a ρ-admissible curve on a linear anchored bundle, which proves the above theorem. □

Let (ν, ρ) denote a linear anchored bundle and let $c : [a, b] \to N$ denote a smooth ρ-admissible curve, with base curve $\widetilde{c}$. We now prove that any "reparametrisation" of $\widetilde{c}$ is again the base curve of a ρ-admissible curve. Assume that $\phi : [a, b] \to [a', b']$ is a diffeomorphism satisfying $\phi(a) = a'$ and $\phi(b) = b'$. Consider the following curve $c' : [a', b'] \to N$ defined by $c'(s) = \phi'^{-1}(s)c(\phi^{-1}(s))$, with $\phi'(s) = d\phi/ds(s) \in \mathbb{R}$. From elementary calculations it is easily seen that c' is ρ-admissible, and that its base curve equals $\widetilde{c}(\phi^{-1}(s))$, i.e. a reparametrisation of $\widetilde{c}$. From now on we agree on saying that every ρ-admissible curve in a linear anchored bundle is defined on the interval $[0, 1]$.

3 $\pm\rho$-admissible loops

Consider a point $x \in M$ and let $C(x, N)$ denote the set of all $\pm\rho$-admissible curves taking x to itself. Elements of $C(x, N)$ are called, with some abuse of terminology, *$\pm\rho$-admissible loops with base point x*. Indeed, in general a $\pm\rho$-admissible loop need not be continuous, nor closed.

Let $\pi_1(x, M)$ denote the fundamental group of M with reference point x and consider the map $C(x, N) \to \pi_1(x, M)$, associating to the base curve of a $\pm\rho$-admissible loop c, its homotopy class in $\pi_1(x, M)$, i.e. if $\widetilde{c}$ is the base curve of $c = c_\ell \cdot \ldots \cdot c_1 \in C(x, N)$, then c is mapped onto $[\widetilde{c}]$. It is easily seen that the image of $C(x, N)$ determines a subgroup of $\pi_1(x, M)$, which is denoted by $\pi_1^N(x, M)$. Indeed, assume that $c = c_\ell \cdot \ldots \cdot c_1$ and $d = d_{\ell'} \cdot \ldots \cdot d_1$ are elements of $C(x, N)$, with first homotopy classes $[\widetilde{c}]$ and $[\widetilde{d}]$, respectively. Then, the product $[\widetilde{c}] \cdot [\widetilde{d}]$ in $\pi_1(x, M)$ is the homotopy class of the base curve of $c_\ell \cdot \ldots \cdot c_1 \cdot d_{\ell'} \cdot \ldots \cdot d_1$. On the other hand, if $c = c_\ell \cdot \ldots \cdot c_1$ is a $\pm\rho$-admissible loop with base point x, then the curve $c^* = (c_1)^* \cdot \ldots \cdot (c_\ell)^*$ is also contained in $C(x, N)$, and the homotopy class of the base curve of c^* is precisely the inverse $[\widetilde{c}]^{-1}$ of $[\widetilde{c}]$. Therefore, the $\pm\rho$-admissible loops generate a subgroup of $\pi_1(x, M)$ which is denoted by $\pi_1^N(x, M)$. Note that, if (ν, ρ) is linear, then $\pi_1^N(x, M)$ is generated by the set of ρ-admissible loops with base point x, i.e. ρ-admissible curves taking x to itself. From Theorem 2.4, we know that any two points in L_x can be connected by the base curve of a $\pm\rho$-admissible curve. This implies, using standard arguments, that we can omit the reference point: we sometimes write $\pi_1^N(L_x, M)$ instead of $\pi_1^N(y, M)$ for any $y \in L_x$.

We now elaborate on how the above defined structures on anchored bundles behave under homomorphisms. From Section 2, we already know that $\pm\rho$-admissible curves are mapped into $\pm\rho$-admissible curves by an anchored bundle morphism $f : (\nu', \rho') \to (\nu, \rho)$. In particular, $\pm\rho$-admissible loops are preserved by such morphisms which, therefore, induce a group morphism between the corresponding subgroups of the first fundamental group of the base manifolds. More precisely, assume that f denotes a homomorphism between two anchored bundles (ν', ρ') and (ν, ρ), fibred over $\overline{f}$. Then, if $[\overline{f}]$

denotes the corresponding group morphism from $\pi_1(x', M')$ to $\pi_1(\overline{f}(x'), M)$, we have that $[\overline{f}]$ can be restricted to a morphism from $\pi_1^{N'}(L_{x'}, M')$ to $\pi_1^N(L_{\overline{f}(x')}, M)$.

Consider the pull-back setting under $i : L_x \hookrightarrow M$, and let $\pi_2 : N' = i^*N \to N$ denote the associated anchored bundle morphism. From the above, we know that $[i]$ maps the subgroup $\pi_1^{N'}(y, L_x)$ of $\pi_1(L_x)$ to the subgroup $\pi_1^N(y, M)$ of $\pi_1(y, M)$ (note that L_x is connected, allowing us to omit the reference point in the first homotopy group of L_x). We now prove that $[i] : \pi_1^{N'}(y, L_x) \to \pi_1^N(y, M)$ is onto. Consider an arbitrary element $[\widetilde{c}]$ of $\pi_1^N(y, M)$ associated with some $c \in C(y, N)$. From Lemma 2.1 we know that $\widetilde{c}$ is contained in the leaf $L_y = L_x$, which in turn implies that there exists a $\pm\rho$-admissible loop $c' \in C(y, N')$ such that $\pi_2 \circ c' = c$. In particular, we have that $[i]([\widetilde{c'}]) = [\widetilde{c}]$, and, hence, $[i]$ is onto, when restricted to $\pi_1^{N'}(y, L_x)$. Finally, again using the fact that any two points in L_x can be connected using a $\pm\rho$-admissible curve, we use the shorthand notation $\pi_1^{N'}(L_x)$ for $\pi_1^{N'}(y, L_x)$.

II
Generalised connections

The theory of connections undoubtedly constitutes one of the most beautiful and most important chapters of differential geometry, which has been widely explored in the literature (see e.g. [19, 22, 24, 39], and references therein). Besides its purely mathematical interest, connection theory has also become an indispensable tool in various branches of theoretical and mathematical physics.

Consider an arbitrary linear bundle $\pi : E \to M$, with total space E and base space M, and let $V\pi$ denote the canonical vertical distribution, i.e. the subbundle of TE consisting of all vectors tangent to the fibres of π. A connection on π (or E) is then given by a smooth distribution $H\pi$ on E, called a horizontal distribution, which is complementary to $V\pi$ and projects onto TM and which is invariant under the flow of the dilation vector field on E. This leads to a direct sum decomposition of TE, i.e. $TE = H\pi \oplus V\pi$. Note that there exist other ways of characterising a connection. For instance, a connection on π is sometimes defined as a global section of the first jet bundle $J^1\pi$ over E, or also as a splitting of the short exact sequence

$$0 \longrightarrow V\pi \xrightarrow{i} TE \xrightarrow{\widetilde{\pi}} \pi^*TM \longrightarrow 0,$$

i.e. a smooth map $h : \pi^*TM \to TE$ such that $\widetilde{\pi} \circ h$ is the identity map on the pull-back bundle π^*TM, where i denotes the natural injection and $\widetilde{\pi}$ the projection of TE onto π^*TM(cf. [19, 39, 49]). The main consequence of considering a connection on E is that, given any curve connecting two points in M, a parallel transport operator can be defined which determines a "connection" between the elements of the fibres of E over these base points.

In the literature one can find several generalisations of the concept of connection introduced above, obtained by relaxing the conditions on $H\pi$. First of all, we are thinking here of the so-called *partial connections*, where the horizontal distribution $H\pi$ does not determine a full complement of $V\pi$. More precisely, $H\pi$ has zero intersection with $V\pi$, but projects onto a subbundle of TM, rather than onto the full tangent bundle (see e.g. [21]). Of special interest are partial connections projecting onto an integrable subbundle of

TM, which play an important role in the study of the geometry of regular foliations.

Secondly, there also exists a notion of *pseudo-connection*, introduced under the name of quasi-connection in a paper by Y.C. Wong [59]. A fundamental role in the definition of a pseudo-connection on a manifold M is played by a type $(1,1)$-tensor field on M which simply becomes the unit tensor field in case of an ordinary connection. Pseudo-connections, and generalisations of it, have been studied by many authors (see [12] for a coordinate free definition of a pseudo-connection on a fibre bundle, and for more references to the subject).

The inspiration for the generalisation presented in this chapter mainly stems from some recent work by R.L. Fernandes on a notion of 'contravariant connection' in the framework of Poisson geometry (cf. [14]). Given a Poisson manifold (M, Λ), with Poisson tensor Λ, not necessarily of constant rank, and a principal G-bundle $\pi : P \to M$, a contravariant connection on π is defined as a G-invariant bundle map $h : \pi^*(T^*M) \to TP$ over the natural vector bundle morphism $\sharp_\Lambda : T^*M \to TM$ induced by the Poisson tensor. This concept of connection significantly deviates from the standard one, in that the 'horizontal' distribution $\mathrm{im}(h)$ may have nonzero intersection with the vertical subbundle $V\pi$ and, as for partial connections, projects onto a distribution of TM, namely $\sharp_\Lambda(T^*M)$. It is demonstrated in [14] that this definition of connection leads to familiar concepts such as parallelism, holonomy, curvature, etc..., and, therefore, plays an important role in the study of global aspects of Poisson manifolds. In a subsequent paper [15], Fernandes has extended this theory by replacing the cotangent bundle of a Poisson manifold by a Lie algebroid over an arbitrary manifold, and the $\sharp_\Lambda$-map of the Poisson tensor by the anchor map of the Lie algebroid structure. This resulted into a notion of Lie algebroid connection which, in particular, turns out to be appropriate for studying the geometry of singular foliations.

In this chapter we generalise the notion of connection by altering its definition as a horizontal lift. Roughly speaking, a generalised connection lifts elements of a anchored bundle on M towards vectors tangent to a bundle E over M. This agrees with our previous policy of considering an anchored bundle as an alternative to TM. In the first section we give a formal definition of a generalised connection, where we distinguish the case where the anchored bundle is not linear from the case where it is linear. In the former we will talk about a *lift over an anchor map* and in the latter about a *connection over an anchor map*. In subsequent sections, we examine the basic properties of these generalised connections and some additional structures that can be defined. In particular we show how the standard connections, as well as the notions of pseudo-connection, partial connection and Lie algebroid connection, fit into the general scheme presented below.

1 Lifts over an anchor map

Consider an anchored bundle (ν, ρ), with $\nu : N \to M$ and $\rho : N \to TM$. Let $\pi : E \to M$ be a vector bundle over M, with ℓ-dimensional fibres and with local bundle coordinates denoted by (q^i, y^A), where $i = 1, \ldots, n$ and $A = 1, \ldots, \ell$. We can then consider the pull-back bundle $\pi^* N = \{(e, s) \in E \times N \,|\, \pi(e) = \nu(s)\}$ which can be regarded as being fibred over E as well as over N, with natural projections given in coordinates by, respectively,

$$\widetilde{\pi}_1 : \pi^* N \longrightarrow E, \ (q^i, y^A, u^a) \longmapsto (q^i, y^A)$$

and

$$\widetilde{\pi}_2 : \pi^* N \longrightarrow N, \ (q^i, y^A, u^a) \longmapsto (q^i, u^a).$$

In particular, for each point $e \in E$, the fibre $(\widetilde{\pi}_1)^{-1}(e)$ can be identified with the typical fibre $N_{\pi(e)} = \nu^{-1}(\pi(e))$. Next, since E is assumed to be linear, one can consider the dilation vector field Δ on E, which locally reads

$$\Delta = y^A \frac{\partial}{\partial y^A}.$$

The flow of Δ is denoted by $\{\lambda_t\}$. Finally, let $\mathcal{D}$ denote the family of vector fields on M associated with the given anchored bundle (ν, ρ) (cf. Chapter I). We now have all ingredients at hand to introduce the main concept of the present chapter.

Definition 1.1. A *lift on π defined over the anchor map ρ*, shortly a *ρ-lift on π*, is a smooth bundle map $h : \pi^* N \to TE$ from $\widetilde{\pi}_1$ to τ_E over the identity on E such that, in addition, the following two conditions hold:

1. $T\pi \circ h = \rho \circ \widetilde{\pi}_2$,
2. $T\lambda_t(h(e, s)) = h(\lambda_t(e), s)$, for all $(e, s) \in \pi^* N$ and for all t.

From the definition of a ρ-lift, we have that the following two diagrams are commutative:

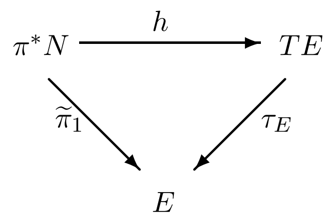

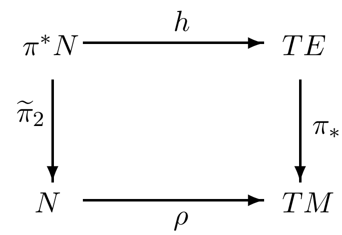

In the case where (ν, ρ) is a linear anchored bundle, the bundle $\widetilde{\pi}_1 : \pi^*N \to E$ inherits a linear structure. This leads us to the following definition.

Definition 1.2. A lift h over ρ on the bundle π is called a *connection on* π *over* ρ, shortly a ρ*-connection on* π if the map h is a linear bundle map from $\widetilde{\pi}_1$ to τ_E, i.e. conditions 1 and 2 from Definition 1.1 hold and, in addition,

3. $h(e, \lambda s + \lambda' s') = \lambda h(e, s) + \lambda' h(e, s')$, for $\lambda, \lambda' \in \mathbb{R}$ and $s, s' \in N_{\pi(e)}$.

Remark 1.1. It should be noted, for the sake of completeness, that one can also define a lift over ρ in the case where E is not necessarily a vector bundle. The definition should then be modified by leaving out condition 2. Since we do not use this generalisation in the remaining chapters, we refer to [5] for further details. However, we shall consider the specific case where E is a principal fibre bundle, since these structures are important when defining leafwise holonomy.

Let us now proceed with the case where h is a ρ-lift on a linear bundle π. In the remainder of this section, we first consider coordinate expressions of ρ-lifts and ρ-connections, and study how they behave under coordinate transformations. Next, we continue with defining additional structures associated with ρ-lifts.

In terms of the bundle coordinates introduced above, and taking into account the local form of ρ from Equation 2.1 in Chapter I, we can write h as

$$h(q^i, y^A, u^a) = (q^i, y^A, \gamma^i(q, u^a), \Gamma^A_B(q, u)y^B). \tag{1.1}$$

Note that the linearity in y^B of the functions Γ^A_B follows from Condition 2 in Definition 1.1. The functions Γ^A_B play the role similar to "connection coefficients", and will be called the *lift coefficients of the* ρ*-lift* h. Note that in the particular case where h is a ρ-connection, these functions are linear in u^a and can be written as

$$\Gamma^A_B(q, u) = \Gamma^A_{aB}(q)u^a.$$

In order to see how these functions behave under natural coordinate transformations, take any point $(e, s) \in \pi^*N$, with $\pi(e) = \nu(s) = x$, and consider a change of coordinates $(q^i, y^A, u^a) \to (\bar{q}^i, \bar{y}^A, \bar{u}^a)$ in a neighbourhood of (e, s), compatible with the underlying bundle structures:

$$\bar{q}^i = \bar{q}^i(q), \quad \bar{y}^A = \Xi^A_B(q) y^B, \quad \bar{u}^a = \bar{u}^a(q, u).$$

Note, first of all, that with respect to the bundle coordinates $(\bar{q}^i, \bar{u}^a)$ on N, the map ρ can be written as $(\bar{q}^i, \bar{u}^a) \mapsto (\bar{q}^i, \bar{\gamma}^i(\bar{q}, \bar{u}^a))$, with

$$\bar{\gamma}^i(\bar{q}(q), \bar{u}(q, u)) = \frac{\partial \bar{q}^i}{\partial q^j}(q) \gamma^j(q, u).$$

Next, representing the lift $h(e, s)$ over ρ in both coordinate systems by $(q^i, y^A, \gamma^i(q, u), \Gamma^A_B(q, u) y^B)$ and $(\bar{q}^i, \bar{y}^A, \bar{\gamma}^i(\bar{q}, \bar{u}), \bar{\Gamma}^A_B(\bar{x}, \bar{u}) \bar{y}^B)$, respectively, and taking into account the natural coordinate transformation on TE, induced by the transformation $(q^i, y^A) \to (\bar{q}^i, \bar{y}^A)$ on E, one finds after a tedious, but straightforward computation, the following transformation law for the lift coefficients associated to a general ρ-lift:

$$\bar{\Gamma}^A_B(\bar{q}(q), \bar{u}(q, u)) = \left(\frac{\partial \Xi^A_D}{\partial q^j}(q) \gamma^j(q, u) + \Xi^A_C(q) \Gamma^C_D(q, u) \right) (\Xi^{-1})^D_B.$$

In the specific case of a ρ-connection, the transformation rules become with $\bar{u}^a(q, u) = \Lambda^a_b(q) u^b$,

$$\bar{\Gamma}^A_{aB}(\bar{q}(q)) = \left(\frac{\partial \Xi^A_D}{\partial q^j}(q) \gamma^j_b(q) + \Xi^A_C(q) \Gamma^C_{bD}(q) \right) (\Xi^{-1})^D_B (\Lambda^{-1})^b_a\,.$$

Let us now return to the general case and let h denote a ρ-lift. It is easily seen from the definition of h that $(\widetilde{\pi}_1, h)$ determines an anchored bundle and that the bundle morphism $\widetilde{\pi}_2 : \pi^* N \to N$, which is fibred over $\pi : E \to M$, determines an anchored bundle morphism between $(\widetilde{\pi}_1, h)$ and (ν, ρ). Moreover, if h is a ρ-connection, we have that $(\widetilde{\pi}_1, h)$ is a linear anchored bundle and that $\widetilde{\pi}_2$ is a linear anchored bundle morphism. This is represented in the following diagram:

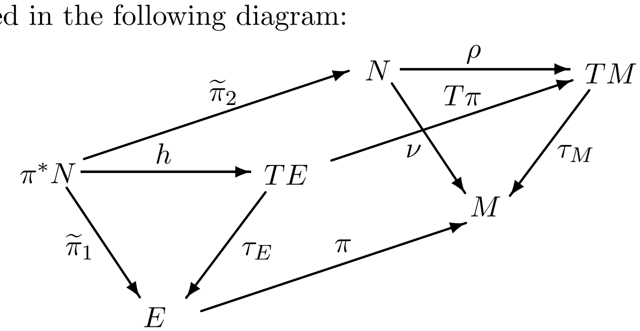

We will now apply the tools from Chapter I to the study of ρ-lifts. We first fix some notations. The everywhere defined family of vector fields generated by $(\widetilde{\pi}_1, h)$ on E is denoted by $\mathcal{Q}$, and correspondingly, the distribution on E generated by $\mathcal{Q}$ is denoted by Q. We refrain from calling Q a horizontal distribution, even in the case of a ρ-connection, since for arbitrary $e \in E$ it may be that Q_e has non-zero intersection with $V_e\pi$. Moreover, in general $\mathcal{Q}_e + V_e\pi \neq T_eE$, i.e. $\mathcal{Q}_e$ and $V_e\pi$ do not necessarily span the full tangent space T_eE. The smallest integrable distribution containing Q is, as usual, denoted by $\widetilde{Q}$. The leaf of $\widetilde{Q}$ through an arbitrary point $e \in E$ is written as $H(e)$. The ρ-lift h can be used to lift several kinds of objects living on the anchored bundle (ν, ρ) to the anchored bundle $(\widetilde{\pi}_1, h)$. For instance, given any (local) section σ of ν, we can define a mapping $\sigma^h : E \to TE$ by

$$\sigma^h(e) = h(e, \sigma(\pi(e))). \qquad (1.4)$$

It is seen that, by construction, σ^h is smooth and verifies $\tau_E(\sigma^h(e)) = e$, i.e. σ^h is a (local) vector field on E, called the *lift of the section σ with respect to h*, or simply *the lift of σ* if no confusion can arise. Let us denote the everywhere defined family of lifts of (local) sections of ν by $\mathcal{D}^h$.

Theorem 1.2. *Given any ρ-lift h, then the following properties hold:*

1. *the family $\mathcal{D}^h$ generates the distribution Q, and, hence, also the integrable distribution $\widetilde{Q}$;*

2. *any h-admissible curve is mapped by $\widetilde{\pi}_2$ onto a ρ-admissible curve;*

3. *given any ρ-admissible curve c taking x to y and a point $e \in E_x$, then there exists a unique h-admissible curve through e, projecting onto c by $\widetilde{\pi}_2$.*

Proof. Property 1 follows easily from the following observation. Take any element $w = h(e, s)$ and fix a section $\sigma \in \Gamma(\nu)$, such that $\sigma(\pi(e)) = s$. The lift σ^h of σ, is a vector field in Q, satisfying $\sigma^h(e) = w$. According to the definitions in Chapter I, this implies that $\mathcal{D}^h$ generates Q since Q is spanned by tangent vectors of the form $w = h(e, s)$.

We now prove 2. Let $c' = (\widetilde{c}', c)$ denote a h-admissible curve in π^*N, with base curve $\widetilde{c}'$ in E. From Chapter I (page 9) we know that $\widetilde{\pi}_2 \circ c' = c$ is a ρ-admissible curve, with base $\pi \circ \widetilde{c}'$.

The proof of Property 3 requires some more effort. Let $c : [a, b] \to N$ denote a ρ-admissible curve, with base curve $\widetilde{c}$ taking $x = \widetilde{c}(a)$ to $y = \widetilde{c}(b)$. First, consider a coordinate neighbourhood $U \subset M$ which is locally trivialising with respect to both vector bundle structures ν and π. Coordinates on $\nu^{-1}(U)$ and on $\pi^{-1}(U)$ are denoted by (q^i, u^a) and (q^i, y^A), respectively. Assume now that the image of the given ρ-admissible curve c is contained in $\nu^{-1}(U)$, with $c(a) = x = (q^i_0, u^a_0)$. Then, putting $c(t) = (q^i(t), u^a(t))$, the ρ-admissibility of c is expressed by the relation $\dot{q}^i(t) = \gamma^i(q^j(t), u^a(t))$ for all $t \in [a, b]$. Next, take any point $e_0 = (q^i_0, y^A_0) \in E_x$ and consider the following system of linear first-order ordinary differential equations with time-dependent coefficients:

$$\dot{y}^A = \Gamma^A_B(q^j(t), u^a(t))y^B,$$

It follows from the theory of linear differential equations that this system admits a unique solution $y^A(t)$ with $y^A(a) = y^A_0$ and which, moreover, is defined for all $t \in [a, b]$. The curve $c'(t) = (q^i(t), y^A(t), u^a(t))$ then clearly satisfies all the requirements of the proposition.

The proof for the more general case, with $\operatorname{im} c$ not necessarily contained in a single bundle chart, follows by taking a partition $a = t_0 < t_1 < \cdots < t_n = b$ of $[a, b]$ in such a way that the previous construction can be applied to the restriction of c to each subinterval $[t_i, t_{i+1}]$, and then gluing the results together. □

Let c denote a ρ-admissible curve and let c' denote the unique h-admissible curve through $e \in \pi^{-1}(\widetilde{c}(a))$ constructed in the above theorem. The base curve $\widetilde{c}'$ in E of c' is called the *lift of c through e with respect to h*, and is denoted by c^h_e. The induced map c^h from $\pi^{-1}(\widetilde{c}(a))$ to $\pi^{-1}(\widetilde{c}(b))$ is called the *h-displacement along* c and is defined by $c^h(e) = c^h_e(b)$. It is easily seen that c^h is linear on the fibres of E.

We now consider the inverse anchored bundles $(\nu, -\rho)$ and $(\widetilde{\pi}_1, -h)$. It is straightforward to check that $(-h)$ is a $(-\rho)$-lift, which allows us to apply the results obtained above to the $(-\rho)$-lift $-h$. Assume that $x, y \in M$ and that $x \xrightarrow{c} y$, with $c : [a, b] \to N$ a ρ-admissible curve. Then, from elementary calculations we find that $(c^*)^{-h}(t) = (c^h)^*(t)$, which in turn implies that $(c^*)^{-h}(c^h(e)) = e$. Thus, the h-displacement is an isomorphism on the fibres of E. Using the above theorem on h and $-h$, we have that any $\pm h$-admissible curve is projected onto a $\pm\rho$-admissible curve, and that any $\pm\rho$-admissible curve is the projection of a $\pm h$-admissible curve. The following corollary is then straightforward.

Corollary 1.3. $\pi(H(u)) = L_{\pi(u)}$.

For the sake of completeness, we mention some trivial properties of ρ-lifts of sections of the anchored bundle $\nu : N \to M$.

Proposition 1.4. *Given a ρ-lift h on π, we have for any $\sigma \in \Gamma(\nu)$ that:*

1. $\pi_* \circ \sigma^h = (\rho \circ \sigma) \circ \pi$, *i.e. the vector fields $\sigma^h \in \mathfrak{X}(E)$ and $\rho \circ \sigma \in \mathfrak{X}(M)$ are π-related;*
2. $[\Delta, \sigma^h] = 0$*;*
3. *if ν is a linear anchored bundle and if h is a ρ-connection then $(f\sigma)^h = (\pi^* f)\sigma^h$ and $(\sigma_1+\sigma_2)^h = \sigma_1^h+\sigma_2^h$ for all $\sigma_1, \sigma_2 \in \Gamma(\nu)$ and $f \in C^\infty(M)$.*

We now continue with further definitions.

Regarding TE as a bundle over TM, with projection $T\pi$, we can define the pull-back bundle $\rho^*TE = \{(s, w) \in N \times TE \mid \rho(s) = T\pi(w)\}$. Clearly, if $(s, w) \in \rho^*TE$, with $\tau_E(w) = e$, then $(e, s) \in \pi^*N$ and, given a ρ-lift h on π, one easily verifies that

$$T\pi(w - h(e, s)) = 0.$$

Hence, one can define a mapping $V : \rho^*TE \to V\pi$ by

$$V(s, w) = w - h(e, s) \quad \text{with } e = \tau_E(w). \tag{1.5}$$

Since $\pi : E \to M$ is a vector bundle, it is well-known that there exists a canonical isomorphism between $V\pi$ and the fibred product $E \times_M E$ ($\cong \pi^*E$). Denote by $p_2 : V\pi \cong E \times_M E \to E$ the projection onto the second factor, i.e. in coordinates: $p_2(q^i, y^A, 0, w^A) = (q^i, w^A)$. Given a ρ-lift h on π, we can define a mapping $K : \rho^*TE \to E$ by

$$K(s, w) = (p_2 \circ V)(s, w) \quad \text{for all} \quad (s, w) \in \rho^*TE. \tag{1.6}$$

In coordinates this reads

$$K(q^i, u^a, y^A, w^A) = (q^i, w^A - \Gamma^A_B(q, u)y^B). \tag{1.7}$$

The mapping K will be called the *connection map* (associated to the given ρ-lift), in analogy with the connection map associated to an ordinary connection on a vector bundle (see e.g. [57]).

Remark 1.5. T. Mestdag et al. in [42, 43] recognised that a ρ-connection can be defined alternatively as a splitting of the following short exact sequence of bundles over E:

$$0 \longrightarrow V_\rho E \longrightarrow \rho^*TE \longrightarrow \pi^*N \longrightarrow 0,$$

where $V_\rho E$ is the subbundle of ρ^*TE containing all $(s, w) \in \rho^*TP$ with $s = 0$.

2 Principal ρ-lifts

Consider the specific case where $\pi : P \to M$ is a principal fibre bundle, with structure group G. Let us denote the right action of G onto P by $R_g : P \to P$ with $g \in G$. The definition of a *principal* lift h over the anchor ρ is then given in the following way: condition 2 in Definition 1.1 should be replaced by $TR_g(h(u,s)) = h(ug,s)$, for all $(u,s) \in \pi^*N$.

Definition 2.1. A *principal lift on* π *defined over the anchor map* ρ, shortly a *principal* ρ*-lift on* π, is a smooth bundle map $h : \pi^*N \to TP$ from $\widetilde{\pi}_1$ to τ_P over the identity on P such that, in addition, the following two conditions hold:

1. $T\pi \circ h = \rho \circ \widetilde{\pi}_2$,
2. $TR_g(h(u,s)) = h(ug,s)$, for all $(u,s) \in \pi^*P$ and $g \in G$.

Similarly as in Definition 1.2, a principal ρ-lift h is called a *principal* ρ*-connection* if (ν, ρ) is a linear anchored bundle and if h is a linear bundle map from $\widetilde{\pi}_1$ to τ_P. In this section we will prove that any ρ-lift on a linear bundle determines a principal ρ-lift and vice-versa. This is a well known property from standard connection theory which can be extended to ρ-lifts. In view of these observations, we mainly focus on principal ρ-lifts. However, for making this work more accessible to the reader who is interested in the forthcoming chapters, we shall elaborate in this section on ρ-lifts in such a way that this section and Sections 3 and 4, dealing with the general theory on principal ρ-lifts, can be skipped.

Next, we recall some definitions and results on principal fibre bundles and principal connections from [22], since they will be used extensively in the following sections. Let $\pi : P \to M$ denote a principle fibre bundle with structure group G. The Lie algebra of G is denoted by $\mathfrak{g}$.

Consider for each $u \in P$ the smooth map $\sigma_u : G \to P$ defined by $\sigma_u(g) = ug$. Then we have $T_e\sigma_u : \mathfrak{g} \to V_u\pi$ and $TR_g \circ T_e\sigma_u = T_e\sigma_{ug} \circ Ad_{g^{-1}}$, where $Ad_h : \mathfrak{g} \to \mathfrak{g}$ denotes the adjoint action of G on its Lie algebra. Given any $A \in \mathfrak{g}$, let $\sigma(A)$ denote the fundamental vector field on P corresponding to A, defined by $\sigma(A)(u) = T_e\sigma_u(A)$. It is easily seen that $(R_g)_*\sigma(A) = \sigma(Ad_{g^{-1}}A)$.

A standard principal connection on P is defined by *a connection form* ω on P, i.e. ω is a $\mathfrak{g}$-valued one-form on P satisfying the following two conditions: (1) for any $A \in \mathfrak{g}$, $\omega(\sigma(A)) = A$, and (2) for any $g \in G$, $R_g^*\omega = Ad_{g^{-1}} \cdot \omega$. It is well known that ω is equivalently defined by a *horizontal lift* $h^\omega : P \times_M TM \to TP$, where h^ω and ω are related in the following way: $h^\omega(u,X) = \widetilde{X} - T_e\sigma_u(\omega(\widetilde{X}))$ for any $\widetilde{X} \in T_uP$ satisfying $T\pi(\widetilde{X}) = X$. From (2) it follows that h^ω is right invariant, i.e. $TR_g(h^\omega(u,X)) = h^\omega(ug,X)$ for $X \in T_{\pi(u)}M$ and $g \in G$ arbitrary. For the sake of completeness we mention that, equivalently, a principal connection can be defined by the right invariant distribution spanned by the image of h^ω, determining a direct sum decomposition of TP, i.e. putting $\operatorname{im} h^\omega = H\pi$, then $TP = H\pi \oplus V\pi$.

Before starting our study of principal ρ-lifts, we state the following lemma (see [22, p 69]).

Lemma 2.1. *Let G be a Lie group and $\mathfrak{g}$ its Lie algebra. Let $Y(t)$, for $a \leq t \leq b$, define a continuous curve in $\mathfrak{g}$. Then there exists a unique curve $g(t)$ of class C^1 in G such that $g(a) = e$ and $\dot{g}(t)g(t)^{-1} = Y(t)$ for $a \leq t \leq b$.*

Let us now return to the general treatment of principal ρ-lifts. Let (ν, ρ) denote an anchored bundle on M and let P denote a principal fibre bundle on M with structure group G.

Fix a standard principal connection ω on P. In the following we will use the connection form ω in order to obtain an alternative description for a principal ρ-lift h. This alternative description will allow us to derive easily some properties of lifts of ρ-admissible curves with respect to h (see below) using the theory of standard connections. Thus, let h be a given principal ρ-lift and consider the map $\chi : \pi^*N \to \mathfrak{g}$ defined by $\chi(u,s) = \omega(h(u,s))$ for any $(u,s) \in \pi^*N$. Note that the following relation holds $\chi(ug,s) = Ad_{g^{-1}} \cdot \chi(u,s)$ and that $h(u,s) = T_e\sigma_u(\chi(u,s)) + h^\omega(u,\rho(s))$. We shall sometimes refer to χ as *the coefficient of h with respect to ω*. The pair (ω,χ) uniquely determines the principal ρ-lift h in the following way. Given any connection form ω on P and a map $\chi : \pi^*N \to \mathfrak{g}$, such that χ transforms under the right action in the following way: $\chi(ug,s) = Ad_{g^{-1}} \cdot \chi(u,s)$, then the map $h : \pi^*N \to TP$, defined by $h(u,s) = h^\omega(u,s) + T_e\sigma_u(\chi(u,s))$, determines a principal ρ-lift. Note that the coefficient of h with respect to ω is precisely χ. We are now able to prove Theorem 1.2 in the case of principal fibre bundles.

Theorem 2.2. *Given any principal ρ-lift h, then the following properties hold:*

1. *the family $\mathcal{D}^h$ generates the distribution Q, and, hence, also the integrable distribution $\widetilde{Q}$;*

2. *any h-admissible curve is mapped by $\widetilde{\pi}_2$ onto a ρ-admissible curve;*
3. *given any ρ-admissible curve c taking x to y and a point $u \in P_x$, then there exists a unique h-admissible curve projecting onto c by $\widetilde{\pi}_2$ and such that its base curve in P passes through u.*

Proof. Properties 1 and 2 are trivial. In order to prove 3, we fix a principal connection ω and consider the coefficient χ of h with respect to ω.

We prove that, given any smooth ρ-admissible curve $c : [a, b] \to N$ with base curve $\widetilde{c}$, then there always exists a h-admissible curve whose base curve passes through $u \in P_{\widetilde{c}(a)}$ at $t = a$. We start by considering the horizontal lift $\widetilde{d^\omega}$ of $\widetilde{c}$ with respect to the principal connection ω, i.e. $\widetilde{d^\omega}(t)$ is the unique curve in P, projecting onto $\widetilde{c}$, satisfying $\dot{d}^\omega(t) = h^\omega(d^\omega(t), \dot{\widetilde{c}}(t))$ and $d^\omega(a) = u$. Let $g(t)$ denote the curve in G satisfying the equation $TR_{g(t)^{-1}}\dot{g}(t) = \chi(d^\omega(t), c(t))$ and $g(a) = e$, with e the unit element of G. From Lemma 2.1 we know that the curve $g(t)$ always exists and is unique. We now prove that the curve $(d(t), c(t))$ in $\pi^* N$ with $d(t) = d^\omega(t)g(t)$, is a h-admissible curve. Indeed, we find that:

$$\begin{aligned}
\dot{d}(t) &= TR_{g(t)}(\dot{d}^\omega(t)) + T_e\sigma_{d^\omega(t)}(\dot{g}(t)), \\
&= TR_{g(t)}\Big(h^\omega\big(d^\omega(t), \dot{\widetilde{c}}(t)\big)\Big) + T_e\sigma_{d(t)}\big(TL_{g^{-1}(t)} \cdot \dot{g}(t)\big), \\
&= h^\omega\big(d(t), \dot{\widetilde{c}}(t)\big) + T_e\sigma_{d(t)}\big(Ad_{g^{-1}(t)} \cdot \chi(d^\omega(t), c(t))\big).
\end{aligned}$$

From the definition of χ, it follows that the right hand side equals the desired vector $h(d(t), c(t))$. Clearly $(d(t), c(t))$ projects onto $c(t)$ and its base curve $d(t)$ in P passes through u at $t = a$. It is easily seen that this results also holds for a piecewise ρ-admissible curve.

It easily follows that $d(t)$ is uniquely determined by these conditions, since it satisfies a first order differential equation, i.e. $\dot{d}(t) = h(d(t), c(t))$, with given initial condition $d(a) = u$. It is also clear that the curve $d(t)$, constructed above, is independent of the choice of ω. □

2.1 h-Displacement and holonomy

Using the notations from the above theorem, we have that the component $d(t)$ of the h-admissible curve is uniquely determined by the ρ-admissible curve c and a point u in the fibre $P_{\widetilde{c}(a)}$. The curve $d(t)$ is called the *lift of* *the ρ-admissible curve c through u with respect to h* and from now on we write $c_u^h(t)$ to denote $d(t)$. Similar to standard connection theory, we call the map

$$c^h : \pi^{-1}(\widetilde{c}(a)) \to \pi^{-1}(\widetilde{c}(b)) : u \mapsto c_u^h(b),$$

the h-displacement along c. It is easily seen that c^h commutes with R_g for $g \in G$ arbitrary, i.e. $c^h(ug) = c^h(u)g$. Therefore, c^h determines a morphism on the fibres of P. The lift of a composition of ρ-admissible curves, in the sense of Chapter I, equals the composition of the corresponding h-admissible curves. Following the constructions described in the previous section, we can also consider the inverse anchored bundles $(\nu, -\rho)$ and $(\widetilde{\pi}_1, -h)$ of (ν, ρ) and $(\widetilde{\pi}_1, h)$, respectively. We have that $(c^*)^{(-h)} = (c^h)^{-1}$, i.e. c^h is invertible, that any $\pm h$-admissible curve projects onto a $\pm\rho$-admissible curve and that any $\pm\rho$-admissible curve is the projection of a $\pm h$-admissible curve. Hence, similar to Corollary 1.3, we obtain $\pi(H(u)) = L_{\pi(u)}$. This result is of great importance for the development of a notion of leafwise holonomy for principal ρ-lifts.

Definition 2.2. The set of all $g \in G$ such that $ug \in H(u)$, is called the *holonomy group with reference point u* and is denoted by $\Phi(u)$.

The fact that $\Phi(u)$ is a subgroup follows from the following lemma. First, note that, given any $g \in \Phi(u)$, there exists a $\pm h$-admissible curve taking u to $v = ug$, since $v \in H(u)$. This $\pm h$-admissible curve projects onto a $\pm\rho$-admissible loop with base point $\pi(u) = \pi(v) = x$. This implies that v can be reached from u by composing h-admissible curves and $(-h)$-admissible curves. Since a h-admissible curve is a lift of a ρ-admissible curve and a $(-h)$-admissible curve is a lift of a $(-\rho)$-admissible curve, we obtain that g is determined by composing a finite number of h-displacements along ρ-admissible curves and $(-h)$-displacements along $(-\rho)$-admissible curves. In particular, using the notations from Chapter I (page 12) we can define a map from the loop space $C(x, N)$ to $\Phi(u)$, which is onto. These observations are used in the proof of the following lemma.

Lemma 2.3. *$\Phi(u)$ is a subgroup of G.*

Proof. Given any two elements $g, g' \in \Phi(u)$ and let $ug = ((c^\ell)^{\pm h} \circ \cdots \circ (c^1)^{\pm h})(u)$, and $ug' = ((c^{\ell+\ell'})^{\pm h} \circ \cdots \circ (c^{\ell+1})^{\pm h})(u)$, for some $\pm\rho$-admissible curves c^i, $i = 1, \ldots, \ell + \ell'$, and where $(c^i)^{\pm h}$ stands for the h-displacement $(c^i)^h$ along c^i if c^i is ρ-admissible or for the $(-h)$-displacement $(c^i)^{(-h)}$ along c^i if c^i is $(-\rho)$-admissible.

Then $g'g^{-1} \in \Phi(u)$ since

$$ug'g^{-1} = \Big((c^{\ell+\ell'})^{\pm h} \circ \cdots \circ (c^{\ell+1})^{\pm h} \circ ((c^*)^1)^{\pm h} \circ \cdots \circ ((c^*)^\ell)^{\pm h}\Big)(u),$$

and, hence, $ug'g^{-1}$ belongs to $H(u)$. □

In the above proof, we used the fact that with any $\pm\rho$-admissible loop $c = c^\ell \cdot \ldots \cdot c^1$ with base point $x \in M$, we can associate a map of the fibre $\pi^{-1}(x)$ onto itself, namely $(c^\ell)^{\pm h} \circ \cdots \circ (c^1)^{\pm h}$, which commutes with the right action of G (i.e. such a map is called an automorphism of $\pi^{-1}(x)$). Indeed, for $u \in \pi^{-1}(x)$ and $g \in G$ arbitrary, we have

$$(c^\ell)^{\pm h} \circ \cdots \circ (c^1)^{\pm h}(ug) = \Big((c^\ell)^{\pm h} \circ \cdots \circ (c^1)^{\pm h}(u)\Big)g.$$

Using similar arguments as in the above proof, the set of all such automorphisms of the fibre $\pi^{-1}(x)$ forms a group, which is called the *holonomy group with reference point* x and denoted by $\Phi(x)$. We thus have the following commutative diagram:

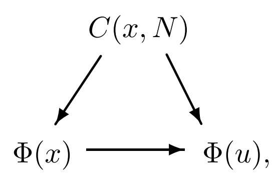

where the map $\Phi(x) \to \Phi(u)$ is defined as $(c^\ell)^{\pm h} \circ \cdots \circ (c^1)^{\pm h}$ is mapped onto the unique $g \in G$ such that

$$\Big((c^\ell)^{\pm h} \circ \cdots \circ (c^1)^{\pm h}(u)\Big) = ug.$$

Remark 2.4. In the specific case where h is a principal ρ-connection, the situation becomes more simple. In order to define the concept of holonomy groups it is sufficient to consider only ρ-admissible loops. Indeed, if c is $(-\rho)$-admissible, then $-c$ is ρ-admissible, and $c^{-h} = (-c)^h$. Moreover, we can consider reparametrisations of ρ-admissible curves and the notion of h-displacement does not depend on the parametrisation of c, in the following sense. Assume that $\phi : [a,b] \to [a',b']$ is a diffeomorphism with $\phi(a) = a'$ and $\phi(b) = b'$, then the curve $c' : [a',b'] \to N$, defined by

$$c'(s) = \frac{d\phi^{-1}}{ds}(s)c(\phi^{-1}(s)),$$

is ρ-admissible and, from elementary calculations, it follows that the h-displacements along c and c' are the same. Recall the definition of the inverse $c^{-1} = -c^*$ of a ρ-admissible curve c. The following identity holds $(c^{-1})^h = (c^h)^{-1}$.

The following properties which are well-known from the standard theory of holonomy, immediately carry over to the present framework.

Proposition 2.5. *(i) Given any $v \in H(u)$, then $\Phi(u) = \Phi(v)$. (ii) Given any $g \in G$, then $\Phi(ug) = I_{g^{-1}}(\Phi(u))$, where, I denotes the inner automorphism on G (i.e. for $h \in G$, $I_h : G \to G : h' \mapsto hh'h^{-1}$).*

Proof. By definition of $H(u)$, we have that $H(ug) = R_g(H(u))$. Indeed, $H(u)$ is the leaf of a foliation of a distribution generated by right invariant vector fields (cf. Theorem 2.2). Thus, if $h \in \Phi(u)$, then $h^{-1} \in \Phi(u)$ and $uh^{-1} \in H(u)$, or $H(uh^{-1}) = H(u) = H(v)$. Acting on the right by h, we obtain $H(u) = H(vh)$. And since $H(u) = H(v)$, we have $h \in \Phi(v)$, proving (i). Since $H(ug) = H(u)g$, we have that, for any $h \in \Phi(u)$, $H(uhg) = H(ug)$. Thus $g^{-1}hg \in \Phi(ug)$, proving (ii). □

2.2 Local lift functions

In this section we study ρ-lifts in terms of local coordinate expressions. Choose an open covering $\{U_\alpha\}$ of M such that each $\pi^{-1}(U_\alpha)$ is provided with a local trivialisation ψ_α, i.e.

$$\psi_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times G : u \mapsto (\pi(u), \phi_\alpha(u))$$

is a diffeomorphism satisfying $\psi_\alpha(ug) = (\pi(u), \phi_\alpha(u)g)$ for all $g \in G$. Consider the canonically defined flat principal connection θ_α on $U_\alpha \times G$, i.e. θ_α is the pull-back by ϕ_α of the canonical left-invariant $\mathfrak{g}$-valued one-form θ on G (see [22]). Consider the lift coefficient χ_α of h with respect to θ_α, then, with $\sigma_\alpha : U_\alpha \to \pi^{-1}(U_\alpha) : x \to \psi_\alpha^{-1}(x, e)$, we define *the local lift coefficients* $\chi^*_\alpha : U_\alpha \times_{U_\alpha} N \to \mathfrak{g}$ by:

$$\chi^*_\alpha(x, s) = \chi_\alpha(\sigma_\alpha(x), s).$$

Assume that $x \in U_\alpha \cap U_\beta \neq \emptyset$ and let $\psi_{\alpha\beta} : U_\alpha \cap U_\beta \to G$ denote the *transition functions*, defined by $(x, \psi_{\alpha\beta}(x)) = \psi_\alpha(\psi_\beta^{-1}(x, e))$. The pull-back of θ by $\psi_{\alpha\beta}$ and ρ, respectively, is denoted by $\theta_{\alpha\beta}$ and equals $\theta_{\alpha\beta}(x, s) = (\psi^*_{\alpha\beta}\theta)(x)(\rho(s))$, for arbitrary $s \in N_x$. Note that $\psi^{-1}_{\alpha\beta}(x)$ is a shorthand notation for $(\psi_{\alpha\beta}(x))^{-1}$.

Lemma 2.6. *The following relation holds:*

$$\chi_\beta^*(x,s) = Ad_{\psi_{\alpha\beta}^{-1}(x)} \cdot \chi_\alpha^*(x,s) - \theta_{\alpha\beta}(x,s).$$

Proof. Let us denote the left multiplication in G by $L_g : G \to G$, with $g \in G$ and let u denote an arbitrary point in $\pi^{-1}(U_\alpha \cap U_\beta)$. Then, from $\phi_\beta(u) = L_{\psi_{\alpha\beta}^{-1}(\pi(u))}\phi_\alpha(u)$, we derive that:

$$T_u\phi_\beta = TL_{\psi_{\alpha\beta}^{-1}(\pi(u))} \circ T_u\phi_\alpha + TR_{\phi_\alpha(u)} \circ T_{\pi(u)}\psi_{\alpha\beta}^{-1} \circ T_u\pi.$$

Using the above identity, we now derive the required equality. Assume in the following that $u = \sigma_\beta(x)$ and $u' = \sigma_\alpha(x)$. Taking into account that, by definition, $\theta_\beta = \phi_\beta^*\theta$, we obtain:

$$\begin{aligned}
\chi_\beta^*(x,s) &= \theta_\beta(u)\big(h(u,s)\big) \\
&= \theta(e)\Big(T_u\phi_\beta\big(h(u,s)\big)\Big) \\
&= \theta\big(\psi_{\alpha\beta}(x)\big)\Big(T_u\phi_\alpha\big(h(u,s)\big)\Big) \\
&\quad + \theta(e)\Big(TR_{\phi_\alpha(u)} \circ T_x\psi_{\alpha\beta}^{-1}\big(\rho(s)\big)\Big) \\
&= \theta_\alpha(u'\psi_{\alpha\beta}(x))\,\Big(TR_{\psi_{\alpha\beta}(x)}(h(u',s))\Big) \\
&\quad - \theta(\psi_{\alpha\beta}(x))\,(T_x\psi_{\alpha\beta}(\rho(s)))\,,
\end{aligned}$$

where we used the fact that $u = u'\psi_{\alpha\beta}(x)$, $\phi_\alpha(u) = \psi_{\alpha\beta}(x)$ and, from $\psi_{\alpha\beta}^{-1}(x)\psi_{\alpha\beta}(x) = e$:

$$TL_{\psi_{\alpha\beta}^{-1}(x)} \circ T_x\psi_{\alpha\beta} = -TR_{\psi_{\alpha\beta}(x)} \circ T_x\psi_{\alpha\beta}^{-1}. \quad \square$$

The computations in the proof of Lemma 2.6 can be used in order to prove that any family of $\mathfrak{g}$-valued functions χ_α^* transforming in the above way, determine a principal ρ-lift h. This result will be of importance in the following section, where we will relate ρ-lifts and principal ρ-lifts.

2.3 Associated bundles

First, we recall the definition of the notion of an associated bundle of a principle bundle P. Let V be a manifold on which the structure group G of

P acts on the left: $(g, \xi) \mapsto g\xi \in V$ (not necessarily free and proper). The *bundle* $\pi_E : E \to M$ *associated with* P and V is defined as the quotient space of $P \times V$ under the right action of G defined by $(g, (u, \xi)) \mapsto (ug, g^{-1}\xi)$. The elements of E are denoted by $u\xi$, where $u\xi$ represents the orbit of $(u, \xi) \in P \times V$. The projection $\pi_E(u\xi) = \pi(u)$ is thus well defined. The bundle E is given a differentiable structure using the local trivialisations of P (see [22]). The natural projection from $P \times V$ onto E is denoted by p_E and equals $p_E(u, \xi) = u\xi$. Every element $u \in P$ induces a bijective map from V to $E_{\pi(u)}$ which is also denoted by u, i.e. $\xi \in V \mapsto u\xi$. Therefore, if V has an algebraic structure which is preserved by the action of G, then the fibres of E inherit this structure by demanding that $u : V \to E$ is an isomorphism with respect to this algebraic structure. The following relation $(ug)\xi = u(g\xi)$ holds for any $g \in G$. The map taking a fibre E_x to another fibre E_y defined by $v \circ u^{-1}$ is an isomorphism, where $u \in \pi^{-1}(x)$ and $v \in \pi^{-1}(y)$. In particular, we have that, if $x = y$, the isomorphisms of the form $u \circ g \circ u^{-1}$, with u an arbitrary fixed point in $\pi^{-1}(x)$, are automorphisms with respect to the induced algebraic structure. In this way, we can consider G as a group of automorphisms of E_x.

As an example, we consider an arbitrary ℓ-dimensional linear bundle $\pi_E : E \to M$, and let $GL(\ell; \mathbb{R})$ denote the general linear group consisting of all non-singular $\ell \times \ell$ real matrices. Consider any point $x \in M$, then $u = (e_1, \dots, e_\ell)$, with $e_A \in E_x$, is said to be *a frame of* E *at* x if it determines a basis of E_x. It can be proven that the set $FR(E)$ of all frames of E is a principal fibre bundle with structure group $GL(\ell; \mathbb{R})$, called the *frame bundle*. The right action of $GL(\ell; \mathbb{R})$ on $FR(E)$ is then defined as follows: given $u = (e_1, \dots, e_\ell)$ and $S \in GL(\ell; \mathbb{R})$ arbitrary, then $uS = (e'_1, \dots, e'_\ell)$ with $e'_B = S^A_B e_A$. Moreover, the initial linear bundle E can be retrieved as the bundle associated with $FR(E)$ and $\mathbb{R}^\ell$, where the action of $GL(\ell; \mathbb{R})$ on $\mathbb{R}^\ell$ is the standard left action. The natural projection $p_E : FR(E) \to E$ then reads, for $u = (e_1, \dots, e_\ell) \in FR(E)$ and $\xi = (\lambda^1, \dots, \lambda^\ell) \in \mathbb{R}^\ell$ arbitrary, $p_E(u, \xi) = e_A \lambda^A \in E_{\pi(u)}$.

From now on, we assume that $\pi_E : E \to M$ denotes a bundle associated with a principal G-bundle P and the vector space $V = \mathbb{R}^\ell$. The action of G on V is assumed linear, i.e. it determines a representation $\rho : G \to GL(\ell, \mathbb{R})$ of G into $GL(\ell, \mathbb{R})$. This implies that E can be given the structure of a linear bundle (see above): every map $u : V \to E_{\pi(u)} : \xi \mapsto u\xi$ is a linear isomorphism.

Assume that we have fixed a principal ρ-lift h on P. We shall now construct a ρ-lift h_E on the bundle E associated with the principal ρ-lift h.

Given any $\xi \in V$, then the map $\overline{f}_\xi : P \to E$ defined by $\overline{f}_\xi(u) = u\xi$ induces a map f_ξ from π^*N to π_E^*N as follows: $(u,s) \in \pi^*N \mapsto (u\xi, s) \in \pi_E^*N$. The ρ-lift h_E can now be defined as the unique anchor map from π_E^*N to TE, such that, for all ξ, f_ξ is a anchored bundle morphism between the anchored bundles $(\pi^*N \to E, h)$ and $(\pi_E^*N \to E, h_E)$. More specifically, let $u\xi$ denote an arbitrary element of E. Then put $h_E(u\xi, s) = T\overline{f}_\xi(h(u,s))$. It is easily seen that h_E is independent of the choice of u and ξ representing the point $u\xi$ and that f_ξ satisfies the proposed property of being an anchored bundle morphism (see also the following section, were we consider mappings between generalised connections). Moreover, given any $\lambda \in \mathbb{R}$, then $f_{\lambda\xi} = \lambda f_\xi$, implying that h_E is a ρ-lift. We can apply the theory from Chapter I: any h-admissible curve $(c_u^h(t), c(t))$ is mapped onto the h_E admissible curve $(c_u^h(t)\xi, c(t))$. It is readily seen that $c_u^h(t)\xi = c_{u\xi}^{h_E}(t)$. On the other hand, assume that $(c_{u\xi}^{h_E}(t), c(t))$ is h_E-admissible. Consider the lift of c through u with respect to h, then $\overline{f}_\xi(c_u^h(t)) = c_{u\xi}^{h_E}(t)$. This is also valid for the inverted anchored bundles, implying that $\overline{f}_\xi(H(u)) = H(u\xi)$.

The above procedure describes how to construct out of a principal ρ-lift, a ρ-lift on associated linear bundles. We now study the converse, i.e. starting from a ρ-lift h_E on a linear bundle E, we shall define a principal ρ-lift on the bundle of frames of E. Assume that we have fixed a ρ-lift h_E on E. Using the local coordinate expressions from Section 1 we know that, on each adapted coordinate system (q^i, u^a) of ν, the ρ-lift h_E on E determines a set of ℓ^2 functions Γ_B^A, transforming in a specific way under coordinate transformations. Using these local functions, we shall define a local $\mathfrak{gl}(\ell,\mathbb{R})$-valued function, that satisfies the transformation rules, expressed in Lemma 2.6. This family of $\mathfrak{gl}(\ell,\mathbb{R})$-valued functions then determines a principal ρ-lift on $FR(E)$, what we wanted to proof. We start with fixing some notations.

The Lie algebra $\mathfrak{gl}(\ell,\mathbb{R})$ of $GL(\ell,\mathbb{R})$ is the set of all real $\ell \times \ell$ matrices. Consider the so-called Weyl basis (E_A^B) of $\mathfrak{gl}(\ell,\mathbb{R})$, i.e. E_A^B is the matrix such that the entry on the A-th row and B-th column is 1 and all other entries equal 0. The natural coordinate functions on $GL(\ell,\mathbb{R})$ are denoted by s_B^A and let t_B^A be the coordinate representation of the inverse map on $GL(\ell,\mathbb{R})$. The canonical one-form θ on $GL(\ell,\mathbb{R})$ then equals:

$$\theta = \sum_{A,B,C} t_B^A ds_C^B E_A^C.$$

Next, assume that $\{V_\alpha\}$ denotes an open covering of π_E^*N, such that every V_α is equipped with a bundle adapted coordinate system (the projection of V_α

on M is denoted by U_α). Assume that $V_\alpha \cap V_\beta \neq \emptyset$ and denote the coordinate functions on V_α and V_β by (q^i, u^a, y^A) and $(\overline{q}^i, \overline{u}^a, \overline{y}^A)$, respectively. Recall the notations from Section 1, i.e. we have $\overline{y}^A = \Xi^A_B(q) y^B$. The transition function $\psi_{\alpha\beta}$, defined on $U_\alpha \cap U_\beta$ and taking values in $GL(\ell, \mathbb{R})$, then equals $(\Xi^{-1})^B_A(q)$ when expressed in the natural coordinates on $GL(\ell, \mathbb{R})$.

The pull-back of θ by $\psi_{\alpha\beta}$ is now easily computed and equals:

$$\begin{aligned} \psi^*_{\alpha\beta}\theta &= \Xi^A_B d\left((\Xi^{-1})^B_C\right) E^C_A \\ &= \Xi^A_B \frac{\partial (\Xi^{-1})^B_C}{\partial q^i} dq^i E^C_A \\ &= -(\Xi^{-1})^B_C \frac{\partial \Xi^A_B}{\partial q^i} dq^i E^C_A. \end{aligned}$$

Substituting this in the definition of $\theta_{\alpha\beta} : \nu^{-1}(U_\alpha \cap U_\beta) \to \mathfrak{gl}(\ell, \mathbb{R})$, we obtain:

$$\theta_{\alpha\beta}(q^i, s) = -(\Xi^{-1})^B_C(q^i) \frac{\partial \Xi^A_B}{\partial q^j}(q^i) \gamma^j(q^i, s) E^C_A,$$

where s is in the fibre of N over the point with coordinates q^i. We define the $\mathfrak{gl}(\ell, \mathbb{R})$-valued functions χ^*_α on $\nu^{-1}(U_\alpha)$ and χ^*_β on $\nu^{-1}(U_\beta)$ by:

$$\chi^*_\alpha(q^i, s) = \Gamma^B_A(q^i, s) E^A_B \text{ and } \chi^*_\beta(\overline{q}^i, s) = \overline{\Gamma}^B_A(\overline{q}^i, s) E^A_B,$$

respectively. Using the above definitions and the local expression for $\theta_{\alpha\beta}$, the right hand side of the transformation rule from Lemma 2.6 becomes, with a slight abuse of notations:

$$\left(\Xi^A_C (\Xi^{-1})^D_B \Gamma^C_D + (\Xi^{-1})^C_B \frac{\partial \Xi^A_C}{\partial q^i} \gamma^i\right) E^B_A.$$

In Section 1, we have proven that the above expression precisely equals $\overline{\Gamma}^A_B E^B_A$, where $\overline{\Gamma}^A_B$ is the lift coefficient of h_E on the coordinate chart V_β.

We now show that the above obtained principal ρ-lift h generates the ρ-lift h_E. If we express $\overline{f}_\xi$ locally as $(q^i, s^A_B) \mapsto (q^i, y^A = s^A_B \lambda^B)$, where s^A_B represent the fibre coordinates on the frame bundle, then the image of (q^i, s^A_B, X^i, X^A_B) under $T\overline{f}_\xi$ equals $(q^i, s^A_B \lambda^B, X^i, X^A_B \lambda^B)$. In this coordinate system, the principal ρ-lift h equals:

$$(q^i, s^A_B, u^a) \mapsto (q^i, s^A_B, \gamma^i(q, u), \Gamma^A_C(q^i, u^a) s^C_B)$$

Therefore, using the above expressions, we see that $T\overline{f}_\xi(h(u,s)) = h_E(u\xi, s)$, for $(u,s) \in \pi^*N$.

Using the above correspondence, we can define the notion of holonomy for ρ-lifts. Let h_E be a ρ-lift on the linear bundle E. The holonomy group $\Phi(x)$ of h on $FR(E)$ can be identified with a subgroup of the linear automorphisms of the fibre E_x. In fact, any automorphism of the fibre $FR(E)_x$ determines a linear automorphism of E_x and vice versa. Indeed, let $\Lambda : FR(E)_x \to FR(E)_x$ be such that $\Lambda(ug) = \Lambda(u)g$, then $\overline{\Lambda} : E_x \to E_x$ defined by $\overline{\Lambda}(u\xi) := \Lambda(u)\xi$ is well defined and $\overline{\Lambda}$ is linear. On the other hand, any automorphism $\overline{\Lambda}$ of E_x determines an automorphism of the linear frames of E_x. Indeed, let $u = (e_1, \ldots, e_\ell) \in FR(E)$, then

$$\Lambda(u) := (\overline{\Lambda}(e_1), \ldots, \overline{\Lambda}(e_\ell)) \in FR(E)$$

and if $S \in GL(\ell, \mathbb{R})$, then $\overline{\Lambda}(e_A S^A_B) = \overline{\Lambda}(e_A) S^A_B$, implying that $\Lambda(R_S u) = R_S \Lambda(u)$.

3 Mappings between generalised connections

We first fix some notations. Let (ν', ρ') and (ν, ρ) denote anchored bundles with base manifolds M' and M, respectively, and consider an anchored bundle morphism $f : N' \to N$ between (ν', ρ') and (ν, ρ), which is fibred over $\overline{f} : M' \to M$. Assume that $\pi' : P' \to M'$ and $\pi : P \to M$ are principal fibre bundles with structure groups G' and G, respectively. Furthermore, we assume that a principal fibre bundle morphism $(F, \overline{F})$ between P' and P is given, i.e. $F : P' \to P$ is a bundle map and $\overline{F} : G' \to G$ is a group morphism between G' and G such that for all $u' \in P'$ and $g' \in G'$, we have $F(u'g') = F(u')\overline{F}(g')$. We assume in the following that the map F is also fibred over $\overline{f} : M' \to M$.

The principal fibre bundle morphism F is called a *morphism between the principal ρ'-lift h' and the principal ρ-lift h* if the map (F, f), defined by $(F, f) : (\pi')^*N' \to \pi^*N : (u', s') \mapsto (F(u'), f(s'))$, is an anchored bundle morphism between $(\widetilde{\pi}'_1, h')$ and $(\widetilde{\pi}_1, h)$. More precisely, we should have that:

$$TF\left(h'(u', s')\right) = h(F(u'), f(s')).$$

In the remainder of this section we are only concerned with principal lifts. However, for the sake of completeness, we will mention here a similar definition for ρ-lifts defined on linear bundles. Assume that E' is a linear bundle over M' and that E is a linear bundle over M, and let h' denote a ρ'-lift on E' and h a ρ-lift on E. If $F : E' \to E$ is a linear bundle morphism, then we say that it determines a *map from the ρ'-lift h' to the ρ-lift h* if $(F, f) : (\pi')^* N' \to \pi^* N : (e', s') \mapsto (F(e'), f(s'))$ is an anchored bundle morphism between $(\widetilde{\pi}'_1, h')$ and $(\widetilde{\pi}_1, h)$. This definition will be used later in Section 6.

Let us now continue with the principal lifts.

Theorem 3.1. *Assume that $f : N' \to N$ is an anchored bundle isomorphism, fibred over the diffeomorphism $\overline{f} : M' \to M$, and that F is a principal fibre bundle morphism from P' to P, also fibred over $\overline{f}$. Let h' be a principal ρ'-lift on P'. Then, there exists a unique principal ρ-lift h such that F is a morphism between h' and h. The holonomy group $\Phi(u')$ corresponding to h' is mapped by $\overline{F}$ onto the holonomy group $\Phi(F(u'))$ corresponding to h.*

Proof. Let u denote an arbitrary point of P, with $\pi(u) = x$. Then fix an element u' in $P'_{\overline{f}^{-1}(x)}$ and an element g in G such that $F(u') = ug$. Define $h(u, s) \in T_u P$, for any $s \in N_{\overline{f}^{-1}(x)}$, by

$$h(u, s) = TR_{g^{-1}} \left(T_{u'} F(h'(u', f^{-1}(s))) \right).$$

This tangent vector in $T_u P$ is well defined, in the sense that it does not depend on the choice of u', since for any other element $v' = u'g'$, we have that $F(v') = F(u')\overline{F}(g') = uh$, with $h = g\overline{F}(g')$. This implies that

$$\begin{aligned} h(u, s) &= TR_{h^{-1}} \left(T_{v'} F(h'(v', f^{-1}(s))) \right) \\ &= TR_{h^{-1}} \left(T_{v'} F(TR_{g'} h'(u', f^{-1}(s))) \right) \\ &= TR_{g^{-1}} TR_{\overline{F}(g'^{-1})} TR_{\overline{F}(g')} \left(T_{u'} F(u', f^{-1}(s)) \right) \\ &= TR_{g^{-1}} \left(T_{u'} F(h'(u', f^{-1}(s))) \right). \end{aligned}$$

In this way, we have constructed a mapping $h : \pi^* N \to TP$, which is clearly right invariant and, by definition, it follows that (F, f) is an anchored bundle morphism between $(\widetilde{\pi}'_1, h')$ and $(\widetilde{\pi}_1, h)$. From the fact that f^{-1} maps any $\pm\rho$-admissible curve onto a $\pm\rho'$-admissible curve, we have that $H(u')$ is mapped by F onto $H(F(u'))$, which concludes the proof. □

In the specific case when P' is a reduced subbundle of P, i.e. F is an injective immersion and $\overline{F}$ is a monomorphism, we say that h is *reducible to a principal*

ρ-lift on P'. This is important for our treatment of holonomy, where we will prove a generalisation of the Reduction Theorem, which, roughly speaking, says that the leaf $H(u)$ in P is a reduced subbundle of P with structure group the holonomy group $\Phi(u)$ and that h is reducible to $H(u)$.

For the following theorem we take for (ν', ρ') the pull-back anchored bundle of (ν, ρ) under $i : L_x \hookrightarrow M$, with L_x the leaf of the integrable distribution $\widetilde{D}$ generated by $\mathcal{D}$ through some $x \in M$. Let $P' = i^*P \to L_x$ denote the pull-back principal bundle of P along i and let $F : P' \to P : (y, u) \mapsto u$ denote the projection onto the second factor. Note that, by definition, the structure group of P' equals G and that F is an injective immersion.

Theorem 3.2. *There exists a unique principal ρ'-lift h' on P' such that F is a morphism between h' and h. Moreover, $F(H(u')) = H(F(u'))$ and, therefore $\Phi(u') = \Phi(F(u'))$.*

Proof. Since F is an injective immersion, we know from Chapter I that a unique anchor map h' on P' can be defined such that F is an anchored bundle morphism between $(\widetilde{\pi}'_1, h')$ and $(\widetilde{\pi}_1, h)$. It is trivial to check that h' satisfies the "right invariance" condition making it into a principal ρ-lift.

The fact that the foliations on P' and P induced by the ρ'-lift h' and the ρ-lift h, respectively, are F-related, follows from the fact that $\pm\rho$-admissible curves are in one-to-one correspondence with $\pm\rho'$-admissible curves. □

In the following section we prove that the holonomy group $\Phi(u)$ of a principal ρ-lift is a Lie subgroup of G. In view of the above theorem, we will assume, without loss of generality in view of the result we wish to obtain, that we are working with the ρ'-lift h' on the pull-back bundle i^*P, with $i : L_x \hookrightarrow M$.

4 Leafwise Holonomy of a principal ρ-lift

In view of the above comment, we restrict ourselves to the case where the anchored bundle (ν, ρ) on M satisfies the additional conditions that M is a connected manifold and that $\widetilde{D} = TM$ (this situation occurs when we are working on the pull-back anchored bundle). The main consequence of these assumptions is that one can prove that the distribution $\widetilde{Q}$ generated by a principal ρ-lift h on a principal bundle P is regular, i.e. has constant rank.

Theorem 4.1. *If $\widetilde{D} = M$ and if M is a connected manifold, then $\widetilde{Q}$ is a regular integrable distribution.*

Proof. We have to show that, given two arbitrary points u, v in P, then $\dim \widetilde{Q}_u = \dim \widetilde{Q}_v$. Let $x = \pi(u)$ and $y = \pi(v)$. Then, since the foliation of $\widetilde{D}$ on M is trivial, there exists a composite flow Φ associated with an ordered set of vector fields $(\rho \circ \sigma^\ell, \dots, \rho \circ \sigma^1)$ on P, belonging to $\mathcal{D}$ and a composite flow parameter T such that $\Phi_T(x) = y$ (cf. Theorem 1.1). Consider the vector fields $(\sigma^i)^h$ in $\mathcal{Q}$. The flow of $(\sigma^i)^h$ and the flow of $\rho \circ \sigma^i$ are π-related by definition and, therefore, if Φ^h is the composite flow of $((\sigma^\ell)^h, \dots, (\sigma^1)^h)$, we have $\pi(\Phi^h_T(u)) = y$. Hence, there exists a $g \in G$ such that $\Phi^h_T(u)g = v$. By definition of $\widetilde{Q}$ we have $T\Phi^h_T(\widetilde{Q}_u) = \widetilde{Q}_{\Phi^h_T(u)}$. On the other hand since $\mathcal{D}^h$ consists of right invariant vector fields and generates $\widetilde{Q}$, we have $TR_h(\widetilde{Q}_w) = \widetilde{Q}_{wh}$ for any $w \in P$ and $h \in G$. Thus, we obtain that $TR_g \circ T\Phi^h_T$ is an isomorphism from $\widetilde{Q}_u$ to $\widetilde{Q}_v$ and, in particular, $\dim \widetilde{Q}_u = \dim \widetilde{Q}_v$. □

For an arbitrary point $u \in P$ and consider the linear subspace $\mathfrak{g}(u)$ of $\mathfrak{g}$ defined by $T_e\sigma_u(\mathfrak{g}(u)) = V_u\pi \cap \widetilde{Q}_u$.

Proposition 4.2. *Let g denote an arbitrary element in G. We have: (i) $\mathfrak{g}(u) = \mathfrak{g}(v)$ for any $v \in H(u)$, (ii) $Ad_{g^{-1}}(\mathfrak{g}(u)) = \mathfrak{g}(ug)$ and (iii) $\mathfrak{g}(u)$ is a Lie subalgebra of $\mathfrak{g}$.*

Proof. (i) follows from the fact that $V\pi$ and $\widetilde{Q}$ are both invariant under the tangent map of a composite flow induced by vector fields in $\mathcal{D}^h$. (ii) follows from $TR_g \circ T_e\sigma_u = T_e\sigma_{ug} \circ Ad_{g^{-1}}$, $TR_g(V_u\pi) = V_{ug}\pi$ and $TR_g(\widetilde{Q}_u) = \widetilde{Q}_{ug}$. (iii) follows from $[\sigma(A), \sigma(B)] = \sigma([A, B])$, for $A, B \in \mathfrak{g}$ and the fact that $\widetilde{Q}$ is involutive (since it is integrable, by definition). □

These properties allow us to consider the connected Lie group $\Phi^0(u)$ generated by the Lie algebra $\mathfrak{g}(u)$, which is called the *restricted holonomy group*. From the preceding proposition, we have that $\Phi^0(u) = \Phi^0(v)$ for $v \in H(u)$ and $\Phi^0(ug) = I_{g^{-1}}(\Phi^0(u))$.

We prove that $\Phi^0(u)$ is a normal subgroup of $\Phi(u)$ and that $\Phi(u)/\Phi^0(u)$ is countable, implying that $\Phi(u)$ is a Lie subgroup of G whose identity component is precisely $\Phi^0(u)$, see [22, p 73]. We first prove that $\Phi^0(u)$ is a normal subgroup of $\Phi(u)$.

Lemma 4.3. *$\Phi^0(u)$ is a normal subgroup of $\Phi(u)$ and $\Phi(u)/\Phi^0(u)$ is a countable set.*

Proof. We start with proving that $\Phi^0(u)$ is a normal subgroup of $\Phi(u)$. Let $h \in \Phi^0(u)$. By construction of the Lie subgroup $\Phi^0(u)$ (i.e. it is defined as the leaf through e of the left invariant distribution on G generated by $\mathfrak{g}(u)$), h can be reached from e by a composite flow associated with an ordered set of left invariant vector fields in $\mathfrak{g}(u)$. Note that if a_t denotes the integral curve through e of the left invariant vector field corresponding to $A \in \mathfrak{g}(u)$, then $ua_t \in H(u)$, since $\sigma(A)$ determines a vector field tangent to $H(u)$, and hence $a_t \in \Phi(u)$. We therefore have $\Phi^0(u)$ is a subgroup of $\Phi(u)$. Since $\Phi^0(ug) = I_{g^{-1}}(\Phi^0(u))$ and $\Phi^0(u) = \Phi^0(ug)$ for any $g \in \Phi(u)$ (i.e. $\mathfrak{g}(u) = \mathfrak{g}(ug)$), we may conclude that $\Phi^0(u)$ is a normal subgroup of $\Phi(u)$.

Next, following a similar reasoning as in [22, p 73], we now prove that $\Phi(u)/\Phi^0(u)$ is countable by constructing a group morphism from $\pi_1^N(M)$ to $\Phi(u)/\Phi^0(u)$ which is onto. Since $\pi_1^N(M) < \pi_1(M)$ and $\pi_1(M)$ is at most countable, we obtain that the quotient is also countable.

Let us first make the following basic observation. In order to prove that the map from $C(x, N)$ to $\Phi(u)$ reduces to a well defined morphism from $\pi_1^N(M)$ to $\Phi(u)/\Phi^0(u)$, we must prove that two $\pm\rho$-admissible loops with homotopic base curves, correspond to elements in $\Phi(u)$ that differ by a multiplication with an element in $\Phi^0(u)$. This is achieved by using some results from standard connection theory. Once we have obtained this morphism $\pi_1^N(M) \to \Phi(u)/\Phi^0(u)$ it is easily seen to be onto, which then concludes the proof.

Consider a connection ω on P, such that $\operatorname{im} h^\omega$ is a subspace of $\widetilde{Q}$. This is always possible since $\widetilde{Q}$ is regular and $T\pi(\widetilde{Q}) = TM$. Consider the coefficient χ of h with respect to ω (see Section 2). Note that $T_e\sigma_u(\chi(u, s)) = h(u, s) - h^\omega(u, \rho(s))$ is contained in $\widetilde{Q} \cap V\pi$ for any $(u, s) \in \pi^*N$. This implies that $\chi(u, s) \in \mathfrak{g}(u)$, for all $s \in N_{\pi(u)}$. On the other hand, the holonomy group with reference point u of the standard connection ω is a subgroup of $\Phi(u)$ and the restricted holonomy group of ω is a subgroup of $\Phi^0(u)$, since the smallest integrable distribution spanned by $\operatorname{im} h^\omega$ must be contained in $\widetilde{Q}$ (see [22]).

In Section 2 (page 25) we have proven that the h-lift $c_u^h(t)$ of a ρ-admissible curve through $u \in \pi^{-1}(x)$ equals $c_u^h(t) = d^\omega(t)g(t)$, where $g(t)$ is a curve in G with $g(a) = e$ and $R_{g(t)^{-1}}\dot{g}(t) = \chi(d^\omega(t), c(t))$, and where $\dot{d}^\omega(t) =$

$h^\omega(d^\omega(t), \dot{\tilde{c}}(t))$, with $d^\omega(a) = u$. In particular we have $g(t) \in \Phi^0(u)$ (since the image of χ is contained in $\mathfrak{g}(u)$). This is also valid for the inverted anchored bundles. Thus we can conclude that any element belonging to $\Phi(u)$ can be written as a product of elements belonging to the holonomy group of ω at u and of elements belonging to $\Phi^0(u)$. Moreover, if the base of a $\pm\rho$-admissible curve is homotopic to zero, then the corresponding product of elements is entirely contained in $\Phi^0(u)$, since the restricted holonomy group of ω is a subgroup of $\Phi^0(u)$. In view of the above remarks, this completes the proof. □

From the above we immediately derive the following result (see [22, p 73] for the analogous result in the case of holonomy groups of principal connections).

Proposition 4.4. *The holonomy group $\Phi(u)$ is a Lie subgroup of the structure group G, with Lie algebra $\mathfrak{g}(u)$.*

We are now able to state a generalisation of the Reduction Theorem for principal h-lifts.

Theorem 4.5. *$H(u)$ is a reduced subbundle of P with structure group $\Phi(u)$ and h reduces to a principal ρ-lift on $H(u)$.*

Proof. We have to show that $H(u) \to M$ is a principal fibre bundle, with structure group $\Phi(u)$. For that purpose, it is sufficient to prove that, given a point $y \in M$, there exists a neighbourhood $U \ni y$ and a section ψ of π defined on U such that $\psi(U) \subset H(u)$. The existence of such a cross-section follows by using a result from [22, p 84] with respect to a connection ω with horizontal distribution contained in the regular integrable distribution $\widetilde{Q}$.

We now prove that the principal ρ-lift h induces a principal ρ-lift on the principle bundle $H(u)$, i.e. that h is reducible. Since $H(u)$ is the leaf of the foliation induced by $\mathcal{Q}$, we can consider the pull-back anchor map of h. Using the fact that $H(u)$ is a principal fibre bundle over M and using Theorem 3.1, it is easily seen that h is reducible to the pull-back of h. □

Assume that $\dim M \geq 2$. Then, since $H(u)$ is connected, there exists a standard principal connection $\overline{\omega}$ on $H(u)$ whose holonomy group is precisely the structure group $\Phi(u)$ (see [22, p 90]). Using Theorem 3.1, $\overline{\omega}$ can then be extended to a connection on P.

Corollary 4.6. *If* $\dim M \geq 2$*, then there exists a connection* ω *on* P *such that the holonomy groups of* ω *coincide with the holonomy groups of the principal* ρ*-lift* h*.*

5 The associated derivative operator

We now return to the case where $\pi : E \to M$ is a linear bundle. It is well known from standard connection theory that a horizontal lift on a linear bundle π determines a covariant derivative operator on the sections of π and vice versa. In this section we derive a similar result for ρ-lifts.

Consider a ρ-lift h on a vector bundle $\pi : E \to M$, with associated connection map K (see Equation 1.6). Take $\sigma \in \Gamma(\nu)$ and $\psi \in \Gamma(\pi)$. For any $x \in \mathrm{Dom}(\sigma) \cap \mathrm{Dom}(\psi)$ one readily verifies that $(\sigma(x), T\psi(\rho(\sigma(x))))$ determines an element of the bundle ρ^*TE. We then define $\nabla_\sigma \psi \in \Gamma(\pi)$ by

$$\nabla_\sigma \psi(x) = K\Big(\sigma(x), T\psi\big(\rho(\sigma(x))\big)\Big)\,. \tag{5.8}$$

Let $U \subset \mathrm{Dom}(\sigma) \cap \mathrm{Dom}(\psi)$ be a trivialising coordinate neighbourhood for both ν and π, with coordinates q^i on U and corresponding local bundle coordinates (q^i, u^a) and (q^i, y^A) on N and E, respectively (see also Section 1). Putting $\psi(q) = (q^i, \psi^A(q))$, we then find, using Equation (1.7):

$$\nabla_\sigma \psi(q) = \left(q^i, \frac{\partial \psi^A}{\partial q^j}(q)\gamma^j(q,\sigma(q)) - \Gamma^A_B(q,\sigma(q))\psi^B(q)\right)\,. \tag{5.9}$$

In terms of the vector field $X = \rho \circ \sigma \in \mathfrak{X}(M)$, we can still rewrite the components of $\nabla_\sigma \psi$ as

$$(\nabla_\sigma \psi)^A(q) = X^j(q)\frac{\partial \psi^A}{\partial q^j}(q) - \Gamma^A_B(q,\sigma(q))\psi^B(q)\,.$$

It is readily seen that this equation, in the case where h is a ρ-connection, takes the following form, with $\sigma(q) = (q^i, \sigma^a(q))$:

$$(\nabla_\sigma \psi)^A(q) = \frac{\partial \psi^A}{\partial q^j}(q)\gamma^j_a(q)\sigma^a(q) - \Gamma^A_{aB}(q)\sigma^a(q)\psi^B(q)\,.$$

The following theorem gives a full characterisation of the operator ∇. It is tacitly assumed that its action is restricted to those pairs $(\sigma,\psi) \in \Gamma(\nu)\times\Gamma(\pi)$ for which $\mathrm{Dom}(\sigma)$ and $\mathrm{Dom}(\psi)$ have nonempty intersection.

Theorem 5.1. *Let h be a ρ-lift on E. The operator $\nabla : \Gamma(\nu)\times\Gamma(\pi) \to \Gamma(\pi)$, defined by Equation (5.8), satisfies the following properties: if $\sigma \in \Gamma(\nu)$*
(i) $\nabla_\sigma : \Gamma(\pi) \to \Gamma(\pi)$ is $\mathbb{R}$-linear;
(ii) for all $\psi \in \Gamma(\pi)$ and $f \in C^\infty(M)$ we have:

$$\nabla_\sigma(f\psi) = f\nabla_\sigma\psi + (\rho \circ \sigma)(f)\psi\,,$$

(iii) $\nabla_\sigma\psi(x)$ only depends on the value of σ at x.

Moreover, ∇ is uniquely determined by the given ρ-lift h. If (ν, ρ) is linear and h a ρ-connection, then ∇ is, in addition, tensorial in the first argument $\Gamma(\nu)$, i.e. (iv) $\nabla_{f\sigma}\psi = f\nabla_\sigma\psi$ and $\nabla\psi : \Gamma(\nu) \to \Gamma(\pi)$ is $\mathbb{R}$-linear.

Proof. The proofs of the properties (i), (ii) and (iv) follow by straightforward computation. Property (iii) follows from the definition of the derivative operator. The fact that ∇ is uniquely determined by h can be easily deduced from Equation (5.8) and the definition of the connection map K. Indeed, different ρ-lifts necessarily induce different maps V (see Equation (1.5)) and, hence, different connection maps K (see Equation (1.6)). Note that, if h is a ρ-connection, then property (iv) induces (iii). □

We will call the operator ∇ the *derivative operator associated to the ρ-lift h*. In case $N = TM$ and ρ is the identity map on TM, we recover the classical notion of covariant derivative operator of a connection on a vector bundle over M. In his treatment of Lie algebroid connections on a vector bundle, where $N = A$ is a Lie algebroid over M with anchor map ρ, Fernandes refers to the ∇-operator as the A-derivative: see [15].

From the definition of ∇ it follows that for a given ψ, $(\nabla_\sigma\psi)(x)$ only depends on the value of σ in $x \in M$, and not on the behaviour of σ in a neighbourhood of x. This allows us to define for each $s \in N$, with $x = \nu(s)$, an operator

$$\nabla_s : \Gamma_x(\pi) \longrightarrow E_x, \psi \longmapsto \nabla_s\psi := \nabla_\sigma\psi(x), \tag{5.10}$$

where σ may be any (local) section of ν for which $\sigma(x) = s$. Alternatively, we could have defined the operator ∇_s directly according to the prescription $\nabla_s\psi = K(s, T_x\psi(\rho(s)))$. The properties of ∇_s immediately follow from Theorem 5.1, i.e. ∇_s is $\mathbb{R}$-linear and for any $f \in C^\infty(M)$ and $\psi \in \Gamma_x(\pi)$, we have that

$$\nabla_s(f\psi) = f(x)\nabla_s\psi + \rho(s)(f)\psi(x).$$

Next, let $c : I \to N$ be a smooth ρ-admissible curve in N, with corresponding base curve $\widetilde{c} = \nu \circ c$. Consider a map $\widetilde{\psi} : I \to E$, i.e. a smooth curve in E, satisfying $\pi \circ \widetilde{\psi} = \widetilde{c}$. It is now readily seen that, for each $t \in I$, $(c(t), \dot{\widetilde{\psi}}(t)) \in \rho^*TE$ and we may then define

$$\nabla_c \widetilde{\psi}(t) := K(c(t), \dot{\widetilde{\psi}}(t)) \ ,$$

which we will call *the derivative of* $\widetilde{\psi}$ *along the admissible curve* c. In coordinates, putting $c(t) = (\widetilde{c}^i(t), c^a(t))$ and $\widetilde{\psi}(t) = (\widetilde{c}^i(t), \widetilde{\psi}^A(t))$, we obtain

$$(\nabla_c \widetilde{\psi}(t))^A = \frac{d\widetilde{\psi}^A}{dt}(t) - \Gamma^A_B(c(t))\widetilde{\psi}^B(t) \ .$$

In the specific case where h is a ρ-connection, this equation becomes

$$(\nabla_c \widetilde{\psi}(t))^A = \frac{d\widetilde{\psi}^A}{dt}(t) - \Gamma^A_{aB}(\widetilde{c}(t))c^a(t)\widetilde{\psi}^B(t) \ .$$

Assume one can find a (local) section $\psi \in \Gamma(\pi)$ such that $\psi(\widetilde{c}(t)) = \widetilde{\psi}(t)$ for all $t \in I$. This will be the case, for instance, if the base curve $\widetilde{c}$ is an injective immersion. A straightforward computation then shows that

$$\begin{aligned}(\nabla_{c(t)}\psi)^A &= \frac{\partial \psi^A}{\partial q^j}(\widetilde{c}(t))\dot{\widetilde{c}}^j(t) - \Gamma^A_B(c(t))\psi^B(\widetilde{c}(t)) \\ &= \frac{d\widetilde{\psi}^A}{dt}(t) - \Gamma^A_B(c(t))\widetilde{\psi}^B(t) \ ,\end{aligned}$$

where, for the second equality, we have used the fact that $\psi^A(\widetilde{c}(t)) \equiv \widetilde{\psi}^A(t)$. We may therefore conclude that the derivative of $\widetilde{\psi}$ along c verifies

$$\nabla_c \widetilde{\psi}(t) = \nabla_{c(t)}\psi,$$

for any $\psi \in \Gamma(\pi)$ such that $\psi(\widetilde{c}(t)) \equiv \widetilde{\psi}(t)$, if such a section ψ exists. The following theorem is interesting since it gives an interpretation of the derivation of $\widetilde{\psi}$ along $c : [a, b] \to N$ in the sense that $\nabla_c \widetilde{\psi}$ "measures" the extent by which $\widetilde{\psi}$ deviates from being a h-admissible curve. We first introduce a specific notation. Let $t \leq t' \in [a, b]$ and let $(c^h)^{t'}_t$ denote the h-displacement along the restriction of c to $[t, t']$. Thus we have: $(c^h)^{t'}_t : E_{\widetilde{c}(t)} \to E_{\widetilde{c}(t')}$. If $t' \leq t$, then we define $(c^h)^{t'}_t = ((c^h)^t_{t'})^{-1}$.

Theorem 5.2. *The following equation holds:*

$$\nabla_c \widetilde{\psi}(t) = \left.\frac{d}{ds}\right|_{s=0} \Big((c^h)^t_{t+s}(\widetilde{\psi}(t+s)) \Big) \ .$$

Proof. The right-hand side of the above equation is well defined, since it is the tangent vector to a curve contained in the linear space $E_{\widetilde{c}(t)}$. The proof follows from $(c^h)_t^{t'}(\overline{e}) = c_e^h(t')$ for $\overline{e} \in E_{\widetilde{c}(t)}$ and $e = ((c^h)_a^t)^{-1}(\overline{e})$, implying that

$$\left.\frac{d}{dt'}\right|_{t'=t} (c^h)_t^{t'}(\overline{e}) = h(\overline{e}, c(t)).$$

To simplify notations, we define $\overline{e}(s) = (c^h)_{t+s}^t(\widetilde{\psi}(t+s))$, which can be regarded as a curve in $E_{\widetilde{c}(t)}$. Using the above, we obtain that the derivative at $s = 0$ of the composed curve $s \mapsto (c^h)_t^{t+s}(\overline{e}(s)) \equiv \widetilde{\psi}(t+s)$ equals

$$h(\overline{e}(0), c(t)) + \dot{\overline{e}}(0) = \dot{\widetilde{\psi}}(t),$$

where we have used the Leibniz rule. The proof is completed using the definition of the connection map K and the following equality:

$$p_2(\dot{\overline{e}}(0)) = \left.\frac{d}{ds}\right|_{s=0} \Big((c^h)_{t+s}^t(\widetilde{\psi}(t+s))\Big). \quad \square$$

Remark 5.3. A special situation occurs, for instance, when ρ is linear and has a nontrivial kernel containing the image of an admissible curve c. In particular, we then know that $c(t)$ necessarily belongs to a fixed fibre of ν and the base curve $\widetilde{c}$ reduces to a point in M, say $\widetilde{c}(t) = \nu(c(t)) = x$ for all t. We then consider a map $\widetilde{\psi} : I = [a,b] \to E_x$. In coordinates, with $x = (q_0^i)$, $\widetilde{\psi}(t) = (q_0^i, y^A(t))$, we then find that, with ∇ the derivative operator of a ρ-connection on E:

$$\nabla_c \widetilde{\psi}(t) = \big(q_0^i, \dot{y}^A(t) - \Gamma_{aB}^A(q_0) c^a(t) y^B(t)\big) \in E_x \, .$$

In particular, if we associate to each point $e_0 = (q_0^i, y_0^A) \in E_x$ the constant map $\widetilde{\psi}(t) \equiv (q_0^i, y_0^A)$, we obtain a time-dependent linear map on the fibre E_x, namely $e_0 \mapsto \nabla_c e_0(t) = (q_0^i, -\Gamma_{aB}^A(q_0) c^a(t) y_0^B)$.

Next, it is easy to see how the action of the derivative operator of a ρ-lift on a vector bundle $\pi : E \to M$, can be extended to sections of the dual vector bundle $\pi^* : E^* \to M$. If, by convention, for $\sigma \in \Gamma(\nu)$ and $f \in C^\infty(M)$ we put $\nabla_\sigma f = (\rho \circ \sigma)(f)$, we can immediately define an action of the operator ∇_σ on $\Gamma(\pi^*)$ as follows: for any $\mathfrak{f} \in \Gamma(\pi^*)$, $\nabla_\sigma \mathfrak{f} \in \Gamma(\pi^*)$ is uniquely determined by

$$\langle \nabla_\sigma \mathfrak{f}, \psi \rangle = \nabla_\sigma \langle \mathfrak{f}, \psi \rangle - \langle \mathfrak{f}, \nabla_\sigma \psi \rangle \, ,$$

for all $\psi \in \Gamma(\pi)$, where $\langle\ ,\ \rangle$ denotes the canonical pairing between sections of π and sections of π^*. Herewith, it is then standard to further extend the action of ∇_σ to sections of any tensor bundle constructed out of E and E^*.

In what precedes we have shown that a ρ-lift on a vector bundle $\pi : E \to M$ gives rise to an operator ∇ verifying the conditions of Theorem 5.1. We now demonstrate that the converse also holds.

Theorem 5.4. *Any operator $\nabla : \Gamma(\nu) \times \Gamma(\pi) \to \Gamma(\pi)$, verifying the properties (i), (ii) and (iii) of Theorem 5.1, is the derivative operator of a unique ρ-lift on π. If (ν, ρ) is a linear anchored bundle, and if (i), (ii) and (iv) are satisfied, then the ρ-lift is a ρ-connection.*

Proof. Take $s \in N$, with $\nu(s) = x$, and $\psi \in \Gamma_x(\pi)$. From (iii) it follows that the given operator ∇ induces an operator ∇_s on $\Gamma_x(\pi)$ such that $\nabla_s\psi \in E_x$. Putting $\psi(x) = e$, and denoting by $\iota_e : E_x \to V_e\pi$ the canonical isomorphism between the vector spaces E_x and $V_e\pi$, we may consider the vector $T\psi(\rho(s)) - \iota_e(\nabla_s\psi) \in T_eE$. It is now straightforward to check that the mapping $\Gamma_x(\pi) \to T_eE : \psi \mapsto T\psi(\rho(s)) - \iota_e(\nabla_s\psi)$ is $C^\infty(M)$-linear in ψ and, hence, only depends on the value of ψ in x. From this we deduce that there exists a well-defined smooth mapping $h : \pi^*N \to TE$, given by

$$h(e, s) = T\psi(\rho(s)) - \iota_e(\nabla_s\psi)\,,$$

for any $\psi \in \Gamma(\pi)$ with $\psi(\nu(s)) = e$. Clearly, $T\pi(h(e, s)) = \rho(s)$, which already shows that $\rho \circ \widetilde{\pi}_2 = \pi_* \circ h$. We now have to check that h commutes with $\{\lambda_t\}$, the flow of the dilation vector field on E. Observe that for any $\psi \in \Gamma(\pi)$ we also have $\lambda_t \circ \psi \in \Gamma(\pi)$ for each $t \in \mathbb{R}$. It is not difficult to verify that

$$\begin{aligned} T(\lambda_t)\Big(T\psi(\rho(s)) - \iota_e(\nabla_s\psi)\Big) &= T(\lambda_t \circ \psi)(\rho(s)) - \iota_{\lambda_t(e)}(\nabla_s(\lambda_t \circ \psi)) \\ &= h(\lambda_t(e), s)\,, \end{aligned}$$

proving that h is indeed a ρ-lift. If ∇ is linear in the first argument, the linearity of $h_e = h(e, .) : N_{\pi(e)} \to T_eE$ is obvious.

It now remains to be shown that the given ∇ is the derivative operator of the constructed ρ-lift h. Let K denote the connection map associated to h (cf. Section 1). Using Equations (1.6) and (1.5), together with the above definition of h, we find for any $s \in N$ and $\psi \in \Gamma_{\nu(s)}(\pi)$, putting $\psi(\nu(s)) = e$:

$$\begin{aligned} K(s, T\psi(\rho(s))) &= p_2(T\psi(\rho(s)) - h(e,s)) \\ &= p_2(\iota_e(\nabla_s \psi)) \\ &= \nabla_s \psi \, . \end{aligned}$$

Since, in view of Equation (5.8), the left-hand side precisely determines the derivative operator associated to h, which completes the proof. □

In view of this result, the derivation on the dual bundle E^* determines a ρ-lift on the bundle E^*, which will be denoted by h^*. It is interesting to know how the transport operator of this "dual ρ-lift" is related to the h-transport operator on E. Therefore, assume that ∇ is a derivative operator associated with a ρ-lift h on E. Let $c : [a,b] \to N$ denote a ρ-admissible curve, with base curve $\widetilde{c}$. The transport operator $(c^h)_t^{\bar{t}}$ is a linear isomorphism from $E_{\widetilde{c}(t)}$ to $E_{\widetilde{c}(\bar{t})}$, for any $t \leq \bar{t}$. Let $\mathfrak{f} : [a,b] \to E^*$ denote the section of E^* along $\widetilde{c}$, defined by $\mathfrak{f}(t) = ((c^h))_t^a)^*(\mathfrak{f}_a)$, with $\mathfrak{f}_a \in E^*_{\widetilde{c}(a)}$ arbitrary. We now prove that $\mathfrak{f}(t)$ satisfies $\nabla_c \mathfrak{f}(t) = 0$ for all t, which implies that the map $((c^h)^{-1})^*$ is the transport operator c^{h^*} determined by the derivative operator ∇ acting on the dual bundle E^*. For that purpose, we fix a $t' \in [a,b]$ and an arbitrary element $e_{t'}$ in $E_{\widetilde{c}(t')}$, and we consider $\psi(t) = (c^h)_{t'}^t(e_{t'})$, for $t \in [a,b]$. In particular $\psi(t)$ is a section of E along $\widetilde{c}$, satisfying $\nabla_c \psi(t) = 0$. We now consider the following equation:

$$\langle \nabla_c \mathfrak{f}(t'), e_{t'} \rangle = \left. \frac{d}{dt} \right|_{t'} \langle \mathfrak{f}(t), \psi(t) \rangle - \langle \mathfrak{f}(t'), \nabla_c \psi(t') \rangle.$$

Since $\langle \mathfrak{f}(t), \psi(t) \rangle = \langle \mathfrak{f}(t'), e_{t'} \rangle$ is constant, we obtain that $\nabla_c \mathfrak{f}(t') = 0$ for all $t' \in [a,b]$.

To close this section we consider the difference between the derivative operators associated to two ρ-lifts in order to characterise all ρ-lifts on a bundle E. Thus, suppose ∇ and $\bar{\nabla}$ are the derivative operators corresponding to two (different) ρ-lifts on the vector bundle π. It follows from Theorem 5.1 that the difference $\nabla - \bar{\nabla}$ is a mapping $S : \Gamma(\nu) \times \Gamma(\pi) \to \Gamma(\pi)$ which is $C^\infty(M)$-linear in the second argument and which locally reads

$$(S(\sigma, \psi)(x))^A = (\Gamma^A_B(\sigma(x)) - \bar{\Gamma}^A_B(\sigma(x)))\psi^B(x) \, .$$

Note that S uniquely determines a section $\mathcal{S}$ of the tensor bundle $(\pi^*N)^* \otimes \pi^*N \to N$, which is well defined since π^*N is a linear bundle over N and $(\pi^*N)^*$ its dual. In the case where ∇ and $\bar{\nabla}$ are derivative operator of ρ-connections, the map S is also $C^\infty(M)$-linear in the first argument and the above equation then becomes, with $x = (q^i)$ and $\sigma(q) = (q^i, \sigma^a(q))$:

$$(S(\sigma,\psi)(q))^A = (\Gamma^A_{aB}(q) - \bar{\Gamma}^A_{aB}(q))\sigma^a(q)\psi^B(q)$$

In this case $\mathcal{S}$ is a section of the tensor product bundle $N^* \otimes E^* \otimes E \to M$.

Conversely, given a derivative operator ∇ and an arbitrary section $\mathcal{S} : N \to (\pi^*N)^* \otimes \pi^*N$, then the operator $\nabla + S$, with $S : \Gamma(\nu) \times \Gamma(\pi) \to \Gamma(\pi)$ associated to $\mathcal{S}$ as above, maps any pair of sections (σ, ψ) onto $\nabla_\sigma\psi + S(\sigma,\psi)$ and also defines a derivative operator verifying the assumptions of Theorem 5.1. Hence, determines a ρ-lift on π.

6 Invariant subbundles

Let (ν, ρ) denote an anchored bundle on M. In this section we study how a ρ-lift h on a vector bundle $\pi : E \to M$ acts on a vector subbundle $\epsilon : F \to M$ of π, (i.e. F_x is a linear subspace of E_x for any $x \in M$). Let $\nabla : \Gamma(\nu) \times \Gamma(\pi) \to \Gamma(\pi)$ denote the derivative operator associated with the ρ-lift.

Definition 6.1. The vector subbundle $\epsilon : F \to M$ is called *invariant under* ∇ if, for any $\psi \in \Gamma(\epsilon)$ and $\sigma \in \Gamma(\nu)$, we have that $(\nabla_\sigma\psi) \in \Gamma(\epsilon)$.

We now prove some elementary properties concerning invariant subbundles. It is easily seen that the derivation ∇, leaving ϵ invariant, determines a derivation on the bundle ϵ, which will be denoted by $\overline{\nabla}$. Moreover, it is easily checked that $\overline{\nabla}$ satisfies all desired properties to determine a ρ-lift $\overline{h}$ on ϵ (cf. Section 5). Assume that $c : [a,b] \to N$ is a ρ-admissible curve taking $x \in M$ to $y \in M$, and let c^h denote the displacement along c with respect to h.

Proposition 6.1. *The natural injection $i : F \to E$ is a morphism between the ρ-lifts $\overline{h}$ and h.*

Proof. Recall the construction of a ρ-lift associated to a given derivative operator (cf. Theorem 5.4). If ψ is a section of ϵ, then we can consider the section $i \circ \psi$ of π. Let $K' : \rho^*TF \to F$ denote the connection map of $\overline{h}$ and $K : \rho^*TE \to E$ the connection map of h. Then, the equation $\nabla_\sigma(i \circ \psi) = i \circ \overline{\nabla}_\sigma \psi$, is equivalent to

$$i\Big(K'\big(\sigma(x), T\psi(\rho(\sigma(x)))\big)\Big) = K\Big(\sigma(x), Ti\big(T\psi(\rho(\sigma(x)))\big)\Big),$$

for all $\sigma \in \Gamma(\nu)$, $\psi \in \Gamma(\epsilon)$. In turn, this equation is equivalent to

$$\begin{aligned} &Ti\Big(T\psi\big(\rho(\sigma(x))\big) - \overline{h}\big(\sigma(x), \psi(x)\big)\Big) = \\ &T\big(i \circ \psi\big)\big(\rho(\sigma(x))\big) - h\big(\sigma(x), i(\psi(x))\big), \end{aligned}$$

or $Ti(\overline{h}(\sigma(x), \psi(x))) = h(\sigma(x), i(\psi(x)))$, which is precisely the condition for i to determine a morphism between the ρ-lifts $\overline{h}$ and h. □

As a consequence of the above theorem and the theory on morphisms between anchored bundles developed in Chapter I, the following property holds. Let $c : [a, b] \to N$ be a ρ-admissible curve with base curve $\widetilde{c}$. Then, given any section ψ of ϵ along $\widetilde{c}$, we have $i(\overline{\nabla}_c \psi(t)) = \nabla_c i(\psi(t))$ and, hence, $c^h|_{F_x} = c^{\overline{h}}$.

A similar result holds for the derivative operations acting on sections of $\epsilon^* : F^* \to M$ and $\pi^* : E^* \to M$. Consider the action of ∇ and $\overline{\nabla}$ on $\pi^* : E^* \to M$ and $\epsilon^* : F^* \to M$, respectively. Then, given any $\mathfrak{f} \in \Gamma(\pi^*)$, we have

$$\langle \nabla_\sigma \mathfrak{f}, \psi \rangle = \rho(\sigma)(\langle \mathfrak{f}, \psi \rangle) - \langle \mathfrak{f}, \nabla_\sigma \psi \rangle,$$

for $\sigma \in \Gamma(\nu)$ and $\psi \in \Gamma(\pi)$ arbitrary. If we denote the dual of $i : F \to E$ by $i^* : E^* \to F^*$, then the section $i^*(\nabla_\sigma \mathfrak{f})$ of ϵ^*, is obtained by considering in the previous equation only sections ψ belonging to $\Gamma(\epsilon)$, which implies $\overline{\nabla}_\sigma i^* \mathfrak{f} = i^*(\nabla_\sigma \mathfrak{f})$. Using similar techniques as in the proof of the above proposition, we obtain that i^* is a mapping between the lifts h^* and $\overline{h}^*$. Using the theory on morphisms between anchored bundles, we obtain that the transport operators of h^* and $\overline{h}^*$ along a ρ-admissible curve c, satisfy: $i^* \circ c^{h^*} = c^{\overline{h}^*} \circ i^*$.

7 General properties on ρ-connections and examples

A ρ-connection on a linear bundle $\pi : E \to M$ admits some interesting properties, which can not be obtained in the general case where (ν, ρ) is not linear. Therefore, in the remainder of this chapter we shall only consider ρ-connections. Recall first of all that any $\pm\rho$-admissible curve can be related to a ρ-admissible curve which, in turn, implies that the holonomy groups can be defined without making use of the additional concept of a $\pm\rho$-admissible curve. In this section we prove some properties on the distribution $Q = \operatorname{im} h$ defined by the ρ-connection h and we consider various situations in differential geometry where generalised connections over a vector bundle map may be considered.

As observed above, the distribution Q defined by a ρ-connection h on a fibre bundle $\pi : E \to M$, in general may have nonzero intersection with the vertical subbundle $V\pi$ of TE. The extent by which Q fails to be a (full) complement of $V\pi$ is characterised by the following proposition.

Proposition 7.1. *For any $x \in M$ and $e \in E_x$ we have*

$$Q_e \cap V_e\pi \cong \ker(\rho_x)/\ker(h_e), \tag{7.11}$$

(where ρ_x and h_e are the linear maps induced by the restrictions of ρ and h, respectively, to the fibre N_x of N), and

$$Q_e + V_e\pi = T_eE \quad \Longleftrightarrow \quad D_x = T_xM. \tag{7.12}$$

Proof. For $w \in T_eE$, with $\pi(e) = x$, we immediately have that $w \in Q_e \cap V_e\pi$ iff $w = h(e,s) = h_e(s)$ for some $s \in N_x$, and $0 = T\pi(w) = T\pi(h(e,s)) = \rho_x(s)$. Hence,

$$w \in \mathcal{Q}_e \cap V_e\pi \iff w \in h_e(\ker(\rho_x)).$$

From the definition of h one can deduce that $\ker(h_e) \subset \ker(\rho_x)$ and it then readily follows that $h_e(\ker(\rho_x)) \cong \ker(\rho_x)/\ker(h_e)$, which completes the proof of (7.11).

Next, assume that $Q_e + V_e\pi = T_eE$, for $e \in E_x$. For any $v \in T_xM$ one can always find a $w \in T_eE$ such that $T\pi(w) = v$. The given assumption implies that w can be written as $w = h_e(s) + \widetilde{w}$, for some $s \in N_x$ and $\widetilde{w} \in V_e\pi$, and this, in turn, gives

$$v = T\pi(w) = T\pi(h_e(s)) = \rho_x(s),$$

i.e. $v \in \operatorname{im}(\rho_x) = D_x$. Since $v \in T_xM$ was chosen arbitrarily, this proves that $D_x = T_xM$. Conversely, assume $D_x = T_xM$. For any $w \in T_eE$ we then have that $T\pi(w) = \rho_x(s)$ for some $s \in N_x$, from which it follows that

$T\pi(w - h_e(s)) = T\pi(w) - \rho_x(s) = 0$, and so $w - h_e(s) \in V_e\pi$. This completes the proof of the equivalence (7.12). □

From this proposition one can readily deduce the following result.

Corollary 7.2. *The distribution Q defines a genuine (Ehresmann) connection on π iff $\mathrm{im}(\rho) = TM$ and $\ker(\rho_x) = \ker(h_e)$ for all $x \in M$ and $e \in E_x$.*

Whereas a ρ-connection h determines a (generalised) distribution Q on E which projects onto D, the converse is certainly not true in general. Moreover, if a distribution Q can be associated to a ρ-connection, the latter need not be uniquely determined. A sufficient condition for a distribution on E to correspond to a unique ρ-connection is that it determines a (not necessarily full) complement of $V\pi$.

Proposition 7.3. *Let F be a smooth generalised distribution on E such that (i) $T\pi(F) = D$, and (ii) $F_e \cap V_e\pi = \{0\}$ for all $e \in E$, then there exists a unique ρ-connection h such that $F = \mathrm{im}\, h = Q$.*

Proof. For each point $e \in E$, we can construct a map $h_e : N_x \to T_eE$, where $x = \pi(e)$, by putting

$$\{h_e(s)\} = F_e \cap \left((T\pi)_{|T_eE}\right)^{-1} (\rho_x(s)),$$

for all $s \in N_x$. From the given assumptions (i) and (ii), it follows that this prescription uniquely determines a point $h_e(s)$. Furthermore, using some simple set-theoretic arguments, it is not difficult to verify that the resulting map h_e is linear. Next, we can 'glue' these linear maps together to a smooth bundle map $h : \pi^*N \to TE$ with $h(e, s) = h_e(s)$. It is then straightforward to see that, by construction, h verifies all properties of a ρ-connection.

Finally, uniqueness of h can be proved as follows. Let $h' : \pi^*N \to TE$ be another ρ-connection for which $\mathrm{im}(h') = F$. Then, for each $(e, s) \in \pi^*N$, with $\pi(e) = \nu(s) = x$, there exists a $s' \in N_x$ such that $h'(e, s) = h(e, s')$. The definition of a ρ-connection then implies that $\rho(s) = \rho(s')$. Since $h(e, s)$ is the unique element in Q that projects onto $\rho(s) = \rho(s')$ and since $\ker(\rho_x) = \ker(h_e)$ it follows that $h'(e, s) = h(e, s') = h(e, s)$, which indeed proves uniqueness of the ρ-connection. □

Note that the ρ-connections referred to in Proposition 3.3 are of a special type in the sense that the corresponding distribution Q is 'transverse' to $V\pi$, i.e. $Q_e \cap V_e\pi = \{0\}$ for all $e \in E$. With a slight abuse of terminology, we will call such a ρ-connection a *partial ρ-connection* on π. If the distribution Q has constant rank it determines indeed a partial connection in the ordinary sense.

Remark 7.4. The notion of partial connection, as defined above, also corresponds to (and reduces to) what Fernandes has called $\mathcal{F}$-connections in his treatment of contravariant connections on Poisson manifolds and connections on Lie algebroids [14, 15].

Assume now that ρ has constant rank. Then, $D = \mathrm{im}(\rho)$ is a vector subbundle of TM, with canonical injection $i : D \hookrightarrow TM$.

Proposition 7.5. *If ρ has constant rank, then for every ρ-connection h on a fibre bundle $\pi : E \to M$ there is a i-connection $\bar{h}$ on π such that* $\mathrm{im}(h) = \mathrm{im}(\bar{h})$ *iff h is a partial connection.*

Proof. If h is a partial connection, we know from the above that $\ker(\rho_x) = \ker(h_e)$ for all $x \in M$ and $e \in E_x$. We can then define a mapping $\bar{h} : \pi^* D \to TE$ as follows: for $s \in N_x$ and $e \in E_x$, put

$$\bar{h}(e, \rho(s)) = h(e, s) \, .$$

From the fact that h is a partial connection it follows that $\bar{h}$ is well defined, and it is straightforward to check that it is a generalised connection over i, determining the same distribution Q on E as h.

Conversely, assume that there exists a i-connection $\bar{h}$ on π, having the same image as a given ρ-connection h. In particular, this implies that for all $(e, s) \in \pi^* N$, with $\nu(s) = \pi(e) = x$, there exists a $s' \in N_x$ such that $h(e, s) = \bar{h}(e, \rho(s'))$. Since, obviously, $\ker(i) = 0$, we also have $\ker(\bar{h}_e) = 0$, from which one can readily deduce that $\ker(h_e) = \ker(\rho_x)$ and, hence, h is a partial connection. □

Next, we shall describe several situations where generalised connections over a vector bundle map may be considered. In particular, we will see how the various types of connections mentioned in the Introduction can be recovered as special cases of the general notion of connection put forward in Definition 1.1.

Some examples

(i) If we put $N = TM$, $\nu = \tau_M$ and $\rho = \mathrm{id}_{TM}$ (the identity map on TM), Definition 1.1 reduces to that of an ordinary connection (an Ehresmann connection) on π, with $h : \pi^*TM \to TE$ defining a splitting of the short exact sequence $0 \to V\pi \to TE \to \pi^*TM \to 0$ and $\mathrm{im}(h) = H\pi$ the horizontal distribution of the connection. In particular, for $E = TM$ we recover the standard notion of connection on a manifold M (see also [57]).

(ii) Let N be a subbundle of TM, $\nu = (\tau_M)_{|N}$, and $\rho = i_N : N \hookrightarrow TM$ the canonical injection. In this case, each ρ-connection h on a fibre bundle $\pi : E \to M$ is a partial connection. Indeed, since for all $x \in M$ we have $\ker((i_N)_x) = \{0\}$, it follows from (7.11) that $Q_e \cap V_e\pi = \{0\}$ for all $e \in E$. Moreover, h is now necessarily injective, implying that Q is a constant rank distribution and, therefore, we are dealing with a partial connection in the ordinary sense. Partial connections are considered in particular in those cases where N defines a regular integrable distribution on M (see e.g. [21]). The horizontal subspaces Q_e then project onto the tangent spaces to the leaves of the induced foliation. But partial connections also make their appearance, for instance, in the framework of sub-Riemannian geometry, where N is a subbundle of TM equipped with a nondegenerate bundle metric (see e.g. [13]). For that purpose, we shall now give a definition of what we call a bundle adapted connection. Let ∇ denote a i_N-connection on the bundle T^*M. Let $\alpha \in \Gamma(N^0)$ denote an arbitrary section of N^0, the annihilator bundle of N, and let $s \in N$. Then ∇ is called *adapted to the bundle N, (shortly N-adapted)* if $\nabla_s\alpha = i_s d\alpha$. It is easily seen that this is well defined, and in general $\nabla_s\alpha \in T^*M$. However, in the specific case where D is integrable, the bundle N^0 is an invariant subbundle of T^*M with respect to ∇. From Section 6, we know that ∇ determines a i_N-connection on N^0. It is precisely this partial i_N-connection that was used by R. Bott et al. in [2] to prove that certain Pontryagin cohomology classes of the bundle N^0 are identically zero, provided that D is integrable. Nevertheless, this mapping naturally pops up in our approach to sub-Riemannian geometry (cf. Chapter IV), and therefore, we define a mapping δ according to

$$\delta : \Gamma(N) \times \Gamma(N^0) \to \mathfrak{X}^*(M) : (X, \eta) \mapsto \delta_X\eta = i_X d\eta.$$

(iii) If $\nu : N \to M$ is a Lie algebroid over M, with anchor map ρ, we recover the notion of *Lie algebroid connection* studied by Fernandes [15]. By definition of a Lie algebroid, the anchor map induces a Lie algebra morphism from the Lie algebra of sections of ν into the Lie algebra of vector fields on M. In this case $\mathrm{im}(\rho) = D$ is an integrable generalised distribution, determining a (possibly singular) foliation of M. Given a ρ-connection h on a fibre bundle $\pi : E \to M$, with associated distribution Q, we have that for each $e \in E$ the subspace Q_e of T_eE projects onto the tangent space at $\pi(e)$ to the leaf passing through $\pi(e)$ of the foliation determined by D. Here, unlike the case of a partial connection, Q may have a nonzero intersection with the vertical distribution $V\pi$.

A particular instance of a Lie algebroid is obtained when M admits a Poisson structure, with Poisson tensor Λ, and $N = T^*M$. The anchor map ρ is then given by the natural vector bundle morphism induced by Λ, i.e. $\sharp_\Lambda : T^*M \to TM,\ \alpha_x \mapsto \Lambda_x(\alpha_x, .)$. This case was also studied extensively by Fernandes [14]. Connections over $\sharp_\Lambda$ were then called *contravariant connections*, following I. Vaisman who introduced a notion of contravariant derivative in the framework of the geometric quantisation of Poisson manifolds [55].

(iv) Let again $N = TM$, $\nu = \tau_M$ and let ρ be the tangent bundle morphism induced by a type $(1,1)$-tensor field A on M. A ρ-connection then corresponds to what is also known as a *pseudo-connection* with fundamental tensor field A (cf. [12, 59]).

Consider the case where A has vanishing Nijenhuis torsion, i.e. $\mathcal{N}_A = 0$, with $\mathcal{N}_A$ the type $(1,2)$-tensor field defined by $1/2\mathcal{N}_A(X,Y) = A^2([X,Y]) + [A(X), A(Y)] - A([A(X),Y]) - A([X,A(Y)])$ for arbitrary $X, Y \in \mathfrak{X}(M)$. The pair (M, A) is sometimes called a *Nijenhuis manifold*, with *Nijenhuis tensor* A. One may then define a new bracket on $\mathfrak{X}(M)$ according to

$$[X,Y]_A := [A(X),Y] + [X,A(Y)] - A([X,Y])\,. \tag{7.13}$$

Using the fact that $\mathcal{N}_A = 0$, it follows after some tedious but straightforward computations that $[\ ,\ ]_A$ is again a Lie bracket on $\mathfrak{X}(M)$ and that, moreover, $A([X,Y]_A) = [A(X), A(Y)]$ and $[X, fY]_A = f[X,Y]_A + A(X)(f)Y$ for all $X, Y \in \mathfrak{X}(M)$ and $f \in C^\infty(M)$ (see e.g. [25]). Consequently, TM becomes a Lie algebroid over M with bracket $[\ ,\ ]_A$ and anchor map A (regarded as a bundle map from TM into itself), and a pseudo-connection whose fundamental tensor field A is a Nijenhuis tensor, is a Lie algebroid connection.

(v) An immediate extension of the previous case is obtained when considering an arbitrary vector valued tensor field $\mathcal{K} \in \mathcal{T}^r_s(M) \otimes \mathfrak{X}(M)$ on M, where $\mathcal{T}^r_s(M)$ denotes the $C^\infty(M)$-module of smooth type (r,s)-tensor fields, i.e. tensor fields of contravariant order r and covariant order s. Putting $N = T^s_r(TM)$, the vector bundle of type (s,r)-tensors on M, and $\rho : T^s_r(TM) \to TM$ the natural bundle morphism over M induced by $\mathcal{K}$, i.e.

$$\rho(v_1 \otimes \cdots \otimes v_s \otimes \alpha_1 \otimes \cdots \otimes \alpha_r) = \mathcal{K}(v_1, \ldots, v_s; \alpha_1, \ldots, \alpha_r),$$

for arbitrary $x \in M$, $v_i \in T_xM$ and $\alpha_j \in T_x^*M$, then one can consider ρ-connections on a fibre bundle E over M as connections which, in some sense, are "parameterised" by (s, r)-tensors. Clearly, the pseudo-connections mentioned above, as well as the contravariant (Poisson) connections, belong to this category.

8 Curvature and torsion for ρ-connections

Clearly, in the case of arbitrary vector bundles $\nu : N \to M$ and $\pi : E \to M$ there is no way, in general, of assigning a notion of torsion or curvature to a ρ-connection. However, let us assume in what follows that the space of sections $\Gamma(\nu)$ is equipped with an algebra structure (over $\mathbb{R}$), with product denoted by $*$, such that the mapping $\Gamma(\nu) \times \Gamma(\nu) \to \Gamma(\nu), (\sigma_1, \sigma_2) \mapsto \sigma_1 * \sigma_2$ is $\mathbb{R}$-bilinear and skew-symmetric and, in addition, verifies a Leibniz-type rule

$$\sigma_1 * (f\sigma_2) = f(\sigma_1 * \sigma_2) + \rho(\sigma_1)(f)\sigma_2 \, , \tag{8.14}$$

for all $\sigma_1, \sigma_2 \in \Gamma(\nu)$ and $f \in C^\infty(M)$. Note that we do not require ρ to induce an algebra morphism between $(\Gamma(\nu), *)$ and $(\mathfrak{X}(M), [\ ,\ ])$.

Whenever the space of sections of the vector bundle $\nu : N \to M$ is equipped with a bilinear operation $*$ verifying the above assumptions, we will follow [17] in saying that N admits the structure of a *pre-Lie algebroid*. When dropping the skew-symmetry assumption of the product $*$, we obtain a so-called pseudo-Lie algebroid (cf. [17]). For a treatment of the differential calculus on pseudo- and pre-Lie algebroids, we refer to [18], where both structures are simply called "algebroids" and "skew algebroids", respectively. The algebraic counterpart of pre-Lie algebroids, namely differential pre-Lie algebras, have also been studied in [25].

In analogy with the Poisson structure that exists on the dual bundle of any Lie algebroid, one can show that on the dual bundle $\mu : N^* \to M$ of any pre-Lie algebroid $\nu : N \to M$ there exists a distinguished bivector field Λ which, in particular, induces an 'almost-Poisson' bracket $\{\ ,\ \}$ on $C^\infty(N^*)$, verifying all properties of a Poisson bracket except for the Jacobi identity. One can show that the Schouten-Nijenhuis bracket of the bivector field Λ with itself vanishes (and, hence, Λ becomes a Poisson tensor) iff the algebra $(\Gamma(\nu), *)$ is a Lie algebra, i.e. the Jacobi identity holds for the $*$-product. In that case one can also prove that ρ induces a Lie algebra homomorphism between $(\Gamma(\nu), *)$ and $(\mathfrak{X}(M), [\ ,\ ])$, and N then becomes a Lie algebroid over M (cf. [17]).

Assume N admits a pre-Lie algebroid structure, with product $*$ on $\Gamma(\nu)$, and consider a ρ-connection on a vector bundle $\pi : E \to M$, with associated derivative operator ∇. We may now define a mapping $R : \Gamma(\nu) \times \Gamma(\nu) \times \Gamma(\pi) \to \Gamma(\pi)$ given by

$$R(\sigma_1, \sigma_2; \psi) := \nabla_{\sigma_1} \nabla_{\sigma_2} \psi - \nabla_{\sigma_2} \nabla_{\sigma_1} \psi - \nabla_{\sigma_1 * \sigma_2} \psi\ . \tag{8.15}$$

It easily follows that R is $C^\infty(M)$-linear and skew-symmetric in σ_1 and σ_2, but fails to be $C^\infty(M)$-linear in ψ. Indeed, a straightforward computation shows that for arbitrary $\sigma_1, \sigma_2 \in \Gamma(\nu), \psi \in \Gamma(\pi)$ and $f \in C^\infty(M)$,

$$\begin{aligned} R(\sigma_1, \sigma_2; f\psi) &= fR(\sigma_1, \sigma_2; \psi) + (\rho(\sigma_1) \circ \rho(\sigma_2) \\ &\quad - \rho(\sigma_2) \circ \rho(\sigma_1) - \rho(\sigma_1 * \sigma_2))\,(f)\psi\ . \end{aligned}$$

From this expression it is seen that R will be fully tensorial iff ρ induces an algebra homomorphism from $(\Gamma(\nu), *)$ to $(\mathfrak{X}(M), [\ ,\ ])$, i.e.

$$\rho(\sigma_1 * \sigma_2) = [\rho(\sigma_1), \rho(\sigma_2)]. \tag{8.16}$$

In particular, this implies that for all $\sigma_1, \sigma_2, \sigma_3 \in \Gamma(\nu)$ we have

$$\sigma_1 * (\sigma_2 * \sigma_3) + \sigma_2 * (\sigma_3 * \sigma_1) + \sigma_3 * (\sigma_1 * \sigma_2) \in \Gamma(\nu_{|\ker(\rho)})\ ,$$

i.e. the 'Jacobiator' of the $*$-product should take values in the kernel of the vector bundle morphism ρ. (The denomination 'Jacobiator' is sometimes used in the literature to indicate, in an algebra with a skew-symmetric product, the cyclic sum that vanishes in case the Jacobi identity holds). If (8.16) holds, we will call the mapping R, defined by (8.15), the *curvature* of the given ρ-connection.

Remark 8.1. Another important consequence of (8.16) is that the generalised distribution $D(= \mathrm{im}(\rho))$ on M is involutive. Note, however, that since ρ need not be of constant rank, involutivity does not necessarily imply integrability of D. (For integrability conditions of a generalised distribution, see e.g. [52, 56].)

Consider a local coordinate neighbourhood U in M, with coordinates q^i ($i = 1, \ldots, n$), which is also a trivialising neighbourhood for both vector bundles ν and π. Let σ_a ($a = 1, \ldots, k$), respectively p_A ($A = 1, \ldots, \ell$), represent a local basis of sections of ν, respectively π, defined on U. We then have

$$\sigma_a * \sigma_a = c^d_{ab} \sigma_d\ ,$$

for some functions $c_{ab}^d \in C^\infty(U)$. Putting $\rho(\sigma_a) = \gamma_a^i(\partial/\partial q^i)$, the condition (8.16) yields the following relation

$$c_{ab}^d \gamma_d^i = \gamma_a^j \frac{\partial \gamma_b^i}{\partial q^j} - \gamma_b^j \frac{\partial \gamma_a^i}{\partial q^j} \ ,$$

for all a, b, i. Given a ρ-connection on π, let

$$\nabla_{\sigma_a} p_A = \Gamma_{aA}^B p_B \ .$$

Denoting the components of the curvature R with respect to the chosen local bases of sections by R_{abA}^B, i.e. $R(\sigma_a, \sigma_b; p_A) = R_{abA}^B p_B$, a straightforward computation reveals that

$$R_{abA}^B = \gamma_a^i \frac{\partial \Gamma_{bA}^B}{\partial q^i} - \gamma_b^i \frac{\partial \Gamma_{aA}^B}{\partial q^i} + \Gamma_{aC}^B \Gamma_{bA}^C - \Gamma_{bC}^B \Gamma_{aA}^C - c_{ab}^d \Gamma_{dA}^B \ . \tag{8.17}$$

Always under the assumption that (8.16) is satisfied, we will establish a link between the curvature of a ρ-connection h on $\pi : E \to M$ and the (lack of) involutivity of the (generalised) distribution $Q = \mathrm{im}(h)$. Recalling that for any $\sigma \in \Gamma(\nu)$, $\sigma^h \in \mathfrak{X}(E)$ denotes its h-lift (cf. Section 1), we have the following useful property.

Lemma 8.2. *For any $\sigma_1, \sigma_2 \in \Gamma(\nu)$,*

$$[\sigma_1^h, \sigma_2^h](e) - (\sigma_1 * \sigma_2)^h(e) \in V_e\pi,$$

holds for all $e \in E$.

Proof. From the fact that for each $\sigma \in \Gamma(\nu)$, σ^h and $\rho \circ \sigma$ are π-related vector fields, it follows that $[\sigma_1^h, \sigma_2^h]$ and $[\rho(\sigma_1), \rho(\sigma_2)]$ are also π-related. Taking into account (8.16) we then easily find that

$$\begin{aligned} T\pi \circ \left([\sigma_1^h, \sigma_2^h] - (\sigma_1 * \sigma_2)^h\right) &= T\pi \circ [\sigma_1^h, \sigma_2^h] - \rho(\sigma_1 * \sigma_2) \circ \pi \\ &= \big([\rho(\sigma_1), \rho(\sigma_2)] - \rho(\sigma_1 * \sigma_2)\big) \circ \pi \\ &= 0 \ , \end{aligned}$$

from which the result follows. □

We now come to the following important result which tells us that the curvature R can indeed be seen as a measure for the 'non-involutivity' of the (generalised) distribution Q determined by a ρ-connection. (Recall that ι_e denotes the canonical identification between $E_{\pi(e)}$ and $V_e\pi$).

Theorem 8.3. *For any $\sigma_1, \sigma_2 \in \Gamma(\nu)$ we have*

$$\iota_e^{-1}\Big([\sigma_1^h, \sigma_2^h](e) - (\sigma_1 * \sigma_2)^h(e)\Big) = R(\sigma_1, \sigma_2; \psi)(x)\ , \qquad (8.18)$$

for each $e \in E$ for which the left-hand side is defined, and where $x = \pi(e)$ and ψ is any section of π such that $\psi(x) = e$.

Proof. First of all, note that the left-hand side of (8.18) makes sense in view of the previous lemma, and that the 'tensorial character' of R implies that the right-hand side does not depend on the choice of the section ψ for which $\psi(x) = e$. Secondly, using the properties of the h-lift of sections it is not difficult to check that $[\sigma_1^h, \sigma_2^h] - (\sigma_1 * \sigma_2)^h$ is $C^\infty(M)$-linear in both σ_1 and σ_2. Since we already know that the same is true for $R(\sigma_1, \sigma_2; \psi)$, it suffices to verify (8.18) on a local basis of sections $(\sigma_a)_{a=1,\ldots,k}$ of $\Gamma(\nu)$, defined on a suitable coordinate neighbourhood U of $x = \pi(e)$, for some chosen point $e \in E$. There is no loss of generality by assuming that U is also a trivialising neighbourhood for π, and denote the corresponding bundle coordinates on E by (q^i, y^A). In particular, let the coordinates of the point e be given by (q_0^i, y_0^A).

Using the notations introduced above, we find after a rather tedious, but straightforward computation, that $[\sigma_a^h, \sigma_b^h](e) - (\sigma_a * \sigma_b)^h(e)$ equals

$$\left(\gamma_a^i \frac{\partial \Gamma_{bB}^A}{\partial q^i} - \gamma_b^i \frac{\partial \Gamma_{aB}^A}{\partial q^i} + \Gamma_{bC}^A \Gamma_{aB}^C - \Gamma_{bC}^A \Gamma_{aB}^C - c_{ab}^d \Gamma_{dB}^A\right)_{q_0} y_0^B \frac{\partial}{\partial y^A}_{|e}\ .$$

The result now easily follows when comparing the right-hand side with the expression (8.17) for the local components of R, and bearing in mind that ι_e maps each $(q_0^i, y_0^A, 0, w^A) \in V_e\pi$ onto $(q_0^i, w^A) \in E_x$. □

Example 8.4. If (N, ν) is a Lie algebroid over M with anchor map ρ, we recover the notion of curvature defined, for instance, in [15].

Assume again $\nu : N \to M$ is a pre-Lie algebroid, i.e. that $\Gamma(\nu)$ admits an algebra structure, with a skew-symmetric product $*$ satisfying (8.14). We do not require, however, that ρ is an algebra homomorphism. Consider now a ρ-connection on ν, with associated derivative operator ∇ (i.e. we take $E = N$ and $\pi = \nu$). We can then define a mapping $T : \Gamma(\nu) \times \Gamma(\nu) \to \Gamma(\nu)$ given by

$$T(\sigma_1, \sigma_2) = \nabla_{\sigma_1}\sigma_2 - \nabla_{\sigma_2}\sigma_1 - \sigma_1 * \sigma_2\ . \qquad (8.19)$$

It is not difficult to check that T, which may be called the *torsion* of the given ρ-connection, is a $C^\infty(M)$-bilinear and skew-symmetric mapping. Let $(\sigma_a)_{a=1,\dots,k}$ represent a local basis of sections of ν such that

$$\nabla_{\sigma_a}\sigma_b = \Gamma^d_{ab}\sigma_d \quad \text{and} \quad \sigma_a * \sigma_b = c^d_{ab}\sigma_d \ .$$

It then readily follows that

$$T(\sigma_a,\sigma_b) = (\Gamma^d_{ab} - \Gamma^d_{ba} - c^d_{ab})\sigma_d \ .$$

Example 8.5. Let A be a type $(1,1)$-tensor field on M and consider a pseudo-connection on τ_M with fundamental tensor field A (cf. Section 7, Example (iv)). Here we have $N = TM$, $\nu = \tau_M$ and for the product $*$ we may take the bracket $[\ ,\ ]_A$ on $\mathfrak{X}(M)$, defined by (7.13). This bracket satisfies (8.14), but in general will not be a Lie bracket (since A need not be a Nijenhuis tensor). The notion of torsion, defined by (8.19), corresponds to the one encountered in treatments of pseudo-connections (see e.g. [12, 59]).

9 The Ambrose-Singer Theorem for connections over a Lie algebroid

Assume that (ν,ρ) is a Lie algebroid, with bracket operation on $\Gamma(\nu)$ denoted as usual by $[\ ,\ ]$. It is well known that the distribution $D = \operatorname{im}\rho$ is integrable. Let $\pi : P \to M$ denote a principal fibre bundle with structure group G and let h denote a principal ρ-connection. Following the ideas of the preceding section, the following definition is straightforward:

Definition 9.1. The curvature two-form of h is the map $\Omega : \pi^*N \otimes \pi^*N \to \mathfrak{g}$ defined as follows. For arbitrary $(u,s_1),(u,s_2) \in \pi^*N$:

$$T_u\sigma(\Omega_u(s_1,s_2)) = [\sigma_1^h,\sigma_2^h](u) - ([\sigma_1,\sigma_2])^h(u),$$

where $\sigma_i \in \Gamma(\nu)$ such that $\sigma_i\big((\pi(u))\big) = s_i$.

It is easily checked that the right hand side is independent of the choice of σ_1 and σ_2. Note that $\Omega_{ug}(s_1,s_2) = Ad_{g^{-1}} \cdot \Omega_u(s_1,s_2)$ for each $g \in G$.

Remark 9.1. Let $\pi_E : E \to M$ denote a vector bundle, equipped with a ρ-connection h_E and denote the corresponding principal ρ-connection on $FR(E)$ by h. The correspondence between the curvature two-form Ω of h and the curvature tensor R of h_E can be seen from Theorem 8.3. Indeed, for arbitrary

$\sigma \in \Gamma(\nu)$ the vector field σ^h and σ^{h_E} are $\overline{f}_\xi$ related, where $\xi \in \mathbb{R}^\ell$ is arbitrary. Therefore,

$$T\overline{f}_\xi\Big([\sigma_1^h, \sigma_2^h](u) - ([\sigma_1, \sigma_2])^h(u)\Big) = [\sigma_1^{h_E}, \sigma_2^{h_E}](u\xi) - ([\sigma_1, \sigma_2])^{h_E}(u\xi)$$

holds.

The following theorem is a generalisation of the Ambrose-Singer theorem, due to Fernandes [15] for principal Lie algebroid ρ-connection.

Theorem 9.2. *The Lie algebra $\mathfrak{g}(u)$ of the holonomy group $\Phi(u)$ is spanned by*

$$\Big\{\Omega_v(s_1, s_2) \;\Big|\; \textit{for all } v \in H(u) \textit{ and } s_1, s_2 \in N_{\pi(v)}\Big\},$$

and by all elements of the form A with $T_v\sigma(A) = h(v, s)$ with $s \in \ker\rho$.

Proof. The element A, associated with $s \in \ker\rho$ and defined by $T_v\sigma(A) = h(v, s)$, for $\pi(v) = \nu(s)$, will be denoted in the following by $\Upsilon_v(s)$. Without loss of generality, we assume here that M is a connected manifold and that the integrable distribution $\widetilde{D}$ on M equals TM (cf. Section 3). In view of these assumptions, it should be noted that the pull-back of (ν, ρ) along i admits a natural Lie algebroid structure, compatible with $[\ ,\ ]$, and that the curvature two-forms on i^*P and P are mapped onto each other.

Fix a point $u \in P$ and let $\mathfrak{g}'(u)$ denote the subalgebra of $\mathfrak{g}(u)$ generated by all $\Omega_v(s_1, s_2)$ and $\Upsilon_v(s)$, with $v \in H(u)$ and $s_1, s_2, \in N_{\pi(v)}$ and where $s \in \ker\rho_{\pi(v)}$. Since $\Upsilon_{vg} = Ad_{g^{-1}}\Upsilon_v$ for any $g \in G$, we obtain, for any $A \in \mathfrak{g}(u)$, that $[A, \Upsilon_v(s)] \in \mathfrak{g}'(u)$. A similar result holds for the curvature two-form (they transform in the same way). Therefore, we have that $\mathfrak{g}'(u)$ is an ideal in $\mathfrak{g}(u)$, and that, for any $g \in G$, $Ad_{g^{-1}}(\mathfrak{g}'(u)) = \mathfrak{g}'(ug)$.

It is easily seen that the distribution S defined by $S_u = Q_u + T_u\sigma(\mathfrak{g}'(u))$ is involutive. The regularity of S follows from the fact that (i) $T\pi(S) = TM$, (ii) $S_u \cap V_u\pi = T_u\sigma(\mathfrak{g}'(u))$, and from the fact that (iii) $\dim\mathfrak{g}'(u) = \dim\mathfrak{g}'(v)$ for all $u, v \in P$. Properties (i) and (iii) are trivial. By definition we have that $S_u \cap V_u\pi \supset T_u\sigma(\mathfrak{g}'(u))$. On the other hand, assume that $h(u, s) \in V_u\pi$, then $s \in \ker\rho$, or $T_u\sigma(\Upsilon_u(s)) = h(u, s)$, or $h(u, s) \in T_u\sigma(\mathfrak{g}'(u))$. This proves (ii).

Since Q is everywhere contained in S, and S is integrable, we have that $\widetilde{Q}$ is everywhere contained in S, or $\mathfrak{g}(u) < \mathfrak{g}'(u)$. On the other hand, since by definition $\mathfrak{g}'(u) < \mathfrak{g}(u)$, we have that $\mathfrak{g}'(u) = \mathfrak{g}(u)$ (and $S = \widetilde{Q}$). □

III

Optimal control theory

The results presented in this chapter find their origin in recent work on sub-Riemannian geometry (cf. Chapter IV), and are also strongly inspired by some ideas developed in the book by L.S. Pontryagin et al. [47]. The main purpose is to provide a comprehensive and coordinate-free proof of the maximum principle and, at the same time, to present a version of this principle that may be readily accessible to researchers studying the variational approach to dynamical systems subjected to nonholonomic constraints, also called "vakonomic dynamics". Applications of our results can be found, for instance, in sub-Riemannian geometry, where the problem of characterising length-minimising curves (see Chapter IV) can be solved by means of the maximum principle. Also the construction of a Lagrangian and Hamiltonian dynamics on Lie-algebroids (see, for instance, [6, 40, 58] and Section 9) can be solved using the formalism described in the present chapter.

1 A geometric framework for control theory

We now proceed towards the construction of a differential geometric setting for certain control problems. It should be emphasised that, although our formulation is not the most general one, if only for the rather strong smoothness conditions we impose, it occurs to us that there is a sufficiently large and relevant class of control problems that fit within the framework described below (see for instance [54] for a different approach).

Definition 1.1. A *geometric control structure* is a triple (τ, ν, ρ) consisting of (i) a fibre bundle $\tau : M \to \mathbb{R}$ over the real line, where M is called the *event space*, (ii) a fibre bundle $\nu : U \to M$, called the *control space*, and (iii) a bundle morphism $\rho : U \to J^1\tau$ over the identity on M, such that $\tau_{1,0} \circ \rho = \nu$.

In the above, $J^1\tau$ is the first jet bundle of $\tau : M \to \mathbb{R}$, with projections $\tau_1 : J^1\tau \to \mathbb{R}$ and $\tau_{1,0} : J^1\tau \to M$. The typical fibre of M plays the role of configuration space and will be denoted by Q, whereas the typical fibre of the control space U is denoted by C and is called the control domain. It follows from the definition that we have the commutative diagram:

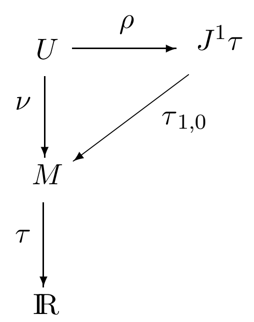

Let c denote a (local) section of $\tau \circ \nu$, i.e. $c : I \subseteq \mathbb{R} \to U$ with $\tau(\nu(c(t))) = t$. With c we can associate a section $\widetilde{c}$ of τ, called the *base section of* c and defined by $\widetilde{c} = \nu \circ c$.

Definition 1.2. A smooth section $c \in \Gamma(\tau \circ \nu)$ is said to be a *smooth control* if $\rho \circ c = j^1\widetilde{c}$, where $\widetilde{c}$ denotes the base section of c and $j^1\widetilde{c}$ its first jet extension. A smooth section $\widetilde{c} \in \Gamma(\tau)$ is called a *smooth controlled section* if $\widetilde{c}$ is the base section of a smooth control c.

Let (t, q^i, u^a) denote an adapted coordinate system on U (i.e. adapted to both fibrations τ and ν). The condition for $c \in \Gamma(\tau \circ \nu)$ to be a smooth control is expressed in coordinates as follows: putting $c(t) = (t, q^j(t), u^a(t))$ and $\rho(t, q^i, u^a) = (t, q^i, \gamma^i(t, q^j, u^a))$ we must have that

$$\dot{q}^i(t) = \gamma^i(t, q^j(t), u^a(t)),$$

for all t. Note that these equations are in agreement with the definition of a control as given in [47, p 56], where $M = \mathbb{R} \times \mathbb{R}^n$ and U is an (open) subset of $M \times \mathbb{R}^k$.

We now discuss the connection between the notion of a smooth control and the notion of an admissible curve (cf. Chapter I). Recall the definition of the the total time derivative operator $\mathbf{T}$ on $J^1\tau$, which is a vector field along $\tau_{1,0}$, i.e. $\mathbf{T} : J^1\tau \to TM$, with $\tau_M \circ \mathbf{T} = \tau_{1,0}$, and which is defined by $\mathbf{T}(j^1_t\widetilde{c}) = (T_t\widetilde{c})(\partial/\partial t)$, for all $j^1_t\widetilde{c} \in J^1\tau$. In coordinates the total time derivative reads, with $j^1_t\widetilde{c} = (t, q^i, \dot{q}^i)$:

$$\mathbf{T}(t, q^i, \dot{q}^i) = \left.\frac{\partial}{\partial t}\right|_{\widetilde{c}(t)} + \dot{q}^i \left.\frac{\partial}{\partial q^i}\right|_{\widetilde{c}(t)} .$$

The total time derivative $\mathbf{T}$ allows us to define an anchored bundle, namely $(\nu, \mathbf{T}\circ\rho)$. Indeed, since $\tau_M \circ \mathbf{T}\circ\rho = \nu$, the map $\mathbf{T}\circ\rho$ determines an anchor.

Note that any section of $\tau\circ\nu$ can be considered as a curve in U. It is easily seen that a smooth control $c : [a,b] \to U$ is a $(\mathbf{T}\circ\rho)$-admissible curve (cf. Chapter I), with base curve $\widetilde{c} = \nu\circ c : [a,b] \to M$. Indeed, we have $\mathbf{T}(\rho(c(t))) = \mathbf{T}(j^1_t\widetilde{c}) = \dot{\widetilde{c}}(t)$, by definition of the total time derivative. On the other hand, any smooth $(\mathbf{T}\circ\rho)$-admissible curve $c : [a,b] \to U$ determines a smooth control if it satisfies $\tau(\nu(c(\bar{t}))) = \bar{t}$. This additional condition is equivalent to saying that c is a section of $\tau\circ\nu$. This is in fact no restriction, since we can always, by considering a simple reparametrisation, transform a smooth $(\mathbf{T}\circ\rho)$-admissible curve into a smooth control. Indeed, consider the projection of the smooth $(\mathbf{T}\circ\rho)$-admissible curve $c(\bar{t})$ on $\mathbb{R}$ by $\tau\circ\nu$, yielding the map $t(\bar{t})$ from $[a,b]$ to $\mathbb{R}$. The equation $\dot{\widetilde{c}}(\bar{t}) = \mathbf{T}(\rho(c(\bar{t})))$ is projected by $T\tau$ onto $\dot{t}(\bar{t}) = 1$. In particular we have that $t(\bar{t}) = \bar{t} - d$ or $\tau(\nu(c(\bar{t}))) = \bar{t} - d$, with d a constant determined by $d = a - \tau(\widetilde{c}(a))$. Thus, if we consider the following reparametrisation of c: put $c' : [a-d, b-d] \to U : t \mapsto c(t+d)$, then $\tau(\nu(c'(t))) = t$. In the remainder of this chapter, we always assume that a smooth $(\mathbf{T}\circ\rho)$-admissible curve is parameterised such that it determines a section of $\tau\circ\nu$.

In Chapter I we defined the composition of smooth admissible curves, which can be applied in particular, using the above correspondence, to smooth controls. In general, a finite composition of smooth controls will simply be called a *control*. Similarly as for $(\mathbf{T}\circ\rho)$-admissible curves, we say that a control $c : [a,b] \to U$ takes a point $x \in M$ to a point $y \in M$ if $x = \nu(c(a))$ and $y = \nu(c(b))$, and we write $x \xrightarrow{c} y$. The set of reachable points from x is denoted by R_x, i.e. $y \in R_x$ iff $x \to y$. Note that the relation $\to$ is an order relation (i.e. reflexive, transitive and non-symmetric) since, if $x \to y$ then $\tau(x) \leq \tau(y)$.

Definition 1.3. A *control* c is a piecewise $(\mathbf{T}\circ\rho)$-admissible curve which is a section of $\tau\circ\nu$.

We use the notation introduced in Chapter I. Let $\mathcal{X} = (X_\ell, \ldots, X_1)$ denote an ordered family of ℓ vector fields in $\mathcal{D} = \{\mathbf{T}\circ\rho\circ\sigma \mid \sigma \in \Gamma(\nu)\}$, with $X_i = \mathbf{T}\circ\rho\circ\sigma_i$, and let $T = (t_\ell, \ldots, t_1) \in \mathbb{R}^\ell$ denote a composite flow parameter. Putting $a_0 = \tau(x)$ and $a_i = a_{i-1} + |t_i|$ for $i = 1, \ldots, \ell$, then we know that the concatenation $\widetilde{c} : [a_0, a_\ell] \to M$ of integral curves associated with $\mathcal{X}$ and T through $x \in M$ is the base curve of the $\pm(\mathbf{T}\circ\rho)$-admissible curve $c = c_\ell \cdot \ldots \cdot c_1$ where $c_i = \sigma_i \circ \widetilde{c}|_{[a_{i-1}, a_i]}$. The curve c is a control (or a

($\mathbf{T} \circ \rho$)-admissible curve) if $T \in \mathbb{R}_+^\ell$, where we define

$$\mathbb{R}_+^\ell = \{(t_\ell, \ldots, t_1) \mid t_i \geq 0\}.$$

In the remaining of this chapter we will say that a curve c constructed this way is *the control associated with* $\mathcal{S}$ *and* T, where $\mathcal{S} = (\sigma_\ell, \ldots, \sigma_1)$ is an ordered family of ℓ sections of ν and T is in $\mathbb{R}_+^\ell$. Note that c takes x to $\Phi_T(x)$ where Φ is the composite flow of $\mathcal{X} = (\mathbf{T} \circ \rho \circ \sigma_\ell, \ldots, \mathbf{T} \circ \rho \circ \sigma_1)$.

In the following, we shall show that, given any control c taking x to y, then there exists an ordered family $\mathcal{S}$ of sections of ν and some $T \in \mathbb{R}_+^\ell$ such that the control associated with $\mathcal{S}$ and T is precisely c. This property is of crucial importance since our notion of 'variation of a control' entirely relies on it (cf. Section 2).

Thus, let $c : I \to U$ be a smooth control with base section $\widetilde{c} := \nu \circ c$. First, assume that the image $c(I)$ is contained in the domain of an adapted coordinate chart V of U with coordinates (t, q^i, u^a). Consider a smooth extension $\hat{c}$ of c, defined on an open interval $\hat{I}$ containing I, i.e. $\hat{c} : \hat{I} \to U$ is a local section of $\tau \circ \nu$ with $\hat{c}(t) = c(t)$ for all $t \in I$. Upon reducing $\hat{I}$ if necessary, we may always assume that $\hat{c}(\hat{I}) \subset V$, and in terms of the adapted coordinates on V we can then write $\hat{c}(t) = (t, q^i(t), u^a(t))$. We can now define a local section σ of ν on the open subset $V' = \nu(V) \cap \tau^{-1}(\hat{I})$ of M as follows: $\sigma(t, q^i) = (t, q^i, u^a(t))$, $\forall (t, q^i) \in V'$. The map $\rho \circ \sigma$ determines a section of $\tau_{1,0}$ satisfying $\rho \circ \sigma(\widetilde{c}(t)) = j^1\widetilde{c}(t)$ for all $t \in I$. This implies that $\widetilde{c}$ is an integral curve of $\mathbf{T} \circ \rho \circ \sigma$. In case the image set $c(I)$ is not fully contained in an adapted coordinate chart, we can always cover the compact set $c(I)$ with a finite number of adapted coordinate charts and choose a subdivision of I such that the image of each subinterval is entirely contained in one of these coordinate charts. The construction above can then be carried out for the restriction of c to each of these subintervals, and it readily follows that the thus obtained family of sections and composite flow parameter $T \in \mathbb{R}_+^\ell$, will induce the control c. As mentioned above, the extension of this proof to the case of general controls is straightforward. Summarising, we have shown that the following property holds.

Proposition 1.1. *Every control c is associated with a finite ordered family $\mathcal{S}$ of (local) sections of ν and with some $T \in \mathbb{R}_+^\ell$.*

Since $(\nu, \mathbf{T} \circ \rho)$ is an anchored bundle, we can consider the leaf L_x through $x \in M$ of the foliation induced by the family $\mathcal{D} = \{\rho \circ \sigma | \sigma \in \Gamma(\nu)\}$, and let R_x denote the set of reachable points from x. We now prove that $\tau(R_x) = [\tau(x), b[$ with $b > \tau(x)$ or $b = \infty$. We have already proven that the projection of the image set of any control c, taking x to some element $y \in M$, is a closed subinterval of $\mathbb{R}$. This implies that $\tau(R_x)$ is an interval in $\mathbb{R}$. In order to prove that it is of the form $[\tau(x), b[$, we will make use of the fact that the family $\mathcal{D}$ is everywhere defined. Indeed, consider any vector field $X \in \mathcal{D}$, with x contained in its domain. Then, if $\{\phi_t\}$ represents the flow of X, the curve $\phi_{t-a}(x)$ is the base curve of a control, and is defined on some interval $[a, a+\epsilon]$, with $\epsilon > 0$. Therefore $b > \tau(x)$. Next, assume that $\tau(R_x)$ is a closed interval, i.e. $\tau(R_x) = [\tau(x), b]$ and consider a control $c : [\tau(x), b] \to M$ taking x to y. Then, fix any vector field X in $\mathcal{D}$ containing y in its domain. Then using similar arguments as above, we find that there exists a control c' defined on the interval $[a, b+\epsilon]$ with $\tilde{c}'(a) = x$, which leads to a contradiction. A straightforward consequence of this property is that $\tau(R_x)$ is a half-open interval. A similar result holds for the inverse anchored bundle $(\nu, -(\mathbf{T} \circ \rho))$ (which can also be considered as the anchored bundle determined by the *inverse* control structure $((-\tau), \nu, \rho)$, where $(-\tau)(x) = -\tau(x)$). Using the above reasoning for the inverse control structure, we have $(-\tau)(R_x^{-1}) = [(-\tau)(x), b[$ or, $\tau(R_x^{-1}) =]-b, \tau(x)]$.

2 The cone of variations

In this section, we consider an arbitrary manifold B, equipped with an everywhere defined family of vector fields $\mathcal{F}$. We introduce a notion of a variation of a concatenation of integral curves of vector fields in the family $\mathcal{F}$. Using the fact that the base curve of a control is such a concatenation, the theory developed in this section will be applicable in control theory. It should be noted that in this section we associate with $\mathcal{F}$ a relation on the points of B which is, in general, different from the relation '$\to$' associated with anchored bundles defined in Chapter I. However, for simplicity we will use the same notation, i.e. $\to$, since in the specific case of a geometric control structure these relations coincide (this is essentially Property 1.1).

2.1 Variations to composite flows

Thus, given an everywhere defined family of vector fields $\mathcal{F}$ on an arbitrary manifold B one can define a quasi-order relation (i.e. a reflexive and transitive relation) R on B as follows: for $x, y \in B$ we put $(x, y) \in R$ if there exists a composite flow Φ, associated with an ordered set $(X_\ell, \ldots, X_1)$, $X_i \in \mathcal{F}$, such that $\Phi_T(x) = y$ for some $T \in \mathbb{R}^\ell_+$. We then also write $x \stackrel{(\Phi,T)}{\longrightarrow} y$ (or simply $x \to y$). As described in Chapter I, the concatenation of integral curves through x, determined by Φ and T, is a piecewise smooth curve $\delta(t)$ satisfying $\dot\delta(t) \in F_{\delta(t)}$ for all t where the derivative exists. As in the previous section, we can then define the set of reachable points from $x \in B$ as the subset $R_x = \{y \in B \mid x \to y\}$. Note that $R_x \neq \emptyset$ for all $x \in B$, since $\mathcal{F}$ is assumed to be everywhere defined.

Let $\widetilde{F}$ denote, as usual, the smallest generalised integrable distribution generated by $\mathcal{F}$, and let us denote the leaf of $\widetilde{F}$ through a given point $x \in B$ by L_x. It is a simple exercise to see that $R_x \subset L_x$. If $\mathcal{F} = -\mathcal{F}$ (i.e. $-X \in \mathcal{F}$ for all $X \in \mathcal{F}$), then the relation R is symmetric. Indeed, if $\Phi_T(x) = y$, with Φ the composite flow determined by $(X_\ell, \ldots, X_1)$ and $T = (t_\ell, \ldots, t_1) \in \mathbb{R}^\ell_+$, then $(\Phi_T)^{-1}(y) = x$ and an elementary computation shows that $(\Phi_T)^{-1} = \Psi_{T^*}$, where Ψ is the composite flow corresponding to the ordered set $(-X_1, \ldots, -X_\ell)$ (where, by assumption, $-X_i \in \mathcal{F}$) and $T^* = (t_1, \ldots, t_\ell)$, i.e. we also have $y \to x$. In this case R determines an equivalence relation for which the equivalence classes are precisely the leafs of the foliation $\widetilde{F}$, i.e. $R_x = L_x$ for any $x \in B$.

Remark 2.1. It should be emphasised here that the everywhere defined family of vector fields $\mathcal{D}$ associated to a control structure (τ, ν, ρ), can never be invariant under multiplication by -1 since, by construction, each vector field belonging to $\mathcal{D}$ is of the form $\mathbf{T} \circ \rho \circ \sigma$ for some $\sigma \in \Gamma(\nu)$ and, therefore, projects onto the fixed vector field $\partial/\partial t$ on $\mathbb{R}$.

We will now investigate the local structure of the set of reachable points R_x for a given $x \in B$. For that purpose we will introduce a special class of variations of a concatenation of integral curves of vector fields in $\mathcal{F}$, connecting x with some $y \in R_x$, such that these variations will lead us from x to points in a neighbourhood of y that also belong to R_x. The following description is merely intended to give a general intuitive idea of the kind of variation we have in mind. We will be more precise later on.

Consider the composite flow Φ corresponding to an ordered set of, say, ℓ vector fields in $\mathcal{F}$, and let $T \in \mathbb{R}^\ell_+$ be such that $\Phi_T(x) = y$. Let $\gamma : [a, b] \to B$ be the concatenation of integral curves induced by Φ and T, with $\gamma(a) = x$ and $\gamma(b) = y$. Roughly speaking, a *variation of* γ consists of a 1-parameter family of piecewise smooth curves $\gamma_\epsilon : I_\epsilon \to B$, defined on the interval $I_\epsilon = [a, b(\epsilon)]$ (i.e. I_ϵ has a variable endpoint) and where ϵ is defined over an open interval containing 0. The following conditions are verified:

1. if $\epsilon = 0$, then $I_\epsilon = [a, b]$ and $\gamma_\epsilon = \gamma$;

2. for all ϵ, $\gamma_\epsilon(a) = x$;
3. for any ϵ we have that γ_ϵ is a concatenation of integral curves of vector fields in $\mathcal{F}$ and if $\epsilon \geq 0$ then $\gamma_\epsilon(b(\epsilon)) \in R_x$ holds;
4. the map $\epsilon \mapsto \gamma_\epsilon(b(\epsilon))$ is a smooth curve through y at $\epsilon = 0$.

The tangent vector to the curve $\epsilon \mapsto \gamma_\epsilon(b(\epsilon))$ at $\epsilon = 0$ is called the *tangent vector to the variation* γ_ϵ. Rather than considering all possible variations satisfying the above conditions, we will mainly deal with a specific class of variations, to be determined below, called *single variations*. It will be shown that the tangent vectors at y to these single variations generate a convex cone in the vector space $\widetilde{F}_y$ (where we recall that $\widetilde{F}$ refers to the smallest integrable distribution generated by $\mathcal{F}$) and, moreover, we will prove that each vector belonging to this cone is in fact a tangent vector to a variation. If we agree to call *dimension of a cone* the dimension of the linear space generated by all vectors belonging to the cone, then the main result of this section can be summarised as follows: if the dimension of the cone of tangent vectors at y to single variations equals the dimension of $\widetilde{F}_y$, say d, then there exists a coordinate chart V on the leaf L_y, with $y \in V$, and coordinate functions denoted by $(q^1, \ldots, q^d)$, such that for any point $z \in V$ for which $q^i(z) \geq 0$ for all $i = 1, \ldots, d$, we have that $z \in R_x$.

Consider again a concatenation of integral curves $\gamma : [a, b] \to B$ associated with the composite flow $\Phi : V \subset \mathbb{R}^\ell \times B \to B$ of an ordered set of ℓ vector fields $(X_\ell, \ldots, X_1)$ in $\mathcal{F}$, and with a given value $T \in \mathbb{R}^\ell_+$ of the corresponding composite flow parameter, such that $\gamma(a) = x$ and $\gamma(b) = \Phi_T(x) = y$ (cf. page 5). We now proceed towards the construction of what will be called a single variation of γ. Let $T = (t_\ell, \ldots, t_1) \in \mathbb{R}^\ell_+$ and put $a_0 = a, a_\ell = b$ and $a_i = a_{i-1} + t_i$ for $i = 1, \ldots, \ell$. Choose an arbitrary point $\tau \in]a_0, a_\ell]$ and let Y be any vector field on B such that $\gamma(\tau)$ belongs to the domain of Y. To fix ideas, let us assume that $a_{i-1} < \tau \leq a_i$. The flow of Y will be denoted by $\{\psi_s\}$ and, as before, $\{\phi^i_s\}$ denotes the flow of X_i. We can then consider the composite flow $\Phi^* : V' \subset \mathbb{R}^{\ell+2} \times B \to B$, associated with the ordered set of $\ell + 2$ vector fields $(X_\ell, \ldots, X_i, Y, X_i, \ldots, X_1)$. Next, define

$$\begin{aligned} T^* : \mathbb{R} \;&\to\; \mathbb{R}^{\ell+2} : \\ \epsilon \;&\mapsto\; T^*(\epsilon) = (t_\ell, \ldots, t_{i+1}, a_i - \tau, \epsilon, \tau - a_{i-1}, t_{i-1}, \ldots, t_1)\,. \end{aligned} \tag{2.1}$$

It is easily seen that there exists an open neighbourhood $\widetilde{I} \subset \mathbb{R}$ of 0, such that x is contained in the domain of the map $\Phi^*_{T^*(\epsilon)}$ for all $\epsilon \in \widetilde{I}$. For each $\epsilon \in \widetilde{I}$, let γ_ϵ denote the concatenation of integral curves through x corresponding to Φ^* and $T^*(\epsilon)$ (note that γ_ϵ is defined on $I_\epsilon = [a, b + |\epsilon| = b(\epsilon)]$). The following sketch visualises the situation for $\tau \in]a_1, a_2]$:

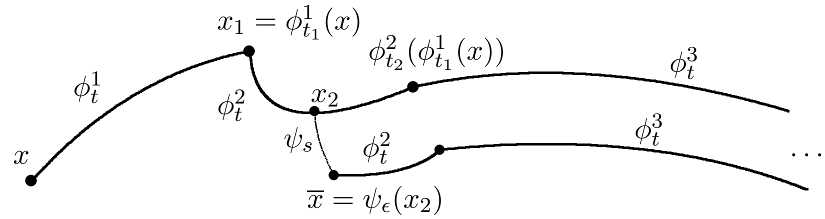

where $x_2 = \phi^2_{\tau - a_1}(x_1)$. The tangent vector to the smooth curve $\epsilon \mapsto \gamma_\epsilon(b(\epsilon)) = \Phi^*_{T^*(\epsilon)}(x)$ at $\epsilon = 0$ is then given by

$$\left.\frac{\partial}{\partial \epsilon}\right|_0 \Phi^*_{T^*(\epsilon)}(x) = T\Phi^{a_\ell}_\tau\Big(Y\big(\gamma(\tau)\big)\Big) \in T_y B\,,$$

where, in order to simplify the notations, we have introduced the mapping $T\Phi^{a_\ell}_\tau : T_{\gamma(\tau)}B \to T_y B$, given by

$$T\Phi^{a_\ell}_\tau(v) = T\phi^\ell_{t_\ell} \circ T\phi^{\ell-1}_{t_{\ell-1}} \circ \cdots \circ T\phi^i_{a_i - \tau}(v)\,, \quad \forall v \in T_{\gamma(\tau)}B\,.$$

Assume now that $Y \in \mathcal{F}$. Then one can see that the 1-parameter family of piecewise smooth curves γ_ϵ satisfies the conditions proposed above for a variation of γ.

Next, suppose we take $Y = -X_i$ and $\tau \in]a_{i-1}, a_i]$ for some $i \in \{1, \dots, \ell\}$, then for $\epsilon > 0$ (but sufficiently small) and for any $t \in]\tau, \tau + \epsilon]$, the tangent vector $\dot{\gamma}_\epsilon(t)$ to the concatenation of integral curves through x, induced by Φ^* and $T^*(\epsilon)$, in general will not belong to $F_{\gamma_\epsilon(t)}$ since $-X_i$ does not have to belong to $\mathcal{F}$. Consequently, if $-X_i \notin \mathcal{F}$, the γ_ϵ resulting from the choice $Y = -X_i$ is, strictly speaking, not a variation in the sense put forward above. However, we can easily remedy the situation by constructing a *reduced* composite flow as follows. Putting $\widehat{T}(\epsilon) = (t_\ell, \dots, t_i - \epsilon, \dots, t_1) \in \mathbb{R}^\ell$, we see that for ϵ sufficiently small, $\Phi_{\widehat{T}(\epsilon)}$ is well-defined in a neighbourhood of x and, moreover, since $\phi^i_{a_i - \tau} \circ \phi^i_{-\epsilon} \circ \phi^i_{\tau - a_{i-1}} = \phi^i_{t_i - \epsilon}$, it follows that $\Phi^*_{T^*(\epsilon)} = \Phi_{\widehat{T}(\epsilon)}$. The concatenation of integral curves determined by Φ and

$\hat{T}(\epsilon)$ does verify the conditions for a variation of γ. The tangent vector at $\epsilon = 0$ to this reduced variation equals

$$\frac{\partial}{\partial\epsilon}\bigg|_0 \Phi^*_{T^*(\epsilon)}(x) = \frac{\partial}{\partial\epsilon}\bigg|_0 \Phi_{\widehat{T}(\epsilon)}(x) = -T\Phi^{a_\ell}_\tau\Big(X_i\big(\gamma(\tau)\big)\Big).$$

We have thus shown that if $\tau \in]a_{i-1}, a_i]$, a variation of the given γ is also determined by the ordered set $(X_\ell, \ldots, X_i, -X_i, X_i, \ldots, X_1)$.

To conclude, if we are given a piecewise smooth curve $\gamma : [a, b] \to B$, with $\gamma(a) = x$, consisting of a concatenation of integral curves determined by the composite flow Φ and composite flow parameter $T = (t_\ell, \ldots, t_1) \in \mathbb{R}^\ell_+$ of an ordered set of vector fields $(X_\ell, \ldots, X_1)$ in $\mathcal{F}$, we introduce the following definition.

Definition 2.1. A *single variation* of γ is a 1-parameter family of piecewise smooth curves $\gamma_\epsilon : [a, b(\epsilon)] \to B$, passing through x, with $\gamma_0 = \gamma$, and such that for each ϵ the corresponding γ_ϵ is the piecewise smooth curve determined by the composite flow Φ^* and composite flow parameter $T^*(\epsilon)$ associated to an ordered set of vector fields of the form $(X_\ell, \ldots, X_i, Y, X_i, \ldots, X_1)$ for some $i \in \{1, \ldots, \ell\}$, with $Y \in \mathcal{F} \cup \{-X_i\}$ and where $T^*(\epsilon)$ is given by (2.1). (We will also briefly refer to γ_ϵ as 'the single variation determined by Φ^* and $T^*(\epsilon)$'.)

For later use we introduce the shorthand notation: $\mathcal{F}_{-X} := \mathcal{F} \cup \{-X_i \,|\, i = 1, \ldots, \ell\}$. Whenever we consider a single variation determined by an ordered set $(X_\ell, \ldots, X_i, Y, X_i, \ldots, X_1)$ for some $Y \in \mathcal{F}_{-X}$, it will always be understood that $Y = -X_j$ can only occur if $i = j$.

Given a single variation γ_ϵ of γ, determined by a composite flow Φ^* and composite flow parameter $T^*(\epsilon)$, one can always obtain a 'new' variation by considering a suitable reparametrisation $\epsilon(\epsilon')$. More precisely, let $\epsilon' \mapsto \epsilon(\epsilon')$ denote a smooth map satisfying $\epsilon(0) = 0$ and $\delta = d\epsilon/d\epsilon'(0) > 0$. Then it is not difficult to verify that Φ^* and $T^*(\epsilon(\epsilon'))$ also determine a variation since $\delta > 0$ implies that, in a neighbourhood of 0, $\mathrm{sgn}(\epsilon) = \mathrm{sgn}(\epsilon')$. The tangent vector to the curve $\epsilon' \mapsto \Phi^*_{T^*(\epsilon(\epsilon'))}(x)$ at $\epsilon' = 0$ equals $\delta\ T\Phi^{a_\ell}_\tau(Y(\gamma(\tau)))$. From this one can easily derive that any positive multiple of a tangent vector to a single variation is again a tangent vector to a (not necessarily single) variation. Note that if $\delta Y \in \mathcal{F}_{-X}$, then $\delta\ T\Phi^{a_\ell}_\tau(Y(\gamma(\tau)))$ is again a tangent vector to a single variation. In general, however, if $Y \in \mathcal{F}_{-X}$, the vector field δY need not be contained in $\mathcal{F}_{-X}$. All this naturally leads to the following definition.

Definition 2.2. Let $y \in R_x$ and fix a composite flow Φ, corresponding to an ordered set $(X_\ell, \ldots, X_1)$ of vector fields in $\mathcal{F}$, such that $\Phi_T(x) = y$ for some $T \in \mathbb{R}^\ell_+$. The *variational cone at* y associated to Φ and T, is the cone $C_y R_x(\Phi, T)$ in $\widetilde{F}_y \subset T_y B$ consisting of all finite linear combinations, with nonnegative coefficients, of tangent vectors to single variations, i.e.

$$C_y R_x(\Phi, T) = \left\{ \sum_{i=1}^{s} \delta^i T\Phi^{a_\ell}_{\tau^i}(Y_i(\gamma(\tau^i))) \;\middle|\; Y_i \in \mathcal{F}_{-X}, \delta^i \geq 0, \tau^i \in]a_0, a_\ell], s \in \mathbb{N} \right\}.$$

Clearly, $C_y R_x(\Phi, T)$ is a cone since for any $v \in C_y R_x(\Phi, T)$, the half-ray λv, for $\lambda \geq 0$, also belongs to it. It then follows that $C_y R_x(\Phi, T)$ is a convex set. Indeed if $v, w \in C_y R_x$, then $(1-t)v + tw \in C_y R_x$, for any $t \in [0, 1]$. If no confusion can arise, we will often drop the explicit reference to Φ and T and simply denote the variational cone by $C_y R_x$. It is easily seen that $C_y R_x$ is a convex set. As a consequence of the next lemma it will be seen that any element of $C_y R_x(\Phi, T)$ can be regarded as a tangent vector to a variation of the piecewise smooth curve through x associated with Φ and T.

The proof of the following result is quite technical. As before, we start from a given piecewise smooth curve $\gamma : [a_0, a_\ell] \to B$, with $\gamma(a_0) = x$, associated to the composite flow of an ordered set of ℓ vector fields $(X_\ell, \ldots, X_1)$ in $\mathcal{F}$, and a fixed value T of the composite flow parameter.

Lemma 2.2. *Consider any finite number of (say, s) tangent vectors* v_i *to single variations of* γ*, namely* $v_i = T\Phi^{a_\ell}_{\tau^i}(Y_i(\gamma(\tau^i)))$*, with* $Y_i \in \mathcal{F}_{-X}$ *and* $\tau^i \in]a_0, a_\ell]$ *for* $i = 1, \ldots, s$*. Then, there exists a composite flow* Φ^* *associated to an ordered set of* $\ell + 2s$ *vector fields formed by the* X_i*'s and the* Y_j*'s, and a smooth mapping* $T^* : \mathbb{R}^s \to \mathbb{R}^{\ell+2s}$, $(\epsilon^1, \ldots, \epsilon^s) \mapsto T^*(\epsilon^1, \ldots, \epsilon^s)$ *such that:*

1. $\Phi^*_{T^*(0)} = \Phi_T$;
2. x *belongs to the domain of* $\Phi^*_{T^*(\epsilon^1,\ldots,\epsilon^s)}$ *for all* $(\epsilon^1, \ldots, \epsilon^s)$ *in some open neighbourhood* $I^{(s)}$ *of* $(0, \ldots, 0) \in \mathbb{R}^s$;
3. *for each fixed* $(\epsilon^1, \ldots, \epsilon^s) \in I^{(s)}$*, with* $\epsilon^i > 0$ *for all* i*, the tangent vector to the concatenation of integral curves through* x *determined by* Φ^* *and* $T^*(\epsilon^1, \ldots, \epsilon^s)$ *is everywhere contained in* $\mathcal{F}$ *(possibly after a 'reduction' of* Φ^* *in the sense described above) such that, in particular,*$\Phi^*_{T^*(\epsilon^1,\ldots,\epsilon^s)}(x) \in R_x$;

4. *the tangent vector at* $\epsilon = 0$ *to the curve* $\epsilon \mapsto \Phi^*_{T^*(\epsilon\delta^1,\dots,\epsilon\delta^s)}(x)$ *equals* $\delta^i v_i$*, for all* $\delta^i \in \mathbb{R}$ *(and where the curve is defined on a sufficiently small interval such that* $(\epsilon\delta^1,\dots,\epsilon\delta^s) \in I^{(s)}$*).*

Proof. Without loss of generality, we may assume that the instants τ^i are ordered in such a way that $\tau^1 \leq \tau^2 \leq \cdots \leq \tau^s$. Moreover, whenever some of the successive τ^i coincide, the ordering should be such that from the corresponding vector fields Y_i, those that do not belong to $\mathcal{F}$ always precede those that do belong to $\mathcal{F}$. More precisely, assume $\tau^i = \cdots = \tau^j$ with $1 \leq i < j \leq s$, and let $\tau^i \in]a_{r-1}, a_r]$ for some $r \in \{1,\dots,\ell\}$. Then we require that if $Y_k = -X_r$ for some $k \in \{i,\dots,j\}$, and $-X_r \notin \mathcal{F}$, we have $k < k'$ for all those $k' \in \{i,\dots,j\}$ for which $Y_{k'} \in \mathcal{F}$. Such an arrangement can always be achieved by simply taking a suitable permutation of the ordered set $(Y_i,\dots,Y_j)$, if necessary. Henceforth, we will always assume, for simplicity, that the Y_i's already appear in the correct ordering.

For $j = 1,\dots,\ell$, let s_j denote the maximum of the set $\{i \mid \tau^i \in]a_{j-1}, a_j]\}$ and put $s_j = s_{j-1}$ if $\{i \mid \tau^i \in]a_{j-1}, a_j]\} = \emptyset$ and $s_0 = 0$. The number of τ^i's belonging to the j-th subinterval is then given by $n_j = s_j - s_{j-1}$. Let $\{\psi^i_s\}$ denote the flow of Y_i (and, as before, $\{\phi^j_s\}$ refers to the flow of X_j). Using the 'star' notation introduced in Chapter I (page 4) to denote the composition of composite flows, we now consider for each $j \in \{1,\dots,\ell\}$, the composite flow $\Phi^*_j : \mathbb{R}^{1+2n_j} \times B \to B$ defined by

$$\Phi^*_j = \begin{cases} \phi^j \star \psi^{s_j} \star \phi^j \star \psi^{s_j-1} \star \cdots \star \phi^j \star \psi^{s_{j-1}+1} \star \phi^j & \text{if} \quad n_j > 0, \\ \phi^j & \text{if} \quad n_j = 0, \end{cases}$$

and a mapping $T^*_j : \mathbb{R}^{n_j} \mapsto \mathbb{R}^{1+2n_j}$ (where it is understood that if $n_j = 0$, then $T^*_j \in \mathbb{R}$), defined by

$$T^*_j(\epsilon^{s_{j-1}+1},\dots,\epsilon^{s_j}) = \begin{cases} (a_j - \tau^{s_j}, \epsilon^{s_j}, \tau^{s_j} - \tau^{s_j-1}, \epsilon^{s_j-1}, \dots, \\ \qquad \tau^{s_{j-1}+2} - \tau^{s_{j-1}+1}, \epsilon^{s_{j-1}+1}, \tau^{s_{j-1}+1} - a_{j-1}) \\ \text{if } n_j > 0, \\ (a_j - a_{j-1}) \text{ if } n_j = 0. \end{cases}$$

Next, by Φ^* we denote the 'composition' of all the composite flows Φ^*_j, i.e. $\Phi^* = \Phi^*_\ell \star \cdots \star \Phi^*_1$. Then, Φ^* itself is a composite flow which can be evaluated at points of $\mathbb{R}^{\ell+2s} \times B$. If we consider the mapping $T^* : \mathbb{R}^s \to \mathbb{R}^{\ell+2s}$ given by

$$\begin{aligned} T^*(\epsilon^1,\ldots,\epsilon^s) &= T^*_\ell(\epsilon^{s_{\ell-1}+1},\ldots,\epsilon^{s_\ell})\star\cdots\star T^*_1(\epsilon^1,\ldots,\epsilon^{s_1}),\\ &= \big(T^*_\ell(\epsilon^{s_{\ell-1}+1},\ldots,\epsilon^{s_\ell}),\ldots,T^*_1(\epsilon^1,\ldots,\epsilon^{s_1})\big), \end{aligned}$$

then it is easily seen that $(T^*(0,\ldots,0),x)\in\mathrm{dom}(\Phi^*)$ and $y=\Phi^*_{T^*(0,\ldots,0)}(x)$. This implies, in particular, that there exists an open neighbourhood $I^{(s)}$ of $(0,\ldots,0)\in\mathbb{R}^s$ for which the map $(\epsilon^1,\ldots,\epsilon^s)\mapsto\Phi^*_{T^*(\epsilon^1,\ldots,\epsilon^s)}(x)$ is well defined and, hence, (2) holds. Note that $\Phi^*_{T^*(\epsilon^1,\ldots,\epsilon^s)}(x)$ can still be written as:

$$\Phi^*_{T^*(\epsilon^1,\ldots,\epsilon^s)}(x)=\Big((\Phi^*_\ell)_{T^*_\ell(\epsilon^{s_{\ell-1}+1},\ldots,\epsilon^{s_\ell})}\circ\cdots\circ(\Phi^*_1)_{T^*_1(\epsilon^1,\ldots,\epsilon^{s_1})}\Big)(x).$$

For $s=1$ the definitions of Φ^* and T^* coincide with those encountered in the construction of a single variation. For any $(\delta^1,\ldots,\delta^s)\in\mathbb{R}^s$ and ϵ varying over a sufficiently small interval centred at 0, such that the image of the map $\epsilon\mapsto(\epsilon\delta^1,\ldots,\epsilon\delta^s)$ is contained in $I^{(s)}$, a straightforward, but rather tedious, computation shows that the tangent vector to the curve $\epsilon\mapsto\Phi^*_{T^*(\epsilon\delta^1,\ldots,\epsilon\delta^s)}(x)$, at $\epsilon=0$, equals $\delta^i v_i$, proving (4). It is also easily seen that when putting $\epsilon^i=0$ for all i, we obtain $\Phi^*_{T^*(0)}=\Phi_T$, proving (1).

The proof of (3) we will be provided for a particular, simplified case from which the idea for the general proof can then be easily deduced. Recall that we have chosen the ordering of the τ^i in such a way that, whenever we have a sequence $\tau^i,\ldots,\tau^j,(i<j)$ with $\tau^i=\tau^{i+1}=\cdots=\tau^j$, those vector fields Y_k which belong to the set $\{-X_1,\ldots,-X_\ell\}$ and which are not contained in $\mathcal{F}$, always appear before all the $Y_{k'}\in\mathcal{F}$ in the sequence $Y_i,\ldots,Y_j$. Consider now the particular case where $a_0<\tau^1=\tau^2=\tau^3<a_1<\tau^4$, $Y_1=-X_1(\notin\mathcal{F})$ and $Y_2,Y_3\in\mathcal{F}$. Then,

$$\begin{aligned} (\Phi^*_1)_{T^*_1(\epsilon^1,\epsilon^2,\epsilon^3)} &= \phi^1_{a_1-\tau^1}\circ\psi^3_{\epsilon^3}\circ\psi^2_{\epsilon^2}\circ\phi^1_{-\epsilon^1}\circ\phi^1_{\tau^1-a_0}\\ &= \phi^1_{a_1-\tau^1}\circ\psi^3_{\epsilon^3}\circ\psi^2_{\epsilon^2}\circ\phi^1_{\tau^1-a_0-\epsilon^1}\,. \end{aligned}$$

Therefore, we can define a new composite flow, associated with vector fields in $\mathcal{F}$, by putting $\widehat{\Phi}_1=\phi^1\star\psi^3\star\psi^2\star\phi^1$, and a new composite flow parameter $\widehat{T}_1(\epsilon^1,\epsilon^2,\epsilon^3)=(a_1-\tau^3,\epsilon^3,\epsilon^2,\tau^1-a_0-\epsilon^1)$. Then $(\Phi_1)^*_{T^*_1(\epsilon^1,\epsilon^2,\epsilon^3)}=(\widehat{\Phi}_1)_{\widehat{T}_1(\epsilon^1,\epsilon^2,\epsilon^3)}$ and, for ϵ^1 sufficiently small, the components of $\widehat{T}_1(\epsilon^1,\epsilon^2,\epsilon^3)$ are positive, from which (3) readily follows for the ‘reduced’ composite flow

$\widehat{\Phi}_1$ and the reduced composite flow parameter $\widehat{T}_1$. A similar reasoning can be applied to the general case, which completes the proof of the lemma. □

The previous lemma implies, among other things, that any v in the variational cone $C_yR_x(\Phi, T)$ can be regarded as a tangent vector to a variation of the piecewise smooth curve γ through x, determined by Φ and T. Indeed, by definition of the cone $C_yR_x(\Phi, T)$ we can always write v (in a non-unique way) as $v = \sum_{i=1}^{s} \delta^i v_i$ for a finite number of tangent vectors to single variations $v_i = T\Phi^{a_\ell}_{\tau^i} Y_i(\gamma(\tau^i))$, with $\delta^i > 0$. We can then associate to these v_i a composite flow Φ^*, and a composite flow parameter $T^*(\epsilon^1, \ldots, \epsilon^s)$, as in the above lemma. Then Φ^* and $\epsilon \mapsto T^*(\epsilon\delta^1, \ldots, \epsilon\delta^s)$ determine a one-parameter family of piecewise smooth curves satisfying the conditions for a variation of γ. Moreover, from the above lemma it follows that the tangent vector to the curve $\epsilon \mapsto \Phi^*_{T^*(\epsilon\delta^1,\ldots,\epsilon\delta^s)}(x)$ at $\epsilon = 0$ precisely equals v, which we wanted to demonstrate.

Note that $C_yR_x (= C_yR_x(\Phi, T))$ is entirely contained in $\widetilde{F}_y$. If the dimension of the cone C_yR_x equals $d = \dim \widetilde{F}_y$, then this is equivalent to saying that the the interior of the convex cone C_yR_x, with respect to the standard vector space topology on $\widetilde{F}_y$, is not empty. Indeed, if we have d independent vectors $v^1, \ldots, v^d \in C_yR_x$, then the interior of the simplex in $\widetilde{F}_y$, determined by the ordered set $(0, v^1, \ldots, v^d)$, is contained in C_yR_x. The converse is an immediate consequence of the fact that any (nonempty) open ball in a vector space spans the full space.

2.2 Basic results on the variational cone

The following proposition, says that the variational cones satisfy an "inclusion" property. The proof is a straightforward consequence of the definition of the variational cone.

Proposition 2.3. *Let Φ^1, Φ^2 denote the composite flows of two ordered families of vector fields in $\mathcal{F}$, such that $\Phi^1_{T^1}(x) = y$ and $\Phi^2_{T^2}(y) = z$, for some $x, y, z \in B$, and where T^1, T^2 represent composite flow parameters for Φ^1 and Φ^2, respectively. Then,*

$$T\Phi^2_{T^2}\Big(C_yR_x(\Phi^1, T^1)\Big) \subset C_zR_x\big(\Phi^2 \star \Phi^1, T^2 \star T^1\big) \text{ and}$$
$$C_zR_y(\Phi^2, T^2) \subset C_zR_x\big(\Phi^2 \star \Phi^1, T^2 \star T^1\big).$$

Before stating the main result of this section, we recall that L_y denotes the leaf of $\widetilde{F}$ passing through y (and, of course for $y \in R_x$ we have that $L_y = L_x$). From the theory of integrable distributions, we know that L_y is an immersed submanifold of B whose dimension equals the rank of $\widetilde{F}$ at y.

Theorem 2.4. *Assume that the dimension of the cone C_yR_x equals the dimension d of $\widetilde{F}_y$. Then there exists a coordinate chart V on the leaf L_y, with $y \in V$ and coordinate functions denoted by $(q^1, \ldots, q^d)$, such that for any point $z \in V$ for which $q^i(z) \geq 0$ for all $i = 1, \ldots, d$, we have that $z \in R_x$.*

Proof. By assumption, the linear space spanned by all elements of C_yR_x equals $\widetilde{F}_y$. We can therefore select a basis $\{v_1, \ldots, v_d\}$ of the linear space $\widetilde{F}_y$, with $v_i \in C_yR_x$ for all i. By definition of C_yR_x, each v_i can then be written as a finite linear combination of tangent vectors to single variations, i.e.

$$v_i = \sum_{j=1}^{s_i} \delta^j_{(i)} v^{(i)}_j, \; i = 1, \ldots, d, \tag{2.2}$$

for some $\delta^j_{(i)} \in \mathbb{R}_+$, and where each $v^{(i)}_j$ is of the form

$$v^{(i)}_j = T\Phi^{a_\ell}_{\tau^j_{(i)}} Y^{(i)}_j(\gamma(\tau^j_{(i)}))$$

for some $Y^{(i)}_j \in \mathcal{F}_{-X}$, $\tau^j_{(i)} \in]a_0, a_\ell]$. Although these decompositions are not uniquely determined, for the remainder of the proof we assume that for each of the given basis vectors v_i one particular decomposition has been singled out, i.e. we make a fixed choice for the $v^{(i)}_j$ and for the positive real numbers $\delta^j_{(i)}$ appearing in (2.2). In total we thus have $s = s_1 + \cdots + s_d$ tangent vectors to single variations $v^{(i)}_j$ which, however, need not all be different and/or linearly independent. For convenience, we introduce the following ordering: $(v^{(1)}_1, \ldots, v^{(1)}_{s_1}, v^{(2)}_1, \ldots, v^{(d)}_{s_d})$ and we denote an arbitrary element of this ordered set by w_α, with $\alpha = 1, \ldots, s$ and such that $w_\alpha = v^{(1)}_\alpha$ for $\alpha = 1, \ldots, s_1$, $w_\alpha = v^{(2)}_{\alpha - s_1}$ for $\alpha = s_1 + 1, \ldots, s_1 + s_2$, *etc.* ... . According to Lemma 2.2 we can associate to the s tangent vectors w_α to single variations, a composite flow Φ^* and a map $T^* : \mathbb{R}^s \to \mathbb{R}^{\ell+2s}$ such that

1. $\Phi^*_{T^*(0)} = \Phi_T$,

2. $\Phi^*_{T^*(\epsilon^1,\dots,\epsilon^s)}(x) \in R_x$ if all $\epsilon^i \geq 0$ and sufficiently small: more precisely $(\epsilon^1,\dots,\epsilon^s) \in I^{(s)})$,
3. for any fixed $(\delta^1,\dots,\delta^s) \in \mathbb{R}^s$, the tangent vector to the curve $\epsilon \mapsto \Phi^*_{T^*(\epsilon\delta^1,\dots,\epsilon\delta^s)}(x)$ at $\epsilon = 0$ equals $\delta^\alpha w_\alpha$.

With the convention that $s_0 := 0$, we have for any $v \in \widetilde{F}_y$ that

$$v = l^i v_i = \sum_{i=1}^{d} \sum_{j=1}^{s_i} l^i \delta^j_{(i)} w_{s_0+\cdots+s_{i-1}+j} \in \widetilde{F}_y \,.$$

Putting

$$(\delta^1,\dots,\delta^s) := (l^1\delta^1_{(1)},\dots,l^1\delta^{s_1}_{(1)}, l^2\delta^1_{(2)},\dots,l^d\delta^{s_d}_{(d)})\,,$$

we can still write v as

$$v = \sum_{\alpha=1}^{s} \delta^\alpha w_\alpha \,.$$

Since the $\delta^j_{(i)}$ in (2.2) have been fixed, it follows that all the coefficients δ^α, appearing in this decomposition of v, are determined unambiguously. Therefore, the following mapping is well-defined:

$$\widetilde{T} : \widetilde{F}_y \to \mathbb{R}^{\ell+2s},\ v \mapsto \widetilde{T}(v) = T^*(\delta^1,\dots,\delta^s)\,,$$

and, clearly, $\widetilde{T}$ is smooth.

From the properties of Φ^* and T^*, one can further deduce that, on a sufficiently small open neighbourhood W of the origin in the linear space $\widetilde{F}_y$, the mapping given by

$$f : W(\subset \widetilde{F}_y) \to B,\, v \mapsto \Phi^*_{\widetilde{T}(v)}(x)$$

is well-defined and smooth. Moreover, by definition of Φ^*, we have that $f(0) = y$ and $\operatorname{im} f \subset L_y$. Let $j : L_y \hookrightarrow B$ denote the natural inclusion and let us write $\widetilde{f}$ for f, regarded as a mapping from W into L_y, such that the following relation holds: $j \circ \widetilde{f} = f$. Since j is an immersion and f is smooth, it follows that $\widetilde{f} : W(\subset \widetilde{F}_y) \to L_y$ is smooth. In view of the natural identification $T_0\widetilde{F}_y \cong \widetilde{F}_y$, it is easily proven, using property (3) of Φ^* and T^*, that the tangent map of f at 0 satisfies, for any $v = \delta^\alpha w_\alpha \in \widetilde{F}_y$,

$$T_0 f(v) = \left.\frac{d}{d\epsilon}\right|_0 f(\epsilon v) = \left.\frac{d}{d\epsilon}\right|_0 \Phi^*_{T^*(\epsilon\delta^1,\dots,\epsilon\delta^s)}(x) = \delta^\alpha w_\alpha = v\,.$$

This, in turn, implies that $T_0\widetilde{f} : \widetilde{F}_y \to T_y(L_y) \equiv \widetilde{F}_y$ is the identity map and, hence, $\widetilde{f}$ induces a diffeomorphism from an open neighbourhood $\widetilde{W} \subset W$ of $0 \in \widetilde{F}_y$ onto an open neighbourhood V of y in L_y. Consequently, to each point $z \in V$ there corresponds a unique $v \in \widetilde{W}$, with $\widetilde{f}(v) = z$ and, with respect to the basis $\{v^i : \ i = 1, \dots, d\}$ of $\widetilde{F}_y$ chosen above, we can write $v = l^i v_i$. The open set V then becomes the domain of a local coordinate chart on L_y, with coordinate functions q^i $(i = 1, \dots, d)$ defined by putting $q^i(z) = l^i$. Finally, from property (2) of Φ^* and T^* it follows that for those vectors $v = l^i v_i \in \widetilde{W}$ for which all $l^i \geq 0$, we have $z = f(v) \in R_x$ since, in this case, all the coefficients δ^a appearing in the decomposition $v = \delta^\alpha w_\alpha$ are also non-negative. This completes the proof of the theorem. □

Observe that the coordinate vector fields on L_y corresponding to the special chart constructed in the previous theorem are such that (using the notations from the proof of the previous theorem) $\partial/\partial q^i\big|_y = T_0\widetilde{f}(v^i) = v^i$. This observation will be used when we prove the following result, which is a straightforward consequence of Theorem 2.4.

Corollary 2.5. *Assume that C_yR_x has a non empty interior (denoted by $\mathrm{int}(C_yR_x)$) with respect to the topology of $\widetilde{F}_y$. Then, for any curve $\theta : [0,1] \to (L_x =)L_y$ with $\theta(0) = y$ and $0 \neq \dot{\theta}(0) \in \mathrm{int}(C_yR_x)$ there exists an $\epsilon > 0$ such that $\theta(t) \in R_x$ for $0 \leq t \leq \epsilon$.*

Proof. As pointed out before, the fact that C_yR_x has nonempty interior implies that the 'dimension' of the cone equals that of $\widetilde{F}_y$ and so the previous theorem applies. One can always fix a basis $\{v_1, \dots, v_d\}$ in $\widetilde{F}_y$, with $v_i \in C_yR_x$, such that the $\dot{\theta}(0)$ is contained in the interior of the simplex spanned by $(0, v^1, \dots, v^d)$. In particular, this means that $\dot{\theta}(0) = k^i v_i$ with $k^i \in]0,1[$ for $i = 1, \dots, d$. Consider the coordinate chart $(q^1, \dots, q^d)$ on L_y, in a neighbourhood of y, associated with the basis $v^1, \dots, v^d$ as constructed in Theorem 2.4. Note, in passing, that $q^i(y) = 0$ for all i. Now, since $\partial/\partial q^i\big|_y = v_i$ for $i = 1, \dots, d$, and putting $\theta^i = q^i \circ \theta$, we find that

$$\left.\frac{d}{dt}\right|_0 \theta^i(t) = k^i, \ \text{for } i = 1, \dots, d.$$

This implies that for all $i = 1, \dots, d$, $\dot{\theta}^i(0) > 0$ and hence, since $\theta^i(0) = 0$, $\theta^i(t) > 0$ for $0 \leq t \leq \epsilon$ and ϵ sufficiently small, i.e. $q^i(\theta(t)) > 0$ for $i = 1, \dots, d$. According to Theorem 2.4 this implies that $\theta(t) \in R_x$ for all $0 \leq t \leq \epsilon$. □

The following picture sketches the situation described in the above corollary for $d = 3$.

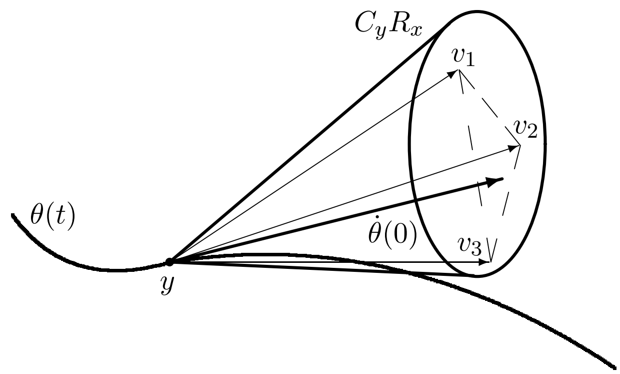

Recall the four conditions we have put forward for the characterisation of a variation γ_ϵ of a concatenation γ of integral curves (cf. Section 2.1 page 64). It will be necessary to enlarge this class of variations when considering "optimal control problems with variable endpoints" (see Section 7). The theorems in the following subsection are proven with similar techniques as the above results.

2.3 Additional theorems for systems with variable initial point

Assume that $i : P \hookrightarrow B$ is an immersed submanifold of B. Let us denote by R_P the set of all points $y \in B$ such that $x \to y$ for some $x \in P$. Note that Corollary 2.5 expresses in some sense how the variational cone C_yR_x is an approximation of the set of reachable points R_x at $y \in R_x$. Similarly, we prove in the next proposition how the set of reachable points R_P at a point $y \in R_P$ can be approximated using (in an indirect way) the variational cone C_yR_x. To simplify the notations, we shall always identify $i(x)$ with x, for arbitrary $x \in P$.

Let C and C' denote two convex cones in a (finite dimensional) linear space $\mathcal{V}$. The convex cone $C * C'$, which is called the cone *generated by C and C'*, is defined as the cone containing all finite linear combinations, with nonnegative coefficients, of vectors in C and C'. The *support plane* of an arbitrary cone C is the subspace of $\mathcal{V}$ spanned by all vectors in C. Assume that $\mathcal{V}$ is equipped with a topological structure, then we can consider the interior of C with respect to the induced topology on the support plane.

Let us fix a composite flow Φ associated with vector fields in $\mathcal{F}$ and a composite flow parameter T, such that $\Phi_T(x) = y$ for some $x \in P$.

Proposition 2.6. *Consider the cone* $K_y = T\Phi_T(T_xP) * C_yR_x$. *Let* $j : L \hookrightarrow B$ *denote an immersed submanifold in* B*, such that* $y \in L$*, and, in addition,*

$$\begin{aligned} &\mathrm{span}(K_y) + T_yL = T_yB \\ &\dim\big(\mathrm{span}(K_y) \cap T_yL\big) \geq 1. \end{aligned}$$

Let $\mathrm{int}(K_y)$ *denote the set of interior points of* K_y*, with respect to its support plane* $\mathrm{span}(K_y)$*. Then, for any* $v \in \mathrm{int}\,(K_y) \cap T_yL$*, there exists a curve* $\theta : [0,1] \to L$*, with* $\theta(0) = y$ *and* $v = \dot{\theta}(0)$*, such that* $j\big(\theta(t)\big) \in R_P$ *for* $0 \leq t \leq \epsilon$*. (Note that in general* $\mathrm{int}\,(K_y) \cap T_yL$ *may be empty.)*

Proof. First of all, consider a finite number of, say s, tangent vectors $v_i \in C_yR_x$ to single variations and s' tangent vectors w_i in T_xP. Then, as proven above, with the vectors v_i one can associate a composite flow Φ^* and a composite flow parameter $T^*(\epsilon^1, \ldots, \epsilon^s)$, verifying the properties stated in Lemma 2.2. Now, for every tangent vector $w_i \in T_xP$ we fix a vector field $W_i \in \mathfrak{X}(P)$, with flow $\{\psi^i_t\}$, such that $W_i(x) = w_i$. The composite flow $\psi^{s'} \star \cdots \star \psi^1$ associated with the ordered set $(W_{s'}, \ldots, W_1)$ is denoted by Ψ. Consider the following map $F : \mathbb{R}^{s'+s} \to B$ defined by

$$(\epsilon^1, \ldots, \epsilon^{s+s'}) \mapsto \Phi_{T(\epsilon^s,\ldots,\epsilon^1)}\Big(\Psi_{(\epsilon^{s+s'},\ldots,\epsilon^{s+1})}(x)\Big).$$

Since $(\epsilon^{s+s'}, \ldots, \epsilon^{s+1})$ is an element of $\mathbb{R}^{s'}$, it can be taken as a composite flow parameter of Ψ. It is also easily seen that this map is well defined on some open interval, say $I^{(s'+s)}$, containing $0 \in \mathbb{R}^{s'+s}$. The following properties hold: (i) if all $\epsilon^i \geq 0$, $i = 1, \ldots, s$ then $F(\epsilon^1, \ldots, \epsilon^{s'+s}) \in R_P$, (ii) $F(0) = y$ and (iii) the tangent vector to the curve $\epsilon \mapsto F(\epsilon\delta^1, \ldots, \epsilon\delta^{s+s'})$ equals

$$\sum_{i=1}^{s}(\delta^i v_i) + \sum_{i=1}^{s'}(\delta^{s+i} w_i),$$

for arbitrary $\delta^i \in \mathbb{R}$.

Assume that $v \in \mathrm{int}\,(K_y) \cap T_yL$ is kept fixed. Denote the dimension of the support plane of K_y by p, i.e. $\dim(\mathrm{span}(K_y)) = p$, and denote the dimension of the manifold L by ℓ. There exists a basis $\{\widetilde{v}_1, \ldots, \widetilde{v}_p\}$ of $\mathrm{span}(K_y)$ such that every $\widetilde{v}_i$ is contained in K_y and v is a linear combination of $\widetilde{v}_i$ with strictly positive coefficients (this is always possible using similar arguments as in the proof of Corollary 2.5). It is readily seen from the definition of K_y, that every $\widetilde{v}_i$ can be written as

$$\widetilde{v}_i = T\Phi_T(w_i) + \sum_{j=1}^{s_i} \delta^j_{(i)} v^{(i)}_j,$$

with all $v^{(i)}_j$ tangent vectors to single variations, $\delta^j_{(i)} \geq 0$ and $w_i \in T_xP$. In total, we thus have $s = s_1 + \cdots + s_p$ tangent vectors to a single variation and $s' = p$ tangent vectors w_i in T_xP. With these tangent vectors we can associate a mapping $F : \mathbb{R}^{s+s'} \to B$ (see the previous paragraph) satisfying the properties (i), (ii) and (iii) formulated above.

By assumption, any element $v' \in \mathrm{span}(K_y)$ can be written as a unique linear combination of the basis vectors $\widetilde{v}_i$: $v' = l^i \widetilde{v}_i$. Using this decomposition, we can associate with v' an element $(\delta^1, \ldots, \delta^{s+s'})$ in $\mathbb{R}^{s+s'}$, as follows,

$$(\delta^1, \ldots, \delta^{s+s'}) = (l^1\delta^1_{(1)}, \ldots, l^1\delta^{s_1}_{(1)}, l^2\delta^1_{(2)}, \ldots, l^p\delta^{s_p}_{(p)}, l^1, \ldots, l^p)$$

We thus obtain the following map:

$$f : W \subset \mathbb{R}^p \cong \mathrm{span}(K_y) \to B : v \mapsto F(\delta^1, \ldots, \delta^{s+s'}),$$

defined on some open interval W containing $0 \in \mathbb{R}^\ell \cong \mathrm{span}(K_y)$. It is easily seen that $T_0 f$ equals the identity map on $\mathrm{span}(K_y)$, implying that f is locally an embedding if $p < n = \dim B$) or a diffeomorphism if $p = n$. This implies that we can fix a coordinate chart on a neighbourhood of $y \in B$, such that, if $(q^1, \ldots, q^n)$ are the coordinate functions, then $q^i(f(l^i\widetilde{v}_i)) = l^i$, for $i = 1, \ldots, p$, and, if $z \in B$ satisfies $q^i(z) = 0$ for $i = p+1, \ldots, n$, then $z \in \mathrm{im}\, f$. In, particular, these coordinates also satisfy: if $q^i(z) \geq 0$ for all $i = 1, \ldots, p$ and $q^i(z) = 0$ for $i = p+1, \ldots, n$, then, using (i) from above, we have $z \in R_P$.

We now use Lemma 2.7 below, which guarantees that, upon replacing L by a neighbourhood of y in L such that L can be regarded as an embedding, the intersection of $\mathrm{im}\, f$ and L is an embedded submanifold which has dimension $\ell + p - n \geq 1$ and for which the tangent space at y equals $T_y(\mathrm{im}\, f) \cap T_y L$. The lemma is proven in the following paragraph. Thus assume that $\theta(t)$ is a curve in the submanifold $\mathrm{im}\, f \cap j(L)$, with $\theta(0) = y$ and with $v = \dot\theta(0)$. Recall that we assumed that $v = \sum_i^p k^i \widetilde{v}_i$ with $k^i > 0$. By construction of the coordinates we have that $\widetilde{v}_i = \partial/\partial q^i|_y$. This implies that

$$q^i(\theta(t)) \geq 0, i = 1, \ldots, p, \text{ and } q^i(\theta(t)) = 0, i = p+1, \ldots, n,$$

for t small enough, say $t \in [0, \epsilon]$. Indeed this follows from $d/dt|_0 q^i(\theta(t)) = k^i > 0$. Therefore, $\theta(t) \in R_P$ for $t \in [0, \epsilon]$, concluding the proof. □

In the proof of the above proposition we have used the following general result.

Lemma 2.7. *Let $g_1 : L_1 \hookrightarrow B$ and $g_2 : L_2 \hookrightarrow B$ denote two embedded submanifolds in a n-dimensional manifold B, with dimension r_1 and r_2, respectively. Assume that there exists a point $x \in g_1(L_1) \cap g_2(L_2)$ such that, in addition, the submanifolds are transversal at x, i.e. $T_xB = Tg_1(T_xL_1) + Tg_2(T_xL_2)$. Then, in a neighbourhood of x, the intersection of L_1 and L_2 is an embedded submanifold of B of dimension $r_1 + r_2 - n$.*

Proof. Let $(q^1, \ldots, q^n)$ denote coordinates on B, defined on $U \subset B$ and centred at x (i.e. the coordinates of x are all zero), such that L_1 is a slice, i.e. if $q^i(y) = 0$ for $i = r_1 + 1, \ldots, n$, then $y \in L_1$. Let $(Q^1, \ldots, Q^{r_2})$ denote coordinates on the embedded submanifold L_2, centred at x and let $q^i = g_2^i(Q^1, \ldots, Q^{r_2})$ denote the coordinate representation of g_2. Then, since g_2 is an embedding, we have that the tangent vectors

$$\frac{\partial g_2^i}{\partial Q^a}(0)\left.\frac{\partial}{\partial q^i}\right|_x, \quad a = 1, \ldots, r_2,$$

are linearly independent. Moreover, we know that

$$\operatorname{span}\left\{\left.\frac{\partial}{\partial q^1}\right|_x, \ldots, \left.\frac{\partial}{\partial q^{r_1}}\right|_x, \frac{\partial g_2^i}{\partial Q^1}(0)\left.\frac{\partial}{\partial q^i}\right|_x, \ldots, \frac{\partial g_2^i}{\partial Q^{r_2}}(0)\left.\frac{\partial}{\partial q^i}\right|_x\right\} = T_xB,$$

with $r_1 + r_2 \geq n$. Fix $n - r_1$ tangent vectors $\partial g_2^i/\partial Q^a(0)\partial/\partial q^i|_x$, such that, together with $\partial/\partial q^i|_x$ ($i = 1, \ldots, r_1$), they form a basis for T_xB. Upon a renumbering of the coordinates we may always take $a = 1, \ldots, n - r_1$. Now, it is easily seen that the $n - r_1$ tangent vectors

$$\sum_{b=1}^{n-r_1} \frac{\partial g_2^{r_1+b}}{\partial Q^a}(0)\left.\frac{\partial}{\partial q^{r_1+b}}\right|_x, \quad a = 1, \ldots, n - r_1$$

are linearly independent. Indeed, if there exists a $(n - r_1)$-tuple $\lambda = (\lambda^1, \ldots, \lambda^{n-r_1}) \in \mathbb{R}^{(n-r_1)}$, with $\lambda \neq 0$ and such that

$$\lambda^a \sum_{b=1}^{n-r_1} \frac{\partial g_2^{r_1+b}}{\partial Q^a}(0)\left.\frac{\partial}{\partial q^{r_1+b}}\right|_x = 0,$$

then

$$\sum_{i=1}^{n} \lambda^a \left(\frac{\partial g_2^i}{\partial Q^a}(0) \left. \frac{\partial}{\partial q^i} \right|_x \right) - \sum_{i=1}^{r_1} \left(\lambda^a \frac{\partial g_2^i}{\partial Q^a}(0) \right) \left. \frac{\partial}{\partial q^i} \right|_x = 0,$$

which is impossible, in view of the assumption that

$$\left\{ \frac{\partial g_2^i}{\partial Q^1} \left. \frac{\partial}{\partial q^i} \right|_x, \ldots, \frac{\partial g_2^i}{\partial Q^{n-r_1}} \left. \frac{\partial}{\partial q^i} \right|_x, \left. \frac{\partial}{\partial q^1} \right|_x, \ldots, \left. \frac{\partial}{\partial q^{r_1}} \right|_x \right\}$$

forms a basis of $T_x B$. We thus obtain that the matrix

$$\left(\frac{\partial g_2^{r_1+b}}{\partial Q^a}(0) \right)_{a,b=1,\ldots,n-r_1}$$

is non-singular. Therefore, applying the implicit function theorem, we have that

$$\begin{aligned} & q^{r_1+a} - g_2^{r_1+a}(Q^1, \ldots, Q^{n-r_1}, Q^{n-r_1+1}, \ldots, Q^{r_2}) = 0, \text{ iff} \\ & Q^a = \hat{g}_2^a(q^{r_1+a}, \ldots, q^n, Q^{(n-r_1)+1}, \ldots, Q^{r_2}), \text{ for } a = 1, \ldots, n - r_1. \end{aligned} \tag{2.3}$$

for some smooth functions $\hat{g}_2^a$ and where, if necessary, we assume that the we restrict the domains of the g_2^i and of the coordinate functions q^i, to some smaller subsets. Consider now the following coordinate transformation in L_2:

$$(Q^1, \ldots, Q^{r_2}) \mapsto (g_2^{r_1+1}(Q), \ldots, g_2^n(Q), Q^{(n-r_1)+1}, \ldots, Q^{r_2}),$$

which is a smooth map from $\mathbb{R}^{r_2}$ to $\mathbb{R}^{r_2}$, whose inverse is given by

$$(Q'^1, \ldots, Q'^{r_2}) \mapsto (\hat{g}_2^1(Q'), \ldots, \hat{g}_2^{n-r_1}(Q'), Q'^{(n-r_1)+1}, \ldots, Q'^{r_2}).$$

We now calculate the coordinate expressions for the map g_2 in this new coordinate system:

$$g_2'^i(Q') = g_2^i(\hat{g}_2^1(Q'), \ldots, \hat{g}_2^{n-r_1}(Q'), Q'^{(n-r_1)+1}, \ldots, Q'^{r_2})$$

In particular, for $i = r_1 + 1, \ldots, n$, then (using 2.3)

$$g_2'^i(Q') = Q'^{i-r_1}.$$

So, if for some point of L_2 we have $Q'^1, \ldots, Q'^{(n-r_1)} = 0$, then we know that the image by g_2 of such a point is contained in L_1, implying that, the r_2+r_1-

n coordinate functions $(Q'^{(n-r_1)+1}, \ldots, Q'^{r_2})$ determine a coordinate system on the intersection of L_1 and L_2. Indeed, assume that $q = (q^1, \ldots, q^n)$ is contained in $U \cap (L_1 \cap L_2)$, then, necessarily, $q^{r_1+1} = \cdots = q^n = 0$ and $q = g_2'(Q')$ for some $Q' \in \mathbb{R}^{r_2}$, implying that $Q'^1 = \cdots = Q'^{(n-r_1)} = 0$. The fact that $U \cap (L_1 \cap L_2)$ is an embedding easily follows from the properties of g_2. Indeed, the composition of the map

$$\mathbb{R}^{r_1+r_2-n} \to \mathbb{R}^{r_2} : \quad (Q'^{r_1-n+1}, \cdots, Q'^{r_2}) \\ \mapsto (0, \ldots, 0, Q'^{r_1-n+1}, \ldots, Q'^{r_2})$$

with the local representation of g_2 in the Q'-coordinates is easily seen to be an embedding. If $n = r_1 + r_2$, the intersection reduces to the point x (where we assumed that, if necessary, the coordinate neighbourhood U is restricted to a smaller open subset). □

A picture representing the above situation, in the case where $B = \mathbb{R}^3$ and $r_1 = 2$ and $r_2 = 2$, may look as follows

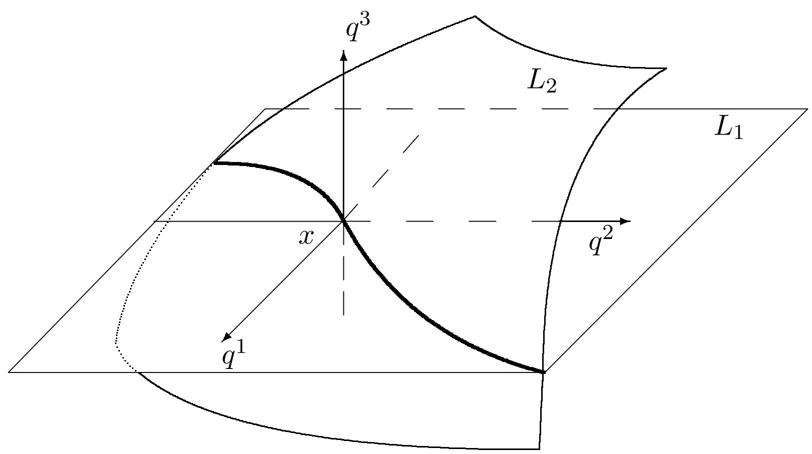

To close this section, we return to the framework of a geometric control structure.

2.4 The vertical variational cone in a geometric control structure

Let (τ, ν, ρ) denote an arbitrary geometric control structure. It is easily seen that the previous definitions and results can be applied, in particular, to the everywhere defined family of vector fields $\mathcal{D} = \{\mathbf{T} \circ \rho \circ \sigma \mid \sigma \in \Gamma(\nu)\}$ on M. Consider a pair $(x, y) \in M \times M$ and a control c taking x to y, i.e. $x \stackrel{c}{\to} y$. In the remainder of this chapter, we shall always assume that we have fixed a finite ordered family $\mathcal{S} = (\sigma_\ell, \ldots, \sigma_1)$ of sections of ν and a composite flow parameter T, inducing the control c in the sense of Proposition 1.1. The composite flow Φ associated with

$$\mathcal{X} = (\mathbf{T} \circ \rho)(\mathcal{S}) = (\mathbf{T} \circ \rho \circ \sigma_\ell, \ldots, \mathbf{T} \circ \rho \circ \sigma_1)$$

allows us to consider the variational cone $C_y R_x(\Phi, T)$.

Since M is fibred over the real line, the kernel of the tangent map $T\tau$ defines a sub-bundle $V\tau = \ker T\tau$ of TM, called the vertical bundle to τ. We will now define a 'sub-cone' of $C_y R_x$ which is vertical in the sense that it is contained in $V_y\tau$. The attention we pay to this "vertical variational cone" will be justified in Section 6, where we prove that it generates the complete variational cone $C_y R_x$.

Definition 2.3. The *vertical variational cone at* y, associated to Φ and T, is given by:

$$V_y R_x(\Phi, T) = \quad \Big\{ \sum_{i=1}^{s} \delta^i T\Phi^{a_\ell}_{\tau^i} (Y_i(\widetilde{c}(\tau^i)) - \dot{\widetilde{c}}(\tau^i)) \mid \delta^i \geq 0, \tau^i \in]a_0, a_\ell], Y_i \in \mathcal{D}, i = 1, \ldots, s \Big\}$$

As for the variational cone, we shall sometimes simply write $V_y R_x$ if there can be no confusion regarding the related Φ and T. The fact that $V_y R_x$ is contained in $V_y\tau$ follows from the fact that any vector field in $\mathcal{D}$ is τ-related to the vector field $\partial/\partial t$ on $\mathbb{R}$, or, that their flows commute with respect to τ. Note that, in the above definition any tangent vector of the form $Y(\widetilde{c}(\tau))$, for $\tau \in]a, b]$, and for some $Y = \mathbf{T} \circ \rho \circ \sigma \in \mathcal{D}$, with $\sigma \in \Gamma(\nu)$, can be written as $\mathbf{T}(\rho(s))$ where $s = \sigma(\widetilde{c}(\tau)) \in U_{\widetilde{c}(\tau)}$.

To fix the ideas we shall now describe an "algorithm" for constructing the vertical variational cone $V_y R_x$. Let c denote a control. Recall the construction in Section 1 (page 62) of an ordered family of sections $\mathcal{S} = (\sigma_\ell, \ldots, \sigma_1)$ of ν and of $T = (t_\ell, \ldots, t_1) \in \mathbb{R}^\ell_+$, such that $\mathcal{S}$ and T induce the control c, i.e. $c = c_\ell \cdot \ldots \cdot c_1$ with $c_i = \sigma_i(\widetilde{c}(t))$ for $t \in [a_{i-1}, a_i]$ and $a_i - a_{i-1} = t_i$. In particular, the sections were defined on local adapted coordinate neighbourhoods, covering $c(I)$, i.e. $\sigma_j(t, q^i) = (t, q^i, u^a(t))$, where $c(t) = (t, q^i(t), u^a(t))$. Let $\Phi = \phi^\ell \star \cdots \star \phi^1$ denote the composite flow of $(\mathbf{T} \circ \rho \circ \sigma_\ell, \ldots, \mathbf{T} \circ \rho \circ \sigma_1)$, where $\{\phi_i\}$ denotes the flow of $X_i = \mathbf{T}(\rho(\sigma_i))$, i.e. locally we have to solve for the integral curves of the vector field

$$X_j = \frac{\partial}{\partial t} + \gamma^i(q^i, u^a(t))\frac{\partial}{\partial q^i},$$

which can be regarded as a (local) time dependent vector field on the typical fibre Q of τ (cf. Remark 2.8). The vertical variational cone $V_y R_x(\Phi, T)$ can then be determined. It is easily seen that the most involved part of this process consists of computing the flows of X_i (i.e. integrating the differential equations: $\dot{q}^i(t) = \gamma^i(t, q^i(t), u^a(t))$.) As we will see in Section 8, in some specific cases (linear autonomous geometric control structures) the vertical variational cone is a linear subspace of $V\tau$, which allows us to "approximate" $V_y R_x$ by iterated lie brackets of time dependent vector fields, which are much easier to compute.

We shall now pay some attention to the case where $\tau : M \to \mathbb{R}$ is a trivial bundle, i.e. $M = \mathbb{R} \times Q$ and τ is the projection onto the first factor. In the following remark we first briefly recall some properties of a time dependent vector field. For further details we refer to [37, p 354].

Remark 2.8. Let B denote an arbitrary manifold. A time dependent vector field $\widetilde{X}$ on B is defined as a smooth map from an open subset Ω of $\mathbb{R} \times B$ into TB, such that $\widetilde{X}(t,x) \in T_x B$, for any (t,x) in the domain of $\widetilde{X}$. We shall sometimes write $\widetilde{X}_t(x)$ for $\widetilde{X}(t,x)$. An integral curve γ of $\widetilde{X}$, is a smooth map from an interval I of $\mathbb{R}$ into M, such that for any $t_0 \in I$, $(t_0, \gamma(t_0))$ is in the domain of $\widetilde{X}$, and such that $\dot{\gamma}(t) = \widetilde{X}(t, \gamma(t))$, for all $t \in I$. Any time dependent vector field on B, determines a vector field on $\mathbb{R} \times B$, which we denote here by X and which is defined by $X = \partial/\partial t + \widetilde{X}$. If we denote the flow of X by $\{\phi_s\}$, then, for fixed t_0, we can consider the following map ψ_{t_0}, defined on some open subset of $\mathbb{R} \times \Omega$ and determined by $(t, \psi^t_{t_0}(x)) = \phi_{t-t_0}(t_0, x)$, for any $(t_0, x) \in \Omega$. This map satisfies the condition that, given any $(t_0, x) \in \Omega$, the curve $t \mapsto \psi^t_{t_0}(x)$ is an integral curve of $\widetilde{X}$. By using the standard properties of the flow $\{\phi_s\}$, we obtain the following properties of $\{\psi^t_{t_0}\}$:

$$\psi^{t_2}_{t_1} \circ \psi^{t_1}_{t_0} = \psi^{t_2}_{t_0},$$
$$(\psi^{t_1}_{t_0})^{-1} = \psi^{t_0}_{t_1}.$$

Let $Y : \mathbb{R} \times B \to TB$ denote a time dependent vector field on B and consider the map $Y_t : B \to TB : x \mapsto Y(t,x)$ for any t. Then $t \mapsto ((\psi^t_{t_0})^* Y_t)(x)$ is a curve in $T_x B$, for which the tangent vector at $t = t_1$ is determined by the following relation, where we use the notation $\dot{Y}_t$ to denote the derivative of the time dependent vector field Y_t with respect to the variable t, while the other variable $x \in B$ is kept fixed, i.e. $\dot{Y}_t(x) = \frac{\partial Y}{\partial t}(t, x)$.

$$\left.\frac{d}{dt}\right|_{t=t_1} \Big((\psi^t_{t_0})^* Y_t\Big)(x) = (\psi^{t_1}_{t_0})^* \Big(\dot{Y}_{t_1} + [\widetilde{X}_{t_1}, Y_{t_1}]\Big)(x).$$

Similar to the notion of composite flow associated to an ordered family of ordinary (time-independent) vector fields, we can consider the composite flow of an ordered family of time dependent vector field, say $(\widetilde{X}_\ell, \ldots, \widetilde{X}_1)$ with $\widetilde{X}_i$ a time dependent vector field on B. Let $\{(\psi^i)^t_{t_0}\}$ denote the flow of $\widetilde{X}_i$, then we can consider, given any $T = (t_\ell, \ldots, t_1) \in \mathbb{R}^\ell$, the following map Ψ_T, defined on some open subset of B:

$$\Psi_T = (\psi^\ell)^{a_\ell}_{a_{\ell-1}} \circ \cdots \circ (\psi^1)^{a_1}_{a_0}, \text{ with } a_i = \sum_{j=0}^{i} t_j.$$

The composite flow Ψ is related to the composite flow Φ of $(X_\ell, \ldots, X_1)$ on $\mathbb{R} \times B$, where $X_i = \partial/\partial t + \widetilde{X}_i$, by the following equality:

$$(t_0 + \cdots + t_\ell, \Psi_T(x)) = \Phi_T(t_0, x).$$

Returning to the case of a control structure (τ, ν, ρ), with $\tau : M = \mathbb{R} \times Q \to \mathbb{R}$ a trivial bundle, every vector field X in $\mathcal{D}$ is of the form $X = \partial/\partial t + \widetilde{X}$, where $\widetilde{X}$ is a time dependent vector field on Q. In particular, the flow $\{\phi_s\}$ of X and the 'flow' $\{\psi^t_{t_0}\}$ of $\widetilde{X}$, satisfy: $T\psi^t_{t_0}(v) = T\phi_{t-t_0}(v)$, given any $v \in T_qQ$ (where we used the identification of T_qQ, with $V_{(t,q)}\tau$ for any $q \in Q$ and $t \in \mathbb{R}$). Let $V_yR_x(\Phi, T)$ denote the vertical variational cone associated to the composite flow Φ of an ordered set $(X_\ell, \ldots, X_1)$, for some $X_i = \partial/\partial t + \widetilde{X}_i \in \mathcal{D}$ and a composite flow parameter $T \in \mathbb{R}^\ell$, where Φ and T induce a control $c : [a, b] \to U$. Let Ψ denote the composite flow of $(\widetilde{X}_\ell, \ldots, \widetilde{X}_1)$. Then, using the above observations, we can alternatively define the vertical variational cone $V_yR_x(\Phi, T)$ as the cone in $T_{q'}Q$ spanned by all tangent vectors of the form:

$$T\Psi^b_t\Big(\mathbf{T}\big(\rho(s)\big) - \mathbf{T}\big(\rho(c(t))\big)\Big),$$

where $c : [a, b] \to U$ denotes the given control and s is an arbitrary element in the fibre of $\nu : U \to M$ at $\widetilde{c}(t)$.

3 Optimal control theory

In this section we give a straightforward application of Corollary 2.5 leading to necessary conditions to be satisfied by an optimal control. We first specify how the notion of optimality of a control can be formulated within the present geometric framework.

Let (τ, ν, ρ) be an arbitrary geometric control structure (with $\tau : M \to \mathbb{R}$, $\nu : U \to M$, $\rho : U \to J^1\tau$, as in Definition 2.1) and let $L \in C^\infty(U)$ denote a smooth function on the control bundle U. If $c : I = [a, b] \to U$ is a control, then the *cost of c with respect to L* is defined by

$$\mathcal{J}(c) = \int_a^b L(c(t))dt.$$

If we put $x = \nu(c(a))$ and $y = \nu(c(b))$, we have, with the notations from Section 1, that $x \stackrel{c}{\to} y$ and, in particular, $y \in R_x$.

Definition 3.1. We say that the control c is *optimal* if $\mathcal{J}(c) \leq \mathcal{J}(c')$ for any other control c' such that $x \stackrel{c'}{\to} y$.

We now specify what we mean by optimal control theory with variable endpoints. Assume that P_i, P_f are two immersed submanifold of M, where, again for notational convenience, we identify points in P_i, P_f with their images in M. Assume that $y \in R_{P_i} \cap P_f$, i.e. there exists a point $x \in P_i$ and a control c taking x to y.

Definition 3.2. We say that the control c is (P_i, P_f)*-optimal* if $\mathcal{J}(c) \leq \mathcal{J}(c')$ for any other control c' taking x to y for some $x \in P_i$ and $y \in P_f$.

For the further discussion, it will be helpful to introduce the following notation:

$$\mathcal{J}_c^{(t_1,t_2)} = \int_{t_1}^{t_2} L(c(t))dt,$$

where $t_1, t_2 \in [a, b]$, with $t_1 \leq t_2$. Note that, in this notation, $\mathcal{J}(c) = \mathcal{J}_c^{(a,b)}$. The function L is sometimes referred to as the *cost function*.

Definition 3.3. A *geometric optimal control structure* (τ, ν, ρ, L) consists of a geometric control structure (τ, ν, ρ) and a cost function L.

We will now show that to every geometric optimal control structure, say (τ, ρ, ν, L), one can associate an *extended geometric control structure*, denoted by $(\overline{\tau}, \overline{\nu}, \overline{\rho})$, in which the cost function is incorporated into the bundle map $\overline{\rho}$. For that purpose, we first introduce the product space $\overline{M} := M \times \mathbb{R}$, the points of which will be denoted by (x, J). For reasons to become clear later on, J will be called the *cost coordinate*. The fibration τ of M over $\mathbb{R}$ induces the fibration $\overline{\tau} : \overline{M} \to \mathbb{R}\,, (x, J) \mapsto \overline{\tau}(x, J) = \tau(x)$. Next, for the extended control bundle we take $\overline{U} = U \times \mathbb{R}$, with projection onto $\overline{M}$ given by $\overline{\nu}(s, J) = (\nu(s), J)$. Finally, we can define a bundle map $\overline{\rho} : \overline{U} \to J^1\overline{\tau}$ as follows: $\overline{\rho}(s, J) = (\rho(s), J, L(s))$, where we have used the canonical identification $J^1\overline{\tau} \cong J^1\tau \times \mathbb{R}^2$, obtained as follows: given any section $\overline{\widetilde{c}}(t) = (\widetilde{c}(t), J(t))$ of $\overline{\tau}$, we map $j^1_t\overline{\widetilde{c}}$ onto $(j^1_t\widetilde{c}, J(t), \dot{J}(t))$. Note that $\overline{\tau}_{1,0}(\overline{\rho}(u, J)) = \overline{\nu}(u, J)$ and, therefore, $(\overline{\tau}, \overline{\nu}, \overline{\rho})$ is indeed a well-defined geometric control structure. Let us denote the projection of $\overline{U}$ onto U by p_U and the projection of $\overline{M}$ onto M by p_M. It is an easy computation to see that p_U is fibred over p_M and determines an anchored bundle morphism between $(\overline{\nu}, \mathbf{T} \circ \overline{\rho})$ and $(\nu, \mathbf{T} \circ \rho)$ (for notational convenience we denote the total time derivative operators on $J^1\overline{\tau}$ and $J^1\tau$ by the same symbol).

Next, we shall prove that any control defined on a geometric optimal control structure (τ, ν, ρ, L) induces a control on the extended structure $(\overline{\tau}, \overline{\nu}, \overline{\rho})$, and vice versa. In one direction the proof easily follows from the fact that, by definition, a control $\overline{c} : [a, b] \to \overline{U}$ in the extended setting is a $\mathbf{T}\circ\overline{\rho}$-admissible curve and since p_U is an anchored bundle morphism, the projection $c = p_U \circ \overline{c}$ is a $\mathbf{T} \circ \rho$-admissible curve. Note that, if $\overline{c}$ takes (x, J_x) to (y, J_y), then $\mathcal{J}(p_U \circ \overline{c}) = J_y - J_x$. This follows from the assumption that $\overline{c}$ is a control: $\dot{J}(t) = L(p_U(\overline{c}(t)))$ is satisfied and, in turn, implies

$$J(b) - J(a) = J_y - J_x = \int_a^b L(p_U(\overline{c}(t)))dt = \mathcal{J}(p_U \circ \overline{c}).$$

Conversely, let $c : I = [a, b] \to U$ be a control in the geometric optimal control structure (τ, ν, ρ, L), with $\nu(c(a)) = x$ and $\nu(c(b)) = y$. We shall construct a control $\overline{c}$ in the associated structure $(\overline{\tau}, \overline{\nu}, \overline{\rho})$ such that for any $J \in \mathbb{R}$ we have $(x, J) \stackrel{\overline{c}}{\to} (y, J + \mathcal{J}_c^{(a,b)})$. More precisely, define the map $\overline{c} : I \to \overline{U}$ by putting

$$\overline{c}(t) = (c(t), J + \mathcal{J}_c^{(a,t)}).$$

It is easily seen that $\overline{\nu}(\overline{c})$ is piecewise smooth. Furthermore, $\overline{c}$ is $\mathbf{T} \circ \overline{\rho}$-admissible. Indeed, the first jet extension of $\overline{\nu} \circ \overline{c}$ equals

$$j^1_t(\overline{\nu} \circ \overline{c}) = j^1_t(\nu \circ c, J + \mathcal{J}_c^{(a,t)}) = (j^1_t(\nu \circ c), J + \mathcal{J}_c^{(a,t)}, L(c(t)))$$

and c being a control, we find that $j^1_t\overline{c} = \overline{\rho}(\overline{c}(t))$. On the other hand, the projections of $\overline{c}(a)$ and $\overline{c}(b)$ onto $\overline{M}$ are given by (x, J) and $(y, J + \mathcal{J}(c))$, respectively. It follows that $(x,J) \stackrel{\overline{c}}{\to} (y, J+\mathcal{J}(c))$ for the extended geometric control problem (and for arbitrary $J \in \mathbb{R}$).

Summarising the preceding discussion, we have proven the following result.

Proposition 3.1. *Let (τ,ν,ρ,L) denote a geometric optimal control structure and $(\overline{\tau},\overline{\nu},\overline{\rho})$ the associated extended geometric control structure. Then for any $x,y \in M$ and $J_x, J_y \in \mathbb{R}$, we have that*

$$x \stackrel{c}{\to} y \text{ and } \mathcal{J}(c) = J_y - J_x \quad \text{iff} \quad (x,J_x) \stackrel{(c,J)}{\to} (y,J_y) \text{ in } (\overline{\tau},\overline{\nu},\overline{\rho}),$$

where $J : [a,b] \to \mathbb{R}$ is given by $J(t) = J_x + \mathcal{J}_c^{(a,t)}$.

The core idea in the book of Pontryagin et al. [47] is to use the (local versions of) Corollary 2.5 and Proposition 2.6 in the extended geometric control structure, in order to obtain necessary conditions for optimal controls that are formulated in the extended setting. These conditions can be rewritten in the initial geometric control structure and turn out to be precisely the conditions of the maximum principle. We are now ready to carry out the first step of this procedure, i.e. we obtain the necessary conditions in the extended geometric control structure. To complete the second step we will use the concept of lifts over a bundle map, which will be treated in Sections 4 and 6.

Remark 3.2. Before proceeding, we first recall some properties and terminology regarding linear spaces and convex cones in a linear space. Let $\mathcal{V}$ be an arbitrary (finite dimensional) linear space. Recall the definition of the annihilator $\mathcal{W}^0$ of a linear subspace $\mathcal{W}$ of $\mathcal{V}$. A hyperplane in $\mathcal{V}$ (i.e. a linear subspace of co-dimension one) can always be defined as the set of all vectors $v \in \mathcal{V}$ satisfying $\langle \eta, v\rangle = 0$ for some (non-zero) co-vector $\eta \in \mathcal{V}^*$. Such a hyperplane divides $\mathcal{V}$ into two 'half-spaces' which are given by the set of all v satisfying $\langle \eta, v\rangle \leq 0$, resp. $\langle \eta, v\rangle \geq 0$, and which are called the 'negative' half-space and the 'positive' half-space, respectively. Let C denote a convex cone in $\mathcal{V}$ (we always assume that the vertex of C is taken in the origin of $\mathcal{V}$). The set of all η for which $\langle \eta, v\rangle \leq 0$, $\forall v \in C$, is called the *dual cone* of C and is denoted by C^*. It is readily seen that C^* is again a convex cone. Indeed, let $\eta, \eta' \in C^*$ and $\lambda, \lambda' \geq 0$, then $\langle \lambda\eta + \lambda'\eta, v\rangle \leq 0$ for all $v \in C$. Intuitively, the cone C^* can be regarded as the set of all hyperplanes with respect to which C is entirely contained in the negative half-space. Note that, if $\mathrm{span}(C) \neq \mathcal{V}$, then $(\mathrm{span}(C))^0 \subset C^*$. Indeed, let $\eta \in (\mathrm{span}(C))^0$, then $\langle \eta, v\rangle = 0$ for all $v \in C$, and therefore, by definition of C^*, the co-vector η is contained in C^*. In the specific case where C is a linear subspace, i.e. $\mathrm{span}(C) = C$, then $C^* = C^0$.

This follows from the fact that, for all $v \in C$, $-v$ also belongs to C and, hence, for any $\eta \in C^*$ we both have $\langle \eta, v\rangle \leq 0$ and $\langle \eta, -v\rangle \leq 0$, which implies that $\langle \eta, v\rangle = 0$ for all $v \in C$, i.e. $\eta \in C^0$.

Assume that C' is another convex cone in $\mathcal{V}$. Recall the definition of $C * C'$, i.e. $C * C'$ is the convex cone generated by all vectors in C and C'. It is easily proven that $(C * C')^* = C^* \cap (C')^*$. We say that C and C' are *separable* if there exists some $\eta \in C^*$ such that $\langle \eta, w\rangle \geq 0$ for all $w \in C'$, i.e. the hyperplane in $\mathcal{V}$ determined by η "separates" the cone C from C', in the sense that C is in the negative half space determined by η, whilst C' is in the positive half space. The following result, which we will use later on, is taken from [47, p 90]. The necessary and sufficient condition for the cones C and C' to be separable is that one of the two following conditions holds: (i) both cones are contained in a hyperplane; (ii) there exists no point which is simultaneously an interior point of C and C', with respect to the subspace topology of the support plane $\mathrm{span}(C)$, resp. $\mathrm{span}(C')$. Let us now denote the closure of a cone C by $\mathrm{cl}(C)$. Then the following properties hold: $C^* = (\mathrm{cl}(C))^*$ and $(C^*)^* = \mathrm{cl}(C)$ (see e.g. [26]). Finally, for any $v \in \mathcal{V}$, the half-ray determined by v, i.e. $\{w \mid w = rv, \forall r \geq 0\}$, is called *the (degenerate) cone generated by* v and is denoted by $C(v)$.

So assume that c is an optimal control, taking x to y. Let $\overline{c}$ denote the control in the extended control structure, taking (x, J_x) to (y, J_y), and let $C_{(y,J_y)}R_{(x,J_x)}$ denote the variational cone in the extended setting (i.e. we have fixed a finite ordered family of sections of $\overline{\nu}$ inducing $\overline{c}$ and with associated composite flow $\overline{\Phi}$). With the notation introduced in the previous remark we also put $C(-\partial/\partial J|_{(y,J_y)})$ the half-ray in $T_{(y,J_y)}\overline{M}$ through the tangent vector $-\partial/\partial J|_{(y,J_y)}$. The proof of the following proposition relies on Corollary 2.5.

Proposition 3.3. *The cones* $C_{(y,J_y)}R_{(x,J_x)}$ *and* $C(-\partial/\partial J|_{(y,J_y)})$ *are separable.*

Proof. We consider two cases. First, assume that $\mathrm{int}(C_{(y,J_y)}R_{(x,J_x)}) = \emptyset$. This is equivalent to saying that $\mathrm{span}(C_{(y,J_y)}R_{(x,J_x)}) \neq T_{(y,J_y)}\overline{M}$ and, in view of the above remark, this implies that $(\mathrm{span}(C_{(y,J_y)}R_{(x,J_x)}))^0 \neq \emptyset$. So, take any element η in the annihilator of $\mathrm{span}(C_{(y,J_y)}R_{(x,J_x)})$. Then, either $\langle \eta, -\partial/\partial J\rangle \geq 0$ or $\langle \eta, -\partial/\partial J\rangle < 0$ holds. In the first case, we have that η separates both cones. In the latter, it suffices to take $-\eta$ instead of η. Next, we assume that $\mathrm{int}(C_{(y,J_y)}R_{(x,J_x)}) \neq \emptyset$. From the above remark, we know that the separability of the cones is equivalent to saying that $(-\partial/\partial J)|_{(y,J_y)} \notin \mathrm{int}(C_{(y,J_y)}R_{(x,J_x)})$. We use an indirect proof. Assume that

$$(-\partial/\partial J)|_{(y,J_y)} \in \mathrm{int}(C_{(y,J_y)}R_{(x,J_x)}).$$

Consider the 'vertical' curve $\theta(t) = (y, J_y - t)$ in $\overline{M}$, defined for $t \in [0,1]$, whose tangent vector at $t = 0$ precisely equals $(-\partial/\partial J)|_{(y,J_y)}$. From Corollary 2.5 it then follows that there exists an $\epsilon > 0$, sufficiently small, such that $\theta(t) \in R_{(x,J_x)}$ for $t \in [0,\epsilon]$. From this, one can deduce that there exists a control $\overline{c}'$ for which $(x, J_x) \xrightarrow{\overline{c}'} (y, J_y - \epsilon)$. In view of previous considerations, this further implies that there exists a control c' on (τ, ν, ρ, L) such that $x \xrightarrow{c'} y$, with cost $\mathcal{J}(c') = \mathcal{J}(c) - \epsilon$ and, hence, $\mathcal{J}(c') < \mathcal{J}(c)$. Since c was assumed to be optimal, this clearly leads to a contradiction. □

We now prove a generalisation of Proposition 3.3, in the case where c is a (P_i, P_f)-optimal control taking $x \in P_i$ to $y \in P_f$. Let $\overline{c}$ denote the corresponding control in the extended setting taking (x, J_x) to (y, J_y) and consider the immersed submanifolds $\overline{P}_i, \overline{P}_f$ in $\overline{M}$, defined by $\overline{P}_i = P_i \times \{J_x\}$ and $\overline{P}_f = P_f \times \{J_y\}$.

Proposition 3.4. *If the control c is (P_i, P_f)-optimal, then the cones $K = T\overline{\Phi}_T(T_{(x,J_x)}\overline{P}_i) * C_{(y,J_y)}R_{(x,J_x)}$ and*

$$K' = C(-(\partial/\partial J)|_{(y,J_y)}) * T_{(y,J_y)}\overline{P}_f$$

in $T_{(y,J_y)}\overline{M}$ are separable.

Proof. Assume that the cones K and K' are not separable. Using Remark 3.2, we know that this is equivalent to saying that $\mathrm{span}(K) + \mathrm{span}(K') = T_{(y,J_y)}\overline{M}$ and that there exists a tangent vector v in the interior of K, with respect to subspace topology of $\mathrm{span}(K)$, which is also contained in the interior of K'. The latter implies that $v_J < 0$, where v_J is the $\partial/\partial J$-component of v. Now, consider the immersed submanifold $g : L \hookrightarrow \overline{M}$, with $L = P_f \times \mathbb{R}$. It is then easily seen that $\mathrm{span}(K') = T_{(y,J_y)}L$. We thus have that $\mathrm{span}(K) + T_{(y,J_y)}L = T_{(y,J_y)}\overline{M}$, and $\dim(\mathrm{span}(K) \cap T_{(y,J_y)}L) \geq 1$, since v, and any positive multiple of it, is contained in $K \cap K'$. From Corollary 2.6, we may conclude that there exists a curve $\overline{\theta}(t) = (\theta(t), \theta^J(t))$ in L, for $t \in [0,1]$, such that $\overline{\theta}(0) = (y, J_y)$, $\dot{\overline{\theta}}(0) = v$ and, for some $\epsilon > 0$, $\overline{\theta}(t)$ is contained in $R_{\overline{P}_i}$ for all $t \in [0,\epsilon]$. Now, since $\dot{\theta}^J(0) < 0$, we have that $\theta^J(t) < \theta^J(0) = J_y$ for t sufficiently small, say $0 < t \leq \epsilon'$. Let $\epsilon'' = \min\{\epsilon, \epsilon'\}$, then, since $\overline{\theta}(t) \in R_{\overline{P}_i}$, there exists a control taking a point $x' \in P_i$ to $\theta(\epsilon'') \in P_f$, with cost $\theta^J(\epsilon'') - J_x < J_y - J_x = \mathcal{J}(c)$. This contradicts the fact that c is assumed to be (P_i, P_f)-optimal. □

In order to relate the previous result to a more familiar formulation of the necessary conditions for an optimal control, in terms of solutions of differential equations, we will use the theory of lifts over anchor maps, developed in Chapter II.

4 The control lift

Let (τ, ν, ρ) denote a geometric control structure. Consider the first-order jet bundle $J^1\nu$ of the bundle $\nu : U \to M$, with associated projections $\nu_1 : J^1\nu \to M$, $\nu_{1,0} : J^1\nu \to U$. Recall that for any two local sections σ and σ' of ν, defined on a neighbourhood of a point $x \in M$, we have that $j^1_x\sigma = j^1_x\sigma' \in J^1\nu$ iff $\sigma(x) = \sigma'(x)$ and $T_x\sigma = T_x\sigma'$ (as linear maps from T_xM into $T_{\sigma(x)}U$). Bearing this in mind, it is easily seen that the following mapping is well-defined:

$$\rho^1 : J^1\nu \to TU,\; j^1_x\sigma \mapsto \rho^1(j^1_x\sigma) = T_x\sigma\Big(\big(\mathbf{T}\circ\rho\big)\big(\sigma(x)\big)\Big). \tag{4.5}$$

Moreover, ρ^1 is a bundle map over the identity on U. In terms of appropriate bundle coordinates (t, q^i, u^a) on U and $(t, q^i, u^a, u^a_0, u^a_i)$ on $J^1\nu$, ρ^1 reads

$$\rho^1(t, q^i, u^a, u^a_0, u^a_i) = \big(t, q^i, u^a, 1, \gamma^j(t, q^i, u^a), u^b_0 + \gamma^j(t, q^i, u^a)u^b_j\big).$$

Clearly $\tau_U \circ \rho^1 = \nu_{1,0}$ and, therefore, $(\nu_{1,0}, \rho^1)$ is an anchored bundle. From the fact that $T\nu\circ\rho^1 = \mathbf{T}\circ\rho\circ\nu_{1,0}$, we can regard $\nu_{1,0}$ as an anchored bundle morphism between $(\nu_{1,0}, \rho^1)$ and $(\nu, \mathbf{T}\circ\rho)$. Thus any ρ^1-admissible curve projects onto a $\mathbf{T}\circ\rho$-admissible curve, i.e. a control. The converse is not true in general. However, if we suitably extend the class of ρ^1-admissible curves, we can prove that any control in (τ, ν, ρ) is the projection of a "piecewise" ρ^1-admissible curve. Therefore, in the sequel, we always assume that a *piecewise* ρ^1*-admissible curve* s is the composition of a finite number of smooth ρ^1-admissible curve, $\mathsf{s}_i : [a_{i-1}, a_i]$ with $i = 1, \ldots, \ell$ such that if $\widetilde{c}_i = \nu_1 \circ \mathsf{s}_i$, then $\widetilde{c}_i(a_i) = \widetilde{c}_{i+1}(a_i)$ for all i (see Equation I-2.2). Roughly speaking, the extension we introduced, consists in allowing ρ^1-admissible curves that may have discontinuities in its base curve c in U. However, the projection onto M is required to be a piecewise smooth curve (which is precisely what we need for c to determine a control).

Lemma 4.1. *The projection onto U of any piecewise ρ^1-admissible curve in $J^1\nu$ is a control, and any control in U can be obtained as the projection of a piecewise ρ^1-admissible curve.*

Proof. The first statement has already been proven above. Conversely, assume that $c : [a_0, a_\ell] \to U$ is a control, with base curve $\widetilde{c} = \nu \circ c$. From Section 1 we know that any control is induced by a finite ordered family $\mathcal{S} = (\sigma_\ell, \ldots, \sigma_1)$ of sections of ν and a composite flow parameter $T = (t_\ell, \ldots, t_1) \in \mathbb{R}^\ell_+$. In particular, the control c then equals the composition $c_\ell \cdot \ldots \cdot c_1$, with $c_i = \sigma_i(\widetilde{c}(t))$ for $t \in [a_{i-1}, a_i]$ and $a_i - a_{i-1} = t_i$. Next, consider the ordered family $\mathcal{S}^1 = (j^1\sigma_\ell, \ldots, j^1\sigma_1)$ of prolongations of sections of ν and define $\mathsf{s}_i(t) = j^1\sigma_i(\widetilde{c}(t))$ for $t \in [a_{i-1}, a_i]$. Then, every s_i is a smooth ρ^1-admissible curve, since $\rho^1(\mathsf{s}_i(t)) = T_{\widetilde{c}(t)}\sigma_i(\dot{\widetilde{c}}(t)) = \dot{c}_i(t)$. Therefore, the composed curve $\mathsf{s} = \mathsf{s}_\ell \cdot \ldots \cdot \mathsf{s}_1$ determines the desired piecewise ρ^1-admissible curve. □

Let σ denote a section of ν. Then the smooth ρ^1-admissible curve $j^1\sigma(\widetilde{c}(t))$, with $\widetilde{c}(t)$ an integral curve of $\mathbf{T} \circ \rho \circ \sigma$, will be called *basic*. The construction in the proof of the previous lemma leads us to the following property.

Proposition 4.2. *Every piecewise ρ^1-admissible curve can be written as the composition of a finite number of smooth basic ρ^1-admissible curves.*

Proof. In order to prove that any (piecewise) ρ^1-admissible curve can be written as a concatenation of smooth basic ρ^1-admissible curve, we shall prove that any smooth ρ^1-admissible curve s whose image is entirely contained in a coordinate chart, is of that form. The general result then follows by a similar argument as the one applied in Section 1 (page 62) when proving that the base curve of any control is a concatenation of integral curves of vector fields in $\mathcal{D}$. So, assume s can be written in bundle adapted coordinates as $\mathsf{s}(t) = (t, q^i(t), u^a(t), u^a_0(t), u^a_i(t))$ for all $t \in I = [a, b]$. Since s is ρ^1-admissible, we then have that

$$\dot{u}^a(t) = u^a_0(t) + u^a_i(t)\dot{q}^i(t) \quad \text{and} \quad \dot{q}^i(t) = \gamma^i(t, q^i(t), u^a(t)).$$

Consider now a smooth extension $\hat{\mathsf{s}}(t) = (t, \hat{q}^i(t), \hat{u}^a(t), \hat{u}^a_0(t), \hat{u}^a_i(t))$ of s, defined on an open interval $\hat{I}$ containing I, such that $\operatorname{im}\hat{\mathsf{s}}$ is still contained in the same coordinate chart, with $\hat{\mathsf{s}}(t) = \mathsf{s}(t)$ for all $t \in I$. Next, we can construct a local section σ of ν, defined on $\tau^{-1}(\widetilde{I})$, as follows: $\sigma(t, q) =$

$(t, q, \sigma^a(t, q))$, with $\sigma^a(t, q) = \hat{u}^a(t) + \hat{u}^a_i(t)(q^i - \hat{q}^i(t))$. For each fixed $t \in I$ we then find that

$$\begin{aligned} \sigma^a(t, q^i(t)) &= u^a(t) \\ \frac{\partial \sigma^a}{\partial t}(t, q^i(t)) &= \dot{u}^a(t) - u^a_i(t)\dot{q}^i(t) = u^a_0(t), \\ \frac{\partial \sigma^a}{\partial q^i}(t, q^i(t)) &= u^a_i(t), \end{aligned}$$

and, hence, we have that $j^1\sigma(t, q(t)) = \mathsf{s}(t)$ for all $t \in I$, which is precisely what we wanted to prove. □

The properties developed above on ρ^1-admissible curves provide us with a complete understanding of how they are related to controls. We now define a natural ρ^1-lift on the bundle $\nu^*TM \to U$. In the following we shall frequently make use of the natural identification $T(\nu^*TM) \cong TU \times_{TM} TTM$, without mentioning it explicitly. We further denote by $\mathfrak{s} : TTM \to TTM$ the canonical involution on TTM. The latter is characterised by the relations $T\tau_M \circ \mathfrak{s} = \tau_{TM}$ and $\tau_{TM} \circ \mathfrak{s} = T\tau_M$.

Remark 4.3. Recall that, given an arbitrary manifold B with local coordinates (q^i), and denoting the natural bundle coordinates on TB and TTB by (q^i, v^i) and $(q^i, v^i, \dot{q}^i, \dot{v}^i)$, respectively, then the canonical involution $\mathfrak{s}$ on TTB reads $\mathfrak{s}\,(q^i, v^i, \dot{q}^i, \dot{v}^i) = (q^i, \dot{q}^i, v^i, \dot{v}^i)$.

Finally, we denote the projections onto the first and second factor of the pull-back bundle $\nu^*TM = U \times_M TM$ by:

$$p_1 : \nu^*TM \to U, \text{ and } \quad p_2 : \nu^*TM \to TM,$$

respectively. Given a geometric control structure (τ, ν, ρ), consider the associated anchored bundle $(\nu_{1,0}, \rho^1)$ with anchor map ρ^1 given by (4.5).

Proposition 4.4. *The map $h^c : p_1^* J^1\nu \to T(\nu^*TM)$, defined by*

$$h^c\big((\sigma(x), v), j^1_x\sigma\big) = \Big(\rho^1(j^1_x\sigma), \mathfrak{s}\left(T(\mathbf{T} \circ \rho)(T_x\sigma(v))\right)\Big),$$

for any $x \in M$, $\sigma \in \Gamma(\nu)$ and $v \in T_xM$, is a lift over ρ^1.

Proof. We first calculate the local coordinate expression of h^c. For that purpose, consider bundle adapted coordinates (t, q^i, u^a) and (t, q^i, v^0, v^j) on

U and TM, respectively. Take $x=(t,q^i)\in M$, $v=(t,q^i,v^0,v^j)\in T_xM$ and $\sigma\in\Gamma(\nu)$ such that $j^1_x\sigma=(t,q^i,u^a,u^a_0,u^a_i)$, then:

$$\begin{aligned}\mathfrak{s}\Big(T\big(\mathbf{T}\circ\rho\big)\big(T_x\sigma(v)\big)\Big) &= \left.\frac{\partial}{\partial t}\right|_v + \gamma^i(t,q^j,u^a)\left.\frac{\partial}{\partial q^i}\right|_v \\ &+\left(v^0\frac{\partial\gamma^k}{\partial t}(t,q^j,u^a)+v^i\frac{\partial\gamma^k}{\partial q^i}(t,q^j,u^a)+(v^0u^b_0+v^iu^b_i)\frac{\partial\gamma^k}{\partial u^b}(t,q^j,u^a)\right)\left.\frac{\partial}{\partial v^k}\right|_v.\end{aligned}$$

Next, using the properties of the canonical involution operator $\mathfrak{s}$, and taking into account (4.5), it is easily seen that

$$T\nu\big(\rho^1(j^1_x\sigma)\big)=T\tau_M\Big(\mathfrak{s}\left(T(\mathbf{T}\circ\rho)(T_x\sigma(v))\right)\Big)\in T_xM$$

which proves indeed that $\operatorname{im}h^c\subset T(\nu^*TM)$.

From its definition it readily follows that h^c is a bundle map fibred over the identity on ν^*TM, and we have that

$$Tp_1\big(h^c(j^1_x\sigma,(\sigma(x),v))\big)=\rho^1(j^1_x\sigma),$$

with $p_1:\nu^*TM\to U$. This already guaranties that $\hat p_2: p_1^*J^1\nu\to J^1\nu$, which is fibred over p_1, determines an anchored bundle morphism between $(\hat p_1,h^c)$ and $(\nu_{1,0},\rho^1)$, with $\hat p_1: p_1^*J^1\nu\to\nu^*TM$. It now remains to check that the local lift functions of h^c are linear in (v^0,v^i). This follows straightforwardly from the above coordinate expression (where v^0,v^i represent the bundle coordinates of p_1). The local lift coefficients take the following form:

$$\begin{aligned}\Gamma^k_i(t,q^i,u^a,u^a_i) &= \frac{\partial\gamma^k}{\partial q^i}(t,q^j,u^a)+u^b_i\frac{\partial\gamma^k}{\partial u^b}(t,q^j,u^a),\\ \Gamma^0_i(t,q^i,u^a,u^a_i) &= 0,\\ \Gamma^0_0(t,q^i,u^a,u^a_i) &= 0,\\ \Gamma^k_0(t,q^i,u^a,u^a_0) &= \frac{\partial\gamma^k}{\partial t}(t,q^j,u^a)+u^b_0\frac{\partial\gamma^k}{\partial u^b}(t,q^j,u^a).\end{aligned}$$

This shows, in particular, that h^c is a ρ^1-lift. □

The above defined ρ^1-lift h^c is called the *control lift* associated to the geometric control structure (τ,ν,ρ). Note that the coefficients Γ^k_i of h^c do not depend on the coordinates u^a_0 of $J^1\nu$. In Remark 5.5 at the end of this section we will return to this point in more detail.

Let us denote the derivative operator associated to h^c by ∇ and let $\mathfrak{X}(\nu)$ denote the set of vector fields along ν, i.e. $\mathfrak{X}(\nu) = \{Z : U \to TM \mid \tau_M(Z(u)) = \nu(u)\}$. Note that, in view of the relation $\nu \circ p_1 = \tau_M \circ p_2$, we have $\mathfrak{X}(\nu) \cong \Gamma(p_1)$.

Proposition 4.5. *Given any $j^1_x\sigma \in J^1\nu$,*

$$\nabla_{j^1_x\sigma} Z = [\mathbf{T} \circ \rho \circ \sigma, Z \circ \sigma](x), \ \textit{for } Z \in \mathfrak{X}(\nu)(\cong \Gamma(p_1))$$

(where the square brackets on the right-hand side denote the ordinary Lie bracket of vector fields on M).

Proof. Recalling the coordinate expression for the ρ^1-derivative associated with the ρ^1-lift h^c (cf. Section II-5), we obtain, with a slight abuse of notation,

$$\left(\nabla_{j^1_x\sigma} Z\right)^i = \left(\frac{\partial Z^i}{\partial t} + \gamma^j \frac{\partial Z^i}{\partial q^j} + (u^a_0 + \gamma^j u^a_j)\frac{\partial Z^i}{\partial u^a} - \Gamma^i_j Z^j - \Gamma^i_0 Z^0\right)(x).$$

The result then easily follows upon substituting $u^a_j = \frac{\partial \sigma^a}{\partial q^j}$ and $u^a_0 = \frac{\partial \sigma^a}{\partial t}$ in the right-hand side, and comparing this with the coordinate expression of the Lie bracket $[\mathbf{T} \circ \rho \circ \sigma, Z \circ \sigma](x)$. □

In the above, we have introduced a ρ^1-lift h^c and considered its associated derivative. In the following we determine the h^c-transport operator. According to the theory developed in Chapter II, the derivative operator ∇ acts on sections X of TM along the projected curve $\widetilde{c} = \nu \circ c$ of the base curve c of a ρ^1-admissible curve s. Indeed, given such an X, then (c, X) is a section of $\nu^*TM \to U$ along c. For notational convenience, we shall write $\nabla_{\mathsf{s}} X$ instead of $\nabla_{\mathsf{s}}(c, X)$. Similarly, we will consider the h^c-transport operator as a map on the fibres of TM, instead of a map on the fibres of ν^*TM. We first consider the case where s is a basic ρ^1-admissible curve, i.e. s takes on the special form $\mathsf{s}(t) = j^1\sigma(\widetilde{c}(t))$ for some section $\sigma \in \Gamma(\nu)$ and with $\widetilde{c} : [a, b] \to M$ an integral curve of $\mathbf{T} \circ \rho \circ \sigma$.

Lemma 4.6. *Let $\mathsf{s} : [a, b] \to J^1\nu$, $t \mapsto \mathsf{s}(t) = j^1\sigma(\widetilde{c}(t))$ be a basic ρ^1-admissible curve, and let $\{\phi_s\}$ denote the flow of $\mathbf{T} \circ \rho \circ \sigma$. Then the h^c-transport operator $\mathsf{s}^b_a : T_{\widetilde{c}(a)}M \to T_{\widetilde{c}(b)}M$ along s is given by $\mathsf{s}^b_a = T\phi_{b-a}$.*

Proof. Take $v_a \in T_{\widetilde{c}(a)}M$ and let $X(t)$ denote the section of TM along $\widetilde{c}(t)$ which is uniquely determined by the conditions $\nabla_{\mathsf{s}} X(t) = 0$ and $X(a) = v_a$.

This is still equivalent to

$$\left.\frac{d}{dt}\right|_t X(t) = \mathfrak{s}\left(T(\mathbf{T}\circ\rho)\big(T_{\widetilde{c}(t)}\sigma(X(t))\big)\right). \tag{4.6}$$

Now $\mathfrak{s}\,(T(\mathbf{T}\circ\rho)T_{\widetilde{c}(t)}\sigma(X(t))) = (\mathbf{T}\circ\rho\circ\sigma)^c(X(t))$, where $(\mathbf{T}\circ\rho\circ\sigma)^c$ denotes the complete lift of the vector field $\mathbf{T}\circ\rho\circ\sigma$ to TM. Taking this into account, Equation (4.6) tells us that $X(t)$ is an integral curve of $(\mathbf{T}\circ\rho\circ\sigma)^c$, passing through v_a. By construction of the complete lift of a vector field, the flow of $(\mathbf{T}\circ\rho\circ\sigma)^c$ is given by $\{T\phi_s\}$ and, therefore, $X(t) = T\phi_{t-a}(X(a))$. The result then follows immediately from the definition of the h^c-transport operator along s. □

Next, we consider the case where $\mathsf{s} : [a,b] \to J^1\nu$ is a piecewise ρ^1-admissible curve whose projection $\widetilde{c} = \nu_1 \circ \mathsf{s}$ onto M is piecewise smooth. Recall, in particular, that $c(t) := \nu_{1,0}(\mathsf{s}(t))$ is a control (see Lemma 4.1). Let $X : [a,b] \to V\tau$ be a piecewise smooth curve projecting onto the base curve $\widetilde{c}(t)$ of s. Thus, assume that $\mathsf{s} = \mathsf{s}_\ell \cdot \ldots \cdot \mathsf{s}_1$, with $\mathsf{s}_i : [a_{i-1}, a_i] \to J^1\nu$ smooth ρ^1-admissible curves. Since the derivative operator ∇_{s} is not defined if $c = \nu_{1,0}\circ\mathsf{s}$ is a discontinuous curve, we now extend the definition simply by putting $\nabla_{\mathsf{s}}X(t) = \nabla_{\mathsf{s}(t)}X$, given any piecewise smooth section X of TM along $\widetilde{c}$. This generalisation is only possible because we are working on the pull-back bundle ν^*TM and because $\widetilde{c}$ is assumed to be piecewise smooth. Now, the equation $\nabla_{\mathsf{s}}X(t) = 0$ for $t \in [a,b]$ also admits a unique solution, provided an initial point $v_a = X(a)$ is given and X is assumed continuous. The notion of a h^c-transport operator is then well-defined, and it is an easy exercise to see that $\mathsf{s}_a^b = (\mathsf{s}_\ell)_{a_{\ell-1}}^b \circ \cdots \circ (\mathsf{s}_1)_a^{a_1}$ and, more generally, that $\mathsf{s}_{\overline{t}}^t = (\mathsf{s}_j)_{a_{j-1}}^t \circ \cdots \circ (\mathsf{s}_i)_{\overline{t}}^{a_i}$, where we assumed that $\overline{t} \leq t$ with $t \in [a_{j-1}, a_j]$ and $\overline{t} \in [a_{i-1}, a_i]$.

Using the above lemma and the fact that every piecewise ρ^1-admissible curve s is the composition of a finite number of basic ρ^1-admissible curves (cf. Proposition 4.2), we can solve the h^c-transport operated explicitly once we have fixed a finite number of basic ρ^1-admissible curves. Indeed, assume that $\mathsf{s}_i = j^1\sigma_i(\widetilde{c}(t))$, for $\sigma_i \in \Gamma(\nu)$ and $t \in [a_{i-1}, a_i]$, with $\widetilde{c}|_{[a_{i-1},a_i]}$ an integral curve of $\mathbf{T}\circ\rho\circ\sigma_i$, and such that $\mathsf{s} = \mathsf{s}_\ell \cdot \ldots \cdot \mathsf{s}_1$, then

$$\begin{aligned}
\mathsf{s}_a^b(X(a)) &:= (\mathsf{s}_\ell)_{a_{\ell-1}}^b \circ \cdots \circ (\mathsf{s}_1)_a^{a_1}(X(a)) \\
&= T\phi^\ell_{a_\ell - a_{\ell-1}} \circ \cdots \circ T\phi^1_{a_1 - a}(X(a)) \\
&= T\Phi_a^b(X(a)) = X(b),
\end{aligned}$$

where $\{\phi_t^i\}$ denotes the flow of $\mathbf{T} \circ \rho \circ \sigma_i$ and where we use the shorthand notation introduced in Section 2.

Returning to the given geometric control structure (τ, ρ, ν), we shall now explain the role of the h^c-transport operator in determining the variational cone.

Consider an arbitrary control $c : [a, b] \to U$, taking x to y and with base curve $\widetilde{c} = \nu \circ c$. In Section 2 we have seen that c is induced by some ordered set $\mathcal{S} = (\sigma_\ell, \ldots, \sigma_1)$ of sections of ν and a composite flow parameter $T = (t_\ell, \ldots, t_1) \in \mathbb{R}_+^\ell$, i.e. $c = c_\ell \cdot \ldots \cdot c_1$, with $c_i(t) = \sigma_i(\widetilde{c}(t))$ for $t \in [a_{i-1}, a_i]$ and $a_i = a_{i-1} + t_i$. If Φ represents the composite flow of the ordered set $\mathcal{X} = (\mathbf{T} \circ \rho \circ \sigma_\ell, \ldots, \mathbf{T} \circ \rho \circ \sigma_1)$, then the variational cone $C_y R_x(\Phi, T)$ was defined as the cone generated by all tangent vectors of the form: $T\Phi_\tau^b(Y(\widetilde{c}(\tau)))$, where $\tau \in [a, b]$ and $Y \in \mathcal{D}$. Using the above definitions, it is easily seen that $T\Phi_\tau^b = \mathsf{s}_\tau^b$, with s the piecewise ρ^1-admissible curve defined as the composition of the basic ρ^1-admissible curves $\mathsf{s}_i(t) = j^1\sigma_i(\widetilde{c}(t))$ for $t \in [a_{i-1}, a_i]$ and $i = 1, \ldots, \ell$. This implies that the variational cone $C_y R_x(\Phi, T)$ only depends on the piecewise ρ^1-admissible curve s. Roughly speaking, one can say that the (piecewise) ρ^1-admissible curve s with base curve the control c, contains sufficient information regarding the sections σ_i in order to determine the variational cone $C_y R_x$. Therefore, we say that the variational cone is associated with s and we write $C_y R_x(\mathsf{s})$, if we want to emphasise that the variational cone can be generated by the h^c-transport operator along the (piecewise) ρ^1-admissible curve s.

5 Properties of variational cones in control theory

In this section, we investigate some properties of the control lift h^c with respect to the variational cone. We first prove that the associated derivative operator ∇ leaves the subbundle $\nu^* V\tau$ of $\nu^* TM$ invariant. Take an arbitrary point $\mathsf{s} = j_x^1\sigma$ in $J^1\nu$. Then, by definition, ∇_{s} acts on the set $\mathfrak{X}(\nu)$. Consider the set $\mathcal{V}(\nu)$ of τ-vertical vector fields along ν, i.e.

$$\mathcal{V}(\nu) = \{Z : U \to V\tau \mid \tau_M(Z(s)) = \nu(s), \text{ for all } s \in U\},$$

which is a subset of $\mathfrak{X}(\nu)$. It is easily seen that $\mathcal{V}(\nu) \cong \Gamma(\widetilde{p}_1)$, with $\widetilde{p}_1 : \nu^* V\tau \to U$, the projection onto the first factor. Fix any $Z \in \mathcal{V}(\nu)$. We now prove that $\nabla_{\mathsf{s}} Z \in V_x\tau$, which implies that ∇ leaves the subbundle $\nu^* V\tau \to U$ of $\nu^* TM \to U$ invariant (cf. Section II-6), and, in turn, this implies that the restriction of the derivative ∇ to $\nu^*V\tau \to U$ determines a ρ^1-lift on the bundle $\nu^*V\tau$.

Proof. Recall that $\nabla_{\mathsf{s}} Z = [\mathbf{T}\circ\rho\circ\sigma, Z\circ\sigma](x)$ where σ is an arbitrary section of ν such that $j^1_x\sigma = \mathsf{s}$. Since the vector field $\mathbf{T}\circ\rho\circ\sigma$ is τ-related to the vector field $\partial/\partial t$ on $\mathbb{R}$, we obtain that the flows of $\mathbf{T}\circ\rho\circ\sigma$ and $\partial/\partial t$ are τ-related. Therefore $T\tau([\mathbf{T}\circ\rho\circ\sigma, Z\circ\sigma](x)) = 0$, or $\nabla_{\mathsf{s}} Z \in V_x\tau$ holds. $\square$

For notational convenience, we denote the ρ^1-lift on $\nu^*V\tau \to U$, determined by the control lift h^c on $\nu^*TM \to U$, and its derivative by the same symbols: h^c and ∇, respectively. The following lemma relates the corresponding 'dual' derivative operators on $\nu^*V^*\tau$ and ν^*T^*M (where $V^*\tau$ is the dual bundle of $V\tau$). For that purpose, we introduce the following section ζ of the fibration $\nu^*T^*M \to \nu^*V^*\tau$. Take $s \in U_x, \theta_x \in V^*_x\tau$ and put $\zeta(s,\theta_x) = (s,\alpha_x)$, where $\alpha_x \in T^*_xM$ is uniquely determined by the conditions $\langle \alpha_x, \mathbf{T}(\rho(s))\rangle = 0$ and α_x projects onto θ_x. The mapping ζ is smooth, as can be easily seen from the following coordinate expression: putting $s = (t, q^i, u^a)$ and $\theta_x = (t, q^i, p_i)$, a straightforward computation gives

$$\zeta(t, q^i, u^a, p_i) = \left(t, q^i, u^a, -\gamma^i(t, q^i, u^a)p_i, p_i\right),$$

i.e. $\zeta(s,\theta_x) = -\gamma^i(t,q^i,u^a)p_i\,dt + p_i\,dq^i_{|x}$.

Lemma 5.1. *Let $(c(t),\alpha(t))$ denote a section of $\nu^*T^*M \to U$ along a control $c : [a,b] \to U$, with $\widetilde{c}$ the base curve of c, such that $\alpha(t)$ is piecewise smooth. Let $\theta(t)$ denote the projection of $\alpha(t)$ onto $V^*\tau$ (i.e. $(c(t),\theta(t))$ is a section of $\nu^*V^*\tau \to U$) and assume that $\zeta(c(t),\theta(t))$ is piecewise smooth. Fix a piecewise ρ^1-admissible curve $\mathsf{s}(t)$ with base curve the control c. The following equivalence holds: $\nabla_{\mathsf{s}}\alpha(t) = 0$ and $\langle\alpha(a), \mathbf{T}(\rho(c(a)))\rangle = 0$ iff $\nabla_{\mathsf{s}}\theta(t) = 0$ and $\zeta(c(t),\theta(t)) = \alpha(t)$ for all $t \in [a,b]$.*

Proof. The condition, saying that $\zeta(c(t),\theta(t))$ is piecewise smooth, is nontrivial since $c(t)$ may have discontinuous points. In coordinates $\nabla_{\mathsf{s}}\alpha(t) = 0$ takes the form, with $\alpha(t) = (t, q^i(t), p_0(t), p_i(t))$:

$$\begin{aligned}
\dot p_0(t) &= -\Gamma^j_0(t, q^i(t), u^a(t), u^a_i(t))p_j(t)\\
\dot p_i(t) &= -\Gamma^j_i(t, q^i(t), u^a(t), u^a_0(t))p_j(t).
\end{aligned}$$

Then, it is easily seen that the projection $\theta(t) = (t, q^i(t), p_i(t))$ of $\alpha(t)$ onto $V^*\tau$ satisfies $\nabla_{\mathsf{s}}\theta(t) = 0$. Recall the expression for the coefficients of ∇, as determined in Proposition 4.4. Consider the function

$$f(t) = -p_j(t)\gamma^j(t, q^i(t), u^a(t)).$$

This function is piecewise smooth since $\zeta(c(t), \theta(t))$ was assumed piecewise smooth. Using the fact that $\dot{q}^j(t) = \gamma^j(t, q^i(t), u^a(t))$ and

$$\dot{u}^a(t) = u_0^a(t) + \gamma^i(t, q^i(t), u^a(t))u_i^a(t),$$

we obtain, with a slight abuse of notation:

$$\begin{aligned}
\dot{f} &= -\dot{p}_j\gamma^j - p_j\frac{\partial\gamma^j}{\partial t} + \frac{\partial\gamma^j}{\partial q^i}\dot{q}^i + \frac{\partial\gamma^j}{\partial u^a}\dot{u}^a, \\
&= -\Gamma_j^k p_k\gamma^j - p_j\Gamma_k^j\gamma^k - p_j\Gamma_0^j, \\
&= -p_j\Gamma_0^j.
\end{aligned}$$

Hence, $\dot{p}_0(t) = \dot{f}(t)$ for all t, i.e. $p_0(t)$ equals $f(t)$ up to a constant, and the latter is determined by the condition that

$$\big\langle \alpha(a), \mathbf{T}(\rho(c(a)))\big\rangle = 0,$$

or, equivalently, $p_0(a) = f(a)$. This, in turn, shows that $\alpha(t) = \zeta(c(t), \theta(t))$. The converse follows by reversing the above arguments. □

The following proposition shows that the dual cones (cf. Remark 3.2) of the variational cone $C_yR_x(\mathsf{s})$ and the vertical variational cone $V_yR_x(\mathsf{s})$ are related by the map ζ (where s is a piecewise ρ^1-admissible curve whose base curve is a given control $c : [a, b] \to U$ taking x to y).

Proposition 5.2.

$\zeta(c(b), \cdot) : (V_yR_x)^* \to (C_yR_x)^*$ *is a one-to-one mapping.*

Proof. We first make the following remark. Assume that $\alpha_y \in (C_yR_x)^*$, then $\big\langle \alpha_y, \mathsf{s}_t^b\big(\mathbf{T}(\rho(c(t)))\big)\big\rangle = 0$. Indeed, this follows from the definition of the dual of a cone (cf. Remark 3.2 page 86) and from the fact that the tangent vectors $\mathsf{s}_t^b(\mathbf{T}(\rho(c(t))))$ and $-\mathsf{s}_t^b(\mathbf{T}(\rho(c(t))))$ are both contained in C_yR_x. Therefore, if θ_y denotes the restriction of α_y to $V_y^*\tau$, then

$$\begin{aligned}
\big\langle \theta_y, \mathsf{s}_t^b\big(\mathbf{T}(\rho(s)) - \mathbf{T}(\rho(c(t)))\big)\big\rangle &= \big\langle \alpha_y, \mathsf{s}_t^b\big(\mathbf{T}(\rho(s)) - \mathbf{T}(\rho(c(t)))\big)\big\rangle \\
&= \big\langle \alpha_y, \mathsf{s}_t^b\big(\mathbf{T}(\rho(s))\big)\big\rangle \leq 0
\end{aligned}$$

holds for all $s \in U_{\widetilde{c}(t)}$. By definition, we have $\theta_y \in (V_yR_x)^*$. We have that $(C_yR_x)^*$ is contained in the image of $\zeta(c(b),\cdot)$ when restricted to $(V_yR_x)^*$.

We now prove that any element $\theta_y \in (V_yR_x)^*$ is mapped by $\zeta(c(b),\cdot)$ onto an element of $(C_yR_x)^*$. Assume that $\theta_y \in (V_yR_x)^*$, i.e. $\langle\theta_y, v\rangle \leq 0$, given any $v \in V_yR_x$. Consider the piecewise smooth section $\theta(t)$ of $V^*\tau$ along $\widetilde{c}$, the base curve of c, such that $\nabla_{\mathsf{s}}\theta(t) = 0$ and $\theta(b) = \theta_y$. More specifically, $\theta(t) = (\mathsf{s}_t^b)^*(\theta_y)$.

We now prove that $\zeta(c(t),\theta(t))$ is continuous (and therefore it is piecewise smooth). Fix a time $t_0 \in]a,b[$ at which c is discontinuous and consider a bundle adapted coordinate system containing the point $c(t_0)$. We have to prove that the local function $-\gamma^i(t,q^i(t),u^a(t))p_i(t)$ is continuous at $t = t_0$. Since, by definition of $\theta(t)$,

$$\begin{aligned}\big\langle\theta(b), \mathsf{s}_t^b\big(\mathbf{T}(\rho(s)) - \mathbf{T}(\rho(c(t)))\big)\big\rangle &= \\ \big\langle\theta(t), \mathbf{T}(\rho(s)) - \mathbf{T}(\rho(c(t)))\big\rangle &\leq 0\end{aligned}$$

for all $s \in U_{\widetilde{c}(t)}$, we have, locally

$$p_i(t)\gamma^i(t,q^i(t),u^a) \leq p_i(t)\gamma^i(t,q^i(t),u^a(t)) \text{ for all } u^a \in \mathbb{R}^k.$$

In the above inequality we substitute for u^a the coordinates $u_+^a(t_0)$ of the limit from the right of the control $c(t)$ for t approaching t_0, i.e. $u_+^a(t_0) = \lim_{t\to t_0^+} u^a(t)$. Then, the above inequality becomes, when t approaches t_0 from the left:

$$p_i(t_0)\gamma^i(t_0,q^i(t_0),u_+^a(t_0)) \leq p_i(t_0)\gamma^i(t_0,q^i(t_0),u^a(t_0)).$$

(Note that c is left continuous by definition). On the other hand, if we substitute for u^a the coordinates of $c(t_0)$ and consider the inequality when t approaches t_0 from the right, we obtain:

$$p_i(t_0)\gamma^i(t_0,q^i(t_0),u^a(t_0)) \leq p_i(t_0)\gamma^i(t_0,q^i(t_0),u_+^a(t_0)).$$

We conclude that $p_i(t_0)\gamma^i(t_0,q^i(t_0),u_+^a(t_0)) = p_i(t_0)\gamma^i(t_0,q^i(t_0),u^a(t_0))$ or that $\zeta(c(t),\theta(t))$ is continuous at all t. This result allows us to apply Lemma 5.1. Putting $\alpha(t) = \zeta(c(t),\theta(t))$, then $\nabla_{\mathsf{s}}\alpha(t) = 0$ for all $t \in [a,b]$ holds.

We now prove that $\alpha(b) \in (C_yR_x)^*$. For that purpose it is sufficient to prove that $\big\langle\alpha(b), \mathsf{s}_t^b(v)\big\rangle \leq 0$, for any $v = \mathbf{T}(\rho(s))$ with $s \in U_{\widetilde{c}(t)}$. By definition of $\alpha(t)$, we know that $\big\langle\alpha(t), \mathbf{T}\big(\rho(c(t))\big)\big\rangle = 0$ and, since $\alpha(t) = (\mathsf{s}_t^b)^*(\alpha(b))$, this is equivalent to $\big\langle \alpha(b), \mathsf{s}_t^b\big(\mathbf{T}(\rho(c(t)))\big)\big\rangle = 0$. Now, assume that $v = \mathbf{T}(\rho(s))$, with $s \in U_{\widetilde{c}(t)}$. Then

$$\begin{aligned}
\langle \alpha(b), \mathsf{s}_t^b(v)\rangle &= \langle \alpha(t), v\rangle \\
&= \Big\langle \alpha(t), v - \mathbf{T}\big(\rho(c(t))\big)\Big\rangle \\
&= \Big\langle \theta(t), v - \mathbf{T}\big(\rho(c(t))\big)\Big\rangle \\
&= \Big\langle \theta_y, \mathsf{s}_t^b\Big(v - \mathbf{T}\big(\rho(c(t))\big)\Big)\Big\rangle \leq 0,
\end{aligned}$$

since ζ is a section of $\nu^*T^*M \to \nu^*V^*\tau$. We conclude that any element of $(V_yR_x)^*$ is mapped onto an element in $(C_yR_x)^*$ by $\zeta(c(b), \cdot)$, for which the inverse is just the restriction of the projection $\nu^*T^*M \to \nu^*V^*\tau$ to $(C_yR_x)^*$. □

In the next proposition, we derive an equivalent characterisation of the dual cone of the vertical variational cone, showing that the dual of $V_yR_x(\mathsf{s})$ in fact only depends on the control $c : [a, b] \to U$, i.e. on the base curve of s. Assume that $\theta(t)$ is a section of $V^*\tau$ along $\widetilde{c}$, the base curve of the control c. Given a time $t_0 \in [a, b]$, let $\zeta_{\theta(t_0)} : U_{\widetilde{c}(t_0)} \to \mathbb{R}$ denote the function defined by $\zeta_{\theta(t_0)}(s) = \big\langle \zeta(c(t_0), \theta(t_0)), \mathbf{T}\big(\rho(s)\big)\big\rangle$.

Proposition 5.3. *The following equivalence holds:*

$$\left.\begin{array}{l}
\nabla_{\mathsf{s}}\theta(t) = 0 \textit{ and,} \\
\zeta_{\theta(t_0)}(s) \leq \zeta_{\theta(t_0)}(c(t_0)) = 0 \\
\textit{for all } t_0 \in [a, b] \textit{ and all } s \in U_{\widetilde{c}(t)}
\end{array}\right\} \iff \theta(b) \in (V_yR_x)^*.$$

Proof. Assume that $\theta_y \in (V_yR_x)^*$ and consider the solution $\theta(t)$ of $\nabla_{\mathsf{s}}\theta(t) = 0$, with $\theta(b) = \theta_y$. We now prove that the maximum condition $\zeta_{\theta(t_0)}(s) \leq 0$ holds for this section $\theta(t)$.

Let $s \in U_{\widetilde{c}(t_0)}$ arbitrary, for some fixed $t_0 \in]a, b]$. Then $\mathsf{s}_{t_0}^b\Big(\mathbf{T}\big(\rho(s)\big) - \mathbf{T}\big(\rho(c(t_0))\big)\Big) \in V_yR_x$, implying that

$$\Big\langle \theta(t_0), \mathbf{T}\big(\rho(s)\big) - \mathbf{T}\Big(\rho\big(c(t_0)\big)\Big)\Big\rangle \leq 0,$$

which is precisely the maximum condition, if we recognise that the left hand side of this inequality equals

$$\Big\langle \zeta\big(c(t_0), \theta(t_0)\big), \mathbf{T}\big(\rho(s)\big)\Big\rangle - \Big\langle \zeta\big(c(t_0), \theta(t_0)\big), \mathbf{T}\Big(\rho\big(c(t_0)\big)\Big)\Big\rangle.$$

It now remains to check that this inequality is valid for $t_0 = a$. This is done by choosing an adapted coordinate system, in a neighbourhood of $c(a)$. Then, the above inequality becomes, for any t_0 sufficiently close to a such that $c|_{[a,t_0]}$ is smooth:

$$-\gamma^i(t_0, q^i(t_0), u^a(t_0))p_i(t_0) + \gamma^i(t_0, q^i(t_0), u^a)p_i(t_0) \leq 0,$$

where u^a is taken arbitrary and where $\theta(t_0) = (q^i(t_0), p_i(t_0))$. By taking t_0 as a parameter and considering the limit from the right for $t_0 \to a$, we obtain that

$$-\gamma^i(a, q^i(a), u^a(a))p_i(a) + \gamma^i(a, q^i(a), u^a)p_i(a) \leq 0.$$

This inequality is still valid for any u^a, and therefore, the maximum condition also holds for $t_0 = a$.

On the other hand, assume that the maximum condition holds for a section $\theta(t)$ of $V^*\tau$ along $\widetilde{c}$ is given, such that $\nabla_{\mathsf{s}}\theta(t) = 0$. Then, by reversing the above arguments, we obtain that

$$\left\langle \theta(b), \mathsf{s}^b_{t_0}\Big(\mathbf{T}\big(\rho(s)\big) - \mathbf{T}\big(\rho(c(t_0))\big)\Big)\right\rangle \leq 0,$$

for any $s \in U_{\widetilde{c}(t_0)}$ and for $t_0 \in [a, b]$. Now, since any element of V_yR_x is a finite linear combination, with *positive* coefficients of tangent vectors of the form $\mathsf{s}^b_{t_0}\Big(\mathbf{T}\big(\rho(s)\big) - \mathbf{T}\big(\rho(c(t_0))\big)\Big)$, the inequality $\langle\theta(b), v\rangle \leq 0$ holds for any $v \in V_yR_x$. This implies, by definition, that $\theta(b) \in (V_yR_x)^*$. □

Recall the expression for the coefficients of the derivative operator ∇ acting on $V^*\tau$. Then, locally, the equations $\nabla_{\mathsf{s}}\theta(t)$ read (using the notations from Lemma 5.1):

$$\dot{p}_i(t) = -\Gamma^j_i(c(t))p_j(t) = -\left(\frac{\partial\gamma^j}{\partial q^i}(c(t)) + \frac{\partial\gamma^j}{\partial u^a}(c(t))u^a_i(t)\right)p_j(t)$$

If $\theta(t) = (q^i(t), p_i(t))$ satisfies the conditions from the above proposition, then the maximum condition implies that:

$$p_i(t)\frac{\partial\gamma^i}{\partial u^a}(c(t)) = 0.$$

If we insert this equality in the differential equations for $p_i(t)$, we obtain that the components $p_i(t)$ of $\theta(t)$ have to satisfy:

$$\dot{p}_i(t) = -\left(\frac{\partial\gamma^i}{\partial q^j}(c(t))\right)p_j(t).$$

Therefore, if $\theta(t)$ satisfies the conditions of the above theorem (i.e. $\nabla_{\mathsf{s}}\theta(t) = 0$ and the maximum condition), then it satisfies a system of differential equations which only depend on the control c and no longer on the first jet coordinates of s. This observation is captured by the following corollary, where ω_0 is the pull-back of the canonical symplectic form on T^*M to $\nu^*V\tau$ along $p_2^* \circ \zeta$, with $p_2^* : \nu^*T^*M \to T^*M$ the projection onto the second factor.

Corollary 5.4. *Let $\theta_y \in V^*_{\widetilde{c}(b)}\tau$. Then $\theta_y \in (V_yR_x)^*$ iff there exists a piecewise smooth section $\theta(t)$ of $V^*\tau$, passing through θ_y at $t = b$, such that the curve $(c(t), \theta(t))$ is a solution of the implicit Hamiltonian differential equation $i_{(\dot{c}(t),\dot{\theta}(t))}\omega_0 = 0$, and such that $\zeta_{\theta(t_0)} : U_{\widetilde{c}(t_0)} \to \mathbb{R}$, defined above, attains a maximum at $c(t_0)$ for any fixed $t_0 \in [a, b]$.*

Proof. We first note that the implicit Hamiltonian equation is only considered for every smooth part of $c(t)$, since $c(t)$ is allowed to be discontinuous at a finite number of points. Using Proposition 5.3, it will be sufficient to prove that $\nabla_{\mathsf{s}}\theta(t) = 0$ is equivalent to the above implicit Hamiltonian system, provided that the maximum condition holds for $\theta(t)$. This follows straightforwardly if we use the above coordinate expression for $\nabla_{\mathsf{s}}\theta(t) = 0$ and compare it with the local expression for the implicit Hamiltonian system. After some tedious calculations, we find that these implicit equations read, with $h_0(c(t), \theta(t)) = p_j\gamma^j(t, q^i, u^a)$:

$$
\begin{aligned}
\dot{q}^i(t) &= \frac{\partial h_0}{\partial p_i}(c(t), \theta(t)) = \gamma^i(c(t)),\\
0 &= \frac{\partial h_0}{\partial u^a}(c(t), \theta(t)) = \frac{\partial \gamma^i}{\partial u^a}(c(t))p_i(t),\\
\dot{p}_i(t) &= -\frac{\partial h_0}{\partial q^i}(c(t), \theta(t)) = -\frac{\partial \gamma^j}{\partial q^i}(c(t))p_j(t),\\
\left.\frac{d}{dt}\right|_t &(h_0(c(t), \theta(t))) = \frac{\partial h_0}{\partial t}(c(t), \theta(t)).
\end{aligned}
$$

The first equation holds since c is assumed to be a control. The second and third equations are shown above. If we recall the definition of the function $f(t)$ introduced in Lemma 5.1, then we have that $f(t) = h_0(t, q^i(t), u^a(t))$. It is easily seen that the fourth equation was proven to hold in Lemma 5.1 if $\nabla_{\mathsf{s}}\theta(t) = 0$. □

The dual cone of $V_yR_x(\mathsf{s})$ will be denoted by $V_y^*R_x(c)$, in order to indicate that it does not depend on the first jet coordinates of s. In the following remark we briefly explain how some of the basic ideas in the treatment of the maximum principle in [47] can be related to our work.

Remark 5.5. The discussion of the maximum principle can be developed for controls that verify the weaker assumption of being measurable and bounded, instead of (piecewise) smooth (see, for instance, L.S. Pontryagin et al. [47]). Using local coordinate expressions, we will roughly sketch how the smoothness conditions we have imposed on the controls can also be relaxed within our framework. The local expression for the equation $\nabla_{\mathsf{s}}X(t) = 0$ reads

$$\dot{X}^k(t) = \left(\frac{\partial\gamma^k}{\partial q^i}(t, q^j(t), u^a(t)) + u_i^b(t)\frac{\partial\gamma^k}{\partial u^b}(t, q^j(t), u^a(t))\right) X^i(t),$$

where, as usual, $\mathsf{s}(t) = (t, q^i(t), u^a(t), u_0^a(t), u_i^a(t))$. The condition, that the functions $u^a(t)$ and $u_i^a(t)$ should be measurable and bounded, suffices to obtain a solution of the equations and, subsequently, to introduce a suitable notion of transport operator. This observation can be translated into our geometric framework as follows. Consider the set $V^1\nu := \cup_{x\in M}\{T_x\sigma_{|V\tau} : V_x\tau \to T_{\sigma(x)}U \mid \sigma \in \Gamma(\nu)\}$. It can be proven by standard arguments that $V^1\nu$ is an affine bundle over U, with coordinates (t, q^i, u^a, u_i^a) (see, for instance, [49]). Note that there exists a natural projection $\mu : J^1\nu \to V^1\nu$, locally expressed by $(t, q^i, u^a, u_0^a, u_i^a) \mapsto (t, q^i, u^a, u_i^a)$. From the fact that the coefficients Γ_i^k of h^c do not depend on the u_0^a (see the proof of Proposition 6.3) it easily follows that the ρ^1-derivative ∇_{s} only depends on $\mu\circ\mathsf{s}$. Now, since s was assumed to be ρ^1-admissible, i.e. $\rho^1(\mathsf{s}) = \dot{c}$, the smoothness condition on c could not be relaxed. However, the curve $\mathsf{s}' = \mu\circ\mathsf{s}$ does not have to satisfy this condition, implying that the smoothness condition can be relaxed without losing the notion of derivative acting on sections of $V\tau$ along $\widetilde{c}$. We can therefore conclude that, in order to define a vertical cone of variations associated with a measurable and bounded control c, we must fix a curve s' in $V^1\nu$. If one works in a coordinate chart, a natural choice for s' is the curve $\mathsf{s}'(t) = (t, q^i(t), u^a(t), u_i^a(t))$ with $u_i^a(t) = 0$. The equations of the derivative associated with s' then reduce to $\dot{X}^k(t) = \frac{\partial\gamma^k}{\partial q^i}(t, q^j(t), u^a(t))X^i(t)$. These equations are precisely the "variational equations" introduced in [47, p79]. By fixing the coordinate chart, one can fix the section $\sigma^a(t,q) = u^a(t)$ and the curve $\mathsf{s}'(t) = (t, q^j(t), u^a(t), 0)$, implying that, respectively, a fixed vertical cone of variations and a fixed derivative associated with s' can be defined. Moreover, this specific choice of s' allowed that the local equations for $\nabla_{\mathsf{s}}\theta(t)$ can be regarded as "Hamiltonian" equations, where the Hamiltonian function equals the function h_0, as can easily be seen from Corollary 5.4. This essentially establishes the link between our approach and the one followed by L.S. Pontryagin et al. However, it should be noted that the approach in [47] is more general, in the sense that he worked with measurable and bounded controls.

To conclude this section, we now adapt the above results to the case where we have an geometric optimal control structure, say (τ, ν, ρ, L). We fix some notations, before stating the main theorem. As in Section 3, we consider the extended geometric control structure $(\overline{\tau}, \overline{\nu}, \overline{\rho})$. First, assume that c is a control taking x to y and assume that $\overline{c}$ is the associated control, in the extended setting, taking $\overline{x} = (x, J_x)$ to $\overline{y} = (y, J_y)$, with $J_y - J_x = \mathcal{J}(c)$. Next, consider the smooth one parameter family of sections ζ_λ, with $\lambda \in \mathbb{R}$, of the fibration $\nu^*T^*M \to \nu^*V^*\tau$, defined by $\zeta_\lambda(s, \theta) = \alpha$, where α is determined by the conditions $\langle \alpha, \mathbf{T}(\rho(s))\rangle + \lambda L(s) = 0$ and α projects onto θ. In a local coordinate system, the section ζ_λ takes the following form, putting $s = (t, q^i, u^a)$ and $\theta = p_i\, dq^i_{|x}$, $\zeta_\lambda(t, q^i, u^a, p_i) = -(p_i\gamma^i + \lambda L)dt + p_i dq^i$. Note that, in the case where $\lambda = 0$, then $\zeta_0 = \zeta$ with ζ defined at the beginning of this section. Therefore, the family of sections ζ_λ can be regarded as a "deformation" of the previously introduced section ζ. Similarly, let ω_λ be the closed two-form on $\nu^*V^*\tau$, determined by the pull-back under $p_2^* : \nu^*T^*M \to T^*M$ and $\zeta_\lambda : \nu^*V^*\tau \to \nu^*T^*M$ of the canonical symplectic two-form on T^*M. For $\lambda = 0$, we recover the closed two-form ω_0 used in Corollary 5.4 and so we can consider ω_λ as a "deformation" of ω_0. Assume that $c : [a, b] \to U$ is a control and that θ is a section of $V^*\tau$ along $\widetilde{c}$, then, with every $t_0 \in [a, b]$ we define a real valued function $\zeta^\lambda_{\theta(t_0)} : U_{\widetilde{c}(t_0)} \to \mathbb{R}$ by $\zeta^\lambda_{\theta(t_0)}(s) = \langle \zeta_\lambda(c(t_0), \theta(t_0)), \mathbf{T}(\rho(s))\rangle + \lambda L(s)$.

Theorem 5.6. *Assume that $\overline{\theta}_y = (\theta_y, \lambda) \in V^*_{\overline{y}}\overline{\tau}$. Then $\overline{\theta}_y \in V^*_{\overline{y}} R_{\overline{x}}(\overline{c})$ iff there exists a piecewise smooth curve $\theta(t)$ passing through θ_y at $t = b$ and such that $(c(t), \theta(t))$ is a solution of the implicit Hamiltonian equation $i_{(\dot{c}(t),\dot{\theta}(t))}\omega_\lambda = 0$ through θ_y at $t = b$, and such that $\zeta^\lambda_{\theta(t_0)} : U_{\widetilde{c}(t_0)} \to \mathbb{R}$ attains a maximum at $c(t_0)$, for all $t_0 \in [a, b]$.*

Proof. From Corollary 5.4 we know that the condition that a vertical covector $\overline{\theta}_y$ is contained in $V^*_{\overline{y}} R_{\overline{x}}$, is equivalent to the existence of a section $\overline{\theta}(t)$ of $V^*\overline{\tau}$ along the base curve of $\overline{c}$ passing through $\overline{\theta}_y$, such that $i_{(\dot{\overline{c}}(t),\dot{\overline{\theta}}(t))}\overline{\omega}_0 = 0$ and $\overline{\zeta}_{\overline{\theta}(t)}$ attains its maximum at $\overline{c}(t)$. Assume that $\overline{\theta}(t) = (\theta(t), \lambda(t))$. Since $\overline{M}$ is a product manifold, we can write the canonical symplectic two-form $\overline{\omega}$ on $T^*\overline{M}$ as $\overline{\omega} = \omega + dp_J \wedge dJ$. The pull-back of $\overline{\omega}$ under $\overline{\zeta}$ then becomes: $\overline{\omega}_0 = -d\overline{h}_0 \wedge dt + dp_i \wedge dq^i + dp_J \wedge dJ$, with $\overline{h}_0 = p_i\gamma^i + p_J L$ and the implicit Hamiltonian system $i_{(\dot{\bar{c}}(t),\dot{\bar{\theta}}(t))}\overline{\omega}_0 = 0$ reads:

$$\begin{aligned}
&\dot{q}^i(t) = \frac{\partial \overline{h}_0}{\partial p_i}(\overline{c}(t),\overline{\theta}(t)) = \gamma^i(c(t)), \qquad \dot{J}(t) = L(c(t)),\\
&0 = \frac{\partial \overline{h}_0}{\partial u^a}(\overline{c}(t),\overline{\theta}(t)) = \frac{\partial \gamma^i}{\partial u^a}(c(t))p_i(t) + p_J(t)\frac{\partial L}{\partial u^a}(c(t)),\\
&\dot{p}_i(t) = -\frac{\partial \overline{h}_0}{\partial q^i}(\overline{c}(t),\overline{\theta}(t)) = -\frac{\partial \gamma^j}{\partial q^i}(c(t))p_j(t) - p_J(t)\frac{\partial L}{\partial q^i}(c(t)),\\
&\dot{p}_J(t) = -\frac{\partial \overline{h}_0}{\partial J}(\overline{c}(t),\overline{\theta}(t)) = 0,\\
&\frac{d}{dt}\bigg|_t \big(h_1(\overline{c}(t),\overline{\theta}(t))\big) = \frac{\partial \overline{h}_0}{\partial t}(\overline{c}(t),\overline{\theta}(t)).
\end{aligned}$$

It is easily seen that the 'cost momentum' $p_J(t)$ is constant and, since $\overline{\theta}(b) = (\theta_y, \lambda)$, $p_J(t) = \lambda$. If we substitute this in the above equations we obtain:

$$\begin{aligned}
&\dot{q}^i(t) = \frac{\partial h_\lambda}{\partial p_i}(\overline{c}(t),\overline{\theta}(t)) = \gamma^i(c(t)),\\
&0 = \frac{\partial h_\lambda}{\partial u^a}(\overline{c}(t),\overline{\theta}(t)) = \frac{\partial \gamma^i}{\partial u^a}(c(t))p_i(t) + \lambda\frac{\partial L}{\partial u^a}(c(t)),\\
&\dot{p}_i(t) = -\frac{\partial h_\lambda}{\partial q^i}(\overline{c}(t),\overline{\theta}(t)) = -\frac{\partial \gamma^j}{\partial q^i}(c(t))p_j(t) - \lambda\frac{\partial L}{\partial q^i}(c(t)),\\
&\frac{d}{dt}\bigg|_t \big(h_\lambda(\overline{c}(t),\overline{\theta}(t))\big) = \frac{\partial h_\lambda}{\partial t}(\overline{c}(t),\overline{\theta}(t)),
\end{aligned}$$

with $h_\lambda = p_i\gamma^i + \lambda L$. It is readily seen that these equations are equivalent to $i_{(\dot{c}(t),\dot{\theta}(t))}\omega_\lambda = 0$. Moreover, since $p_J(t) = \lambda$ is constant, the following are also easily derived:

$$\overline{\zeta}_{\overline{\theta}(t_0)}(s) = \zeta_{\theta(t_0)}(s) + \lambda\big(L(s) - L(c(t_0))\big) = \zeta^\lambda_{\theta(t_0)}(s),$$

for any $s \in U_{\widetilde{c}(t_0)}$ and any fixed $t_0 \in [a,b]$. This implies that the maximum condition holds. The converse is proven by reversing these arguments. □

We shall introduce a special denomination for the curves in $V^*\tau$ which satisfy the implicit Hamiltonian equation of Theorem 5.6.

Definition 5.1. Assume that $c : [a,b] \to U$ is a control taking x to y, with base curve $\widetilde{c}$. Then a pair $(\theta(t), \lambda)$, where $\theta(t)$ is a section of $V^*\tau$ along $\widetilde{c}$ and λ a real number, is called *a multiplier of c* if

1. $i_{(\dot{c}(t),\dot{\theta}(t))}\omega_\lambda = 0$ for all $t \in [a,b]$ where c is smooth, and
2. $\zeta^\lambda_{\theta(t)}(s) \leq \zeta^\lambda_{\theta(t)}(c(t)) = 0$ for any $s \in U_{\tilde{c}(t)}$ and any fixed $t \in [a,b]$.

Corollary 5.7. *Assume that $(\theta(t),\lambda)$ is a multiplier of the control $c : [a,b] \to U$. Then: (i) $(\theta(b),\lambda) \in V^*_{\overline{y}}\overline{\tau}$ is an element in $V^*_{\overline{y}}R_{\overline{x}}(\overline{c})$, and (ii) if $\lambda = 0$, then $\theta(b) \in V^*_y\tau$ determines an element of $V^*_yR_x(c)$.*

Proof. (i) is a straightforward consequence of Theorem 5.6 and the definition of a multiplier. (ii) follows from the fact that $\omega_\lambda = \omega_0$ and $\zeta^\lambda_{\theta(t)} = \zeta_{\theta(t)}$ for $\lambda = 0$. □

Note, that, every linear combination with positive coefficients of a finite number of multipliers, is again a multiplier. In particular, every multiplier can be "normalised" in the sense that given any multiplier $(\theta(t),\lambda)$ of c with $\lambda \neq 0$ then the pair $(|\lambda|^{-1}\theta(t), \mathrm{sgn}(\lambda))$ is also a multiplier of c.

6 The maximum principle and extremal controls

We will now derive the maximum principle by combining the tools developed in Section 5 and the necessary conditions for optimal controls derived in Section 3. In the next section we will derive necessary conditions for optimal controls with variable endpoint conditions.

Assume that a geometric optimal control structure (τ,ν,ρ,L) is given.

Theorem 6.1 (The maximum principle). *Assume that c is an optimal control taking x to y. Then there exists a nontrivial multiplier (θ,λ) with $\lambda \leq 0$.*

Proof. If c is optimal, then, from Proposition 3.3, we know that the cones $V_{\overline{y}}R_{\overline{x}}(\overline{\mathsf{s}})$ and $C(-\partial/\partial J)$, in the extended geometric control structure, are separable, where $\overline{\mathsf{s}}$ is an arbitrary $\overline{\rho}^1$-admissible curve, with base curve the control $\overline{c}$ taking $\overline{x} = (x, J_x)$ to $\overline{y} = (y, J_y)$ (with $J_y - J_x = \mathcal{J}(c)$). This is equivalent to saying that there exists a nontrivial element $\overline{\theta}_{\overline{y}}$ in the dual vertical variational cone $V^*_{\overline{y}}R_{\overline{x}}$, such that $\langle\overline{\theta}_{\overline{y}}, -\partial/\partial J\rangle \geq 0$. If we assume that $\overline{\theta}_{\overline{y}} = (\theta_y,\lambda)$, then this inequality is equivalent to $\lambda \leq 0$. From Theorem 5.6, we know that the element $\overline{\theta}_y \in V^*_{\overline{y}}R_{\overline{x}}$ determines a multiplier (θ,λ), with $\theta(b) = \theta_y$. □

It is a common practice in optimal control theory to call a control c an *extremal control*, if it admits a multiplier $(\theta(t), \lambda)$ with $\lambda \leq 0$. The maximum principle says that an optimal control is an extremal control. The converse is in general not true. If c admits a multiplier $(\theta(t), \lambda)$ with $\lambda = 0$, then c is called an *abnormal extremal*. If $\lambda < 0$, then c is called a *normal extremal*. An extremal control is called *strictly abnormal* if it is abnormal, but not normal. From the definition of a multiplier, it is easily seen that an abnormal extremal satisfies conditions independent of the cost function L, although these conditions are necessary for being an optimal control. For that reason they are called "abnormal". It should be noted that R. Montgomery [45] proved that there exist optimal controls that are strictly abnormal extremals.

The following theorem, gives necessary and sufficient conditions for abnormal and strictly abnormal extremal controls. Recall that the closure of a cone C in a topological vector space equals $(C^*)^*$, which implies that, in the specific case of a geometric control structure, $\mathrm{cl}(V_yR_x) = (V_y^*R_x(c))^*$, implying that the cone $\mathrm{cl}(V_yR_x)$ only depends on the control c.

Theorem 6.2. *The following equivalences hold: (i) a control c taking x to y is an abnormal extremal iff $\mathrm{cl}(V_yR_x)(c) \neq V_y\tau$ and (ii) an extremal control c taking x to y is a strictly abnormal extremal iff $-\partial/\partial J$ belongs to the boundary of $\mathrm{cl}(V_{\overline{y}}R_{\overline{x}})(\overline{c})$.*

Proof. The first statement (i) follows from the fact that every element in the dual cone $(\mathrm{cl}(V_yR_x))^* = V_y^*R_x$ corresponds to a multiplier with $\lambda = 0$ (see Corollary 5.7).

We now prove statement (ii). If an extremal $c : [a, b] \to U$ is strictly abnormal then every element $\overline{\theta}_y = (\theta_y, \lambda)$ in $V_{\overline{y}}^*R_{\overline{x}}$ satisfies $\lambda \geq 0$ (indeed, otherwise it would be normal). Using the definition of the dual of a cone, we obtain that $-\partial/\partial J$ is contained in $\mathrm{cl}(V_{\overline{y}}R_{\overline{x}})$, since $\langle\overline{\theta}_y, -\partial/\partial J\rangle \leq 0$ for any $\overline{\theta}_y \in V_{\overline{y}}^*R_{\overline{x}}$. On the other hand, since c is an extremal we know that $-\partial/\partial J$ is not contained in the interior of the cone $\mathrm{cl}(V_{\overline{y}}R_{\overline{x}})$. Therefore, $-\partial/\partial J$ has to be contained in the boundary of $\mathrm{cl}(V_{\overline{y}}R_{\overline{x}})$. The converse follows by reversing the arguments. □

The above characterisation of an abnormal extremal c implies that $\mathrm{cl}(V_yR_x)$ does not equal the total space of vertical tangent vectors. Since $V_y^*R_x = (C_yR_x)^*$, this is equivalent to saying that $\mathrm{cl}(C_yR_x)$ does not equal the total tangent space T_yM. This result can be intuitively interpreted as follows. An optimal control c is called abnormal if the control does not admit enough "variations" (cf. Section 2, page 64) in order for the associated variational cone to equal the total tangent space. In the case where c is strictly abnormal, we can say that the maximum principle fails since Corollary 2.5 (page 74), from which the maximum principle follows, only gives information on those tangent vectors lying in the interior of a variational cone, and not on those belonging to the boundary. Note that, since the variational cone is entirely contained in the integrable distribution $\widetilde{D}$, the above theorem says that, if $\widetilde{D}_x \neq T_xM$, for some $x \in M$, then any extremal through x, is abnormal. Therefore, it is often assumed that the integrable distribution, generated by D, spans the entire tangent bundle TM. If this is not the case, then we always assume that we are working on the pull back anchored bundle (cf. Chapter I, page 6). This reasoning is no longer valid for optimal control problems with variable endpoints, since the submanifolds P_i and P_f are, in general, not contained in a leaf of the foliation induced by $\widetilde{D}$.

In order to check whether a control c is abnormal or not, one has to determine the vertical variational cone. For that purpose, one has to integrate some set of time dependent vector fields $\mathbf{T}(\rho(\sigma_i))$, where σ_i are a finite number of sections of ν generating the control c (cf. Lemma 4.1, page 90). Indeed, once this is done, one can construct the variational cone (cf. Section 2, page 68), and verify wether or not it equals the whole tangent space. In case one wants to check if a control is strictly abnormal, the calculations become even more involved, since one is obliged to compute the flows of a set of time dependent vector fields in the extended geometric control structure. Moreover, the condition that $-\partial/\partial J$ should belong to the boundary of this cone is difficult to check. In Section 8, where we will consider (linear) autonomous geometric control structures, we will be able to derive some general properties of this cone that might make the calculations easier. Among others, we will give a sufficient condition for the non existence of abnormal extremals, which can be easily computed.

7 Optimal control problems with variable endpoint conditions

In this section we prove a version of the maximum principle for optimal control problems with variable endpoint conditions (cf. Section 3). Assume that a geometric optimal control structure (τ, ν, ρ, L) is given and that P_i and P_f are two immersed submanifolds of M.

Theorem 7.1 (The maximum principle). *Assume that a control* $c : [a,b] \to U$*, taking* $x \in P_i$ *to* $y \in P_f$*, is* (P_i, P_f)*-optimal. Then there exists* *a multiplier* (θ, λ) *with (i)* $\lambda \leq 0$*, (ii)* $\zeta_\lambda(c(a), \theta(a)) \in (T_x P_i)^0$*, and (iii)* $\zeta_\lambda(c(b), \theta(b)) \in (T_y P_f)^0$.

Proof. Recall the notations introduced in Section 3. From Proposition 3.4 we know that if c is (P_i, P_f)-optimal, then the cones K and K' are separable, where $K = \mathsf{\bar{s}}_a^b(T_{\overline{x}}\overline{P}) * C_{\overline{y}} R_{\overline{x}}$ and $K' = C(-\partial/\partial J) * T_{\overline{y}}\overline{Q}$, for some piecewise $\overline{\rho}^1$-admissible curve $\mathsf{\bar{s}}$ with base curve the extended control $\overline{c}$ taking (x, J_x) to (y, J_y). From Remark 3.2, we have that the above condition is equivalent to saying that there exists an element $\overline{\alpha}_y = (\alpha_y, \lambda)$ in the dual cone K^* such that

$$\langle \overline{\alpha}_y, v \rangle \geq 0 \text{ for all } v \in K'. \tag{7.7}$$

Note that, since $K^* = (\mathsf{\bar{s}}_a^b(T_{(x,J_x)}\overline{P}))^* \cap (C_{\overline{y}} R_{\overline{x}})^*$, the element $\overline{\alpha}_y$ in K^* is contained in $(C_{\overline{y}} R_{\overline{x}})^*$. In particular, from Proposition 5.2, we know that $\overline{\alpha}_y = \overline{\zeta}(\overline{c}(b), \overline{\theta}_y)$, where $\overline{\theta}_y = (\theta_y, \lambda)$ is an element of $V_{\overline{y}}^* R_{\overline{x}}$, which is the projection of $\overline{\alpha}_y$ onto $V^*\overline{\tau}$. Since any element in $V_{\overline{y}}^* R_{\overline{x}}$ determines a multiplier of c (cf. Theorem 5.6), we can consider the multiplier $(\theta(t), \lambda)$ for c, with $\theta(b) = \theta_y$.

Since, $C(-\partial/\partial J)$ is a subcone of K', Equation 7.7 tells us that $\lambda \leq 0$, which proves (i). On the other hand, since $T_{\overline{y}}\overline{P}_f$ is also contained in K', we have

$$\langle \overline{\alpha}_y, T_{\overline{y}}\overline{P}_f \rangle = \langle \alpha_y, T_y P_f \rangle \geq 0.$$

Since $T_y P_f$ is a linear subspace, one easily deduces that $\alpha_y \in (T_y P_f)^0$, proving (iii). Condition (ii) follows from a similar argument. Since $\overline{\alpha}_y \in (\mathsf{\bar{s}}_a^b(T_{\overline{x}}\overline{P}_i))^*$, we have $\langle \overline{\alpha}_y, \mathsf{\bar{s}}_a^b(T_{\overline{x}}\overline{P}_i) \rangle \leq 0$. The left-hand side of this inequality can be rewritten as:

$$\langle \overline{\alpha}(a), T_{\overline{x}}\overline{P}_i \rangle = \langle \alpha(a), T_x P_i \rangle \leq 0.$$

Note that, if $\overline{\alpha}(t) = (\mathsf{\bar{s}}_t^b)^*(\overline{\alpha}_y)$, then $\overline{\alpha}(t) = \overline{\zeta}(\overline{c}(t), \overline{\theta}(t))$, with $\overline{\theta}(t) = (\theta(t), \lambda)$ and that $\overline{\zeta}(\overline{c}(t), \overline{\theta}(t)) = (\zeta_\lambda(c(t), \theta(t)), \lambda)$ holds (this easily follows if we use a local coordinate chart). Now, since $T_x P_i$ is a linear subspace, one necessarily has that $\alpha(a) \in (T_x P_i)^0$, which concludes the proof. □

8 Autonomous optimal control problems

With an anchored bundle (ν^A, ρ^A), where $\nu^A : C \to Q$ is a fibre bundle and $\rho^A : C \to TQ$ is a bundle map fibred over the identity on Q, one can associate a geometric control structure (τ, ν, ρ), in the following way:

1. $\tau : M = \mathbb{R} \times Q \to \mathbb{R} : (t, q) \mapsto \tau(t, q) = t$,
2. $\nu : U = \mathbb{R} \times C \to M : (t, s) \mapsto \nu(t, s) = (t, \nu^A(s))$, and
3. $\rho : U \to J^1\tau : (t, s) \mapsto (t, \rho^A(s))$, since $J^1\tau \cong \mathbb{R} \times TQ$.

The anchored bundle (ν^A, ρ^A) is sometimes called an *autonomous geometric control structure*, as will be justified in the following paragraph. In this section we will reformulate the maximum principle in the framework of an autonomous geometric control structure. It is precisely this version of the maximum principle that will be extensively used in further developments. Furthermore, we will reconsider the characterisation of (strictly) abnormal extremals and study the case where (ν^A, ρ^A) is a linear anchored bundle (or *linear autonomous geometric control structure*).

We now translate concepts defined in the anchored bundle (ν^A, ρ^A) to known concepts in the associated geometric control structure (τ, ν, ρ). Note first that, if c is a control in (τ, ν, ρ), then, by definition of ρ, it satisfies the following *autonomous* system of differential equations:

$$\dot{q}^j(t) = \gamma^j(q^i(t), u^a(t)),$$

which justifies the above denomination. We now consider the relation between controls and ρ^A-admissible curves. The natural projection of $U = \mathbb{R} \times C$ onto its second factor C, is denoted by p_C and, similarly, the projection of M onto Q is denoted by p_Q. It is an easy exercise to see that every ρ^A-admissible curve $c^A : [a, b] \to C$ determines a smooth control $c : [a, b] \to U$ with $c(t) = (t, c^A(t))$. On the other hand, since (p_C, p_Q) is an anchored bundle morphism between $(\nu, \mathbf{T} \circ \rho)$ and (ν^A, ρ^A), we have that any control $c(t)$ takes the form $(t, c^A(t))$ where the curve $c^A(t) = p_C(c(t))$ is a ρ^A-admissible curve. In particular, this implies that, given any point $x = (a, q) \in M$, then the set of reachable points R_x in M, determined with respect to the anchored bundle $(\nu, \mathbf{T} \circ \rho)$, is mapped by p_Q onto R_q in Q, with R_q the set of reachable points from q with respect to the anchored bundle (ν^A, ρ^A). Using the above correspondence and the fact that $\tau(R_{(a,q)})$ is a half-open interval $[a, b[$ (cf. page 63), the subset $[a, b[\times R_q$ of M is precisely the set of reachable points $R_{(a,q)}$ defined with respect to the anchored bundle $(\nu, \mathbf{T} \circ \rho)$.

Next, we specify what we mean by optimality in an autonomous geometric control structure (ν^A, ρ^A) and how this notion of optimality is related with the known notion of optimality in the associated geometric control structure (τ, ν, ρ). Consider a smooth real-valued function L^A on C. Using L^A, we can define a cost function on U by putting $L = p_C^* L^A$. A functional on the set of ρ^A-admissible curves is defined by:

$$\mathcal{J}(c^A) = \int_a^b L^A(c^A(t))dt,$$

where we have used the same notation as in the previous section since, clearly, $\mathcal{J}(c^A) = \mathcal{J}(c)$, with $c(t) = (t, c^A(t))$ the control associated with c^A. A ρ^A-admissible curve $c^A : [a,b] \to C$, taking a point q to q', is called *optimal* if, for any ρ^A-admissible curve $c^{A\prime} : [a,b] \to C$ taking q to q', we have $\mathcal{J}(c^A) \leq \mathcal{J}(c^{A\prime})$. Note that c^A is optimal with respect to L^A iff the associated control $c(t) = (t, c^A(t))$ taking (a,q) to (b,q') is optimal with respect to L. We say that c^A is *strong optimal* if, given any other control $c^{A\prime} : [a',b'] \to C$ taking q to q', then $\mathcal{J}(c^A) \leq \mathcal{J}(c^{A\prime})$. Note that the difference between optimality and strong optimality lies within the fact that the interval on which the control is defined may vary. Consider the following two immersed submanifolds P_i and P_f of M defined by $P_i \hookrightarrow \mathbb{R} \times \{p\}$ and $P_f \hookrightarrow \mathbb{R} \times \{q\}$, respectively. It is now easily seen that a ρ^A-admissible curve c^A, taking q to q' is strong optimal iff the associated control c is (P_i, P_f)-optimal with respect to $p_C^* L^A$. The problem of finding the (strong) optimal ρ^A-admissible curves taking q to q', is called the *autonomous (strong) optimal control problem*. The notion of strong optimality is important in time-optimal control problems, i.e. the function L^A is assumed to equal the constant function 1, which implies that, given any ρ^A-admissible curve $c^A : [a,b] \to C$, the cost functional takes the form: $\mathcal{J}(c^A) = b - a$. In this case, a ρ^A-admissible curve, taking q to q', is time-optimal if it connects q with q' in the least amount of time. One can also study autonomous optimal control problems where the endpoints may vary. However, since we will not encounter such optimal control problems with variable endpoints in the remainder of this thesis, these problems will not be discussed.

Finally, we shall consider some relations between objects belonging to the autonomous geometric control structure (ν^A, ρ^A) and objects belonging to the associated geometric control structure (τ, ν, ρ). Let ω_Q denote the canonical symplectic form on T^*Q, and consider the closed two-form ω_Q^A on $C \times_Q T^*Q$, which is the pull-back of ω_Q under the projection $p_{T^*Q} : C \times_Q T^*Q \to T^*Q$. It is easily seen that the $V_x\tau$ can be identified with T_qQ, where $p_Q(x) = q$, and similarly, the fibres of U equal to the fibres of C. In particular, if $s \in U_x$, then we also regard it as an element of C_q and if $\theta_x \in V_x^*\tau$, then we shall also write θ_q for θ_x, with $p_Q(x) = q$. Without mentioning it always explicitly, we shall identify $\nu^*V^*\tau$ with $\mathbb{R} \times \nu^*T^*Q$. We also agree, using these correspondences, that, given a multiplier $(\theta(t), \lambda)$ of a control $c(t) = (t, c^A(t))$, then

$\theta(t)$ can also be regarded as a one-form along c^A. Recall the locally defined function $h_\lambda = \gamma^i p_i + \lambda L$ on $\nu^* V^* \tau$, which appeared in the closed two-form $\omega_\lambda = -dh_\lambda \wedge dt + dp_i \wedge dx^i$. Since the bundle $\tau : M = \mathbb{R} \times Q \to \mathbb{R}$ is trivial, this function is globally well defined, and equals $h_\lambda(s, \theta_q) = \langle \theta_q, \rho^A(s) \rangle + \lambda L^A$.

We can now prove the following version of the maximum principle for autonomous optimal control problems.

Theorem 8.1 (The maximum principle). *If a ρ^A-admissible curve $c^A : [a, b] \to C$ is optimal with respect to L^A then there exists a piecewise smooth one-form $\theta(t)$ along $\widetilde{c}^A(t) = \nu^A(c^A(t))$ and a real number $\lambda \leq 0$ such that:*

1. *The following implicit Hamiltonian system is satisfied*
$$i_{(\dot{c}^A(t), \dot{\theta}(t))} \omega_Q^A = -dh_\lambda(c^A(t), \theta(t))$$
for all t where the ρ^A-admissible curve c^A is smooth,
2. *given any fixed $t_0 \in I$, the function $s \mapsto h_\lambda(s, \theta(t_0))$ on $C_{\widetilde{c}^A(t_0)}$ attains a global maximum for $s = c^A(t_0)$,*
3. *$h_\lambda(c^A(t), \theta(t)) = const.$ for all t and,*
4. *$(\theta(t), \lambda) \neq 0$ for all $t \in I$.*

If c^A is strong optimal then condition 3. should be replaced by the condition that $h_\lambda(c^A(t), \theta(t)) = 0$ for all t.

Proof. We already know that if c^A is optimal, then the associated control c is also optimal in the non-autonomous setting. This correspondence allows us to apply Theorem 6.1. The remainder of this proof consists of translating the necessary conditions from the time-dependent setting to the autonomous setting.

So, assume that (θ, λ) is a multiplier for the optimal control c. By definition this means that $\theta(t)$ is a piecewise smooth one-form along c^A and λ a constant with $\lambda \leq 0$,

1. $i_{(\dot{c}(t), \dot{\theta}(t))} \omega_\lambda = 0$ on every smooth part of the curve $(c, \theta)(t)$,
2. for all $t_0 \in I$, the function $\zeta^\lambda_{\theta(t_0)}$ defined by
$$s \mapsto \langle \zeta_\lambda(c(t_0), \theta(t_0)), \mathbf{T}(\rho(s)) \rangle + \lambda L(s) \text{ on } \nu^{-1}(\widetilde{c}(t_0))$$
attains a global maximum for $s = c(t_0)$,

3. $(\theta(t), \lambda) \neq 0$ for all $t \in I$.

The closed two-form ω_λ can be written (with a slight abuse of notation) as $\omega_Q^A - dh_\lambda \wedge dt$. Moreover, one can verify that $\zeta^\lambda_{\theta(t)}(s) = -h_\lambda(c^A(t), \theta(t)) + h_\lambda(s, \theta(t))$, all $t \in I$ and $s \in \nu^{-1}(\widetilde{c}(t))$. We conclude that the function $s \mapsto h_\lambda(s, \theta(t))$ on the fibres of C attains a global maximum for $s = c(t)$ or, using the above identification, for $s = c^A(t)$.

Given any tangent vector $v = v^A + v^0 \partial/\partial t \in T(U \times_M V^*\tau)$, with $v^A \in T(C \times_Q T^*Q)$ and $v^0 \in \mathbb{R}$, the following equation holds

$$i_v \omega_\lambda = i_{v^A} \omega_Q^A - dh_\lambda(v^A)dt + v^0 dh_\lambda.$$

If we substitute in the above equation $v = (\dot{c}(t), \dot{\theta}(t))$ then, since $v^0 = 1$ and $v^A = (\dot{c}^A(t), \dot{\theta}(t))$, we obtain $i_{v^A}\omega_Q^A = -dh_\lambda(c^A(t), \theta(t))$.

Since $(c^A(t), \theta(t))$ solves the implicit Hamiltonian system with Hamiltonian h_λ, the function h_λ is constant on every smooth part of the curve $(c^A(t), \theta(t))$. To complete the proof that $h_\lambda(c^A(t), \theta(t))$ is constant for all t, we will show that $h_\lambda(c^A(t), \theta(t))$ is continuous, even at those points t here c^A is discontinuous. This is done using similar arguments as in the proof of Proposition 5.2. For that purpose, let us assume that the curve c^A is discontinuous at $t = t_0$ and consider an adapted coordinate chart of C containing the point $c^A(t_0)$. In the coordinates we write $c^A(t) = (q^i(t), u^a(t))$, $\theta(t) = (t, q^i(t), p_i(t))$ and $s = (q^i(t_0), u^a)$. We then have

$$h_\lambda(q^i(t), u^a(t), p_i(t)) \geq h_\lambda(q^i(t), u^a, p_i(t)),$$

for all possible u^a. If we consider this inequality and take successively the limit from the left and from the right, for $t \to t_0$, we find that the function $h_\lambda(q^i(t), u^a(t), p_i(t))$ is indeed continuous at $t = t_0$.

It remains to prove, in the case of strong optimality, that h_λ is zero along the curve $(c^A(t), \theta(t))$. We now make use of Theorem 7.1. From the fact that

$$\zeta_\lambda(c(a), \theta(a)) = -h_\lambda(c^A(a), \theta(a))dt + p_i(a)dq^i \in (TP_i)^0 = V^*\tau,$$

we obtain $h_\lambda(c^A(a), \theta(a)) = 0$ and, hence, in view of Property 3 it follows that h_λ vanishes identically on $(c^A(t), \theta(t))$. □

A pair (θ, λ), with θ a piecewise smooth one-form along a ρ^A-admissible curve c^A and with $\lambda \in \mathbb{R}$, for which the properties mentioned in Theorem

8.1 hold, is also called a multiplier of c^A. By reversing the arguments in the proof of Theorem 8.1, it is easily seen that (θ, λ) is a multiplier for the control c. Similar to the non-autonomous case, we say that a ρ^A-admissible curve c^A is an *extremal* if there exists a multiplier (θ, λ) with $\lambda \leq 0$. If an extremal c^A admits a multiplier (θ, λ) with $\lambda = 0$, then it is called *abnormal*, and if it admits a multiplier with $\lambda < 0$, then it is called *normal*. If c^A is an abnormal extremal, which is not normal, then it is called a *strictly abnormal extremal*.

The closures of the vertical variational cones $\mathrm{cl}(V_y R_x)$ and $\mathrm{cl}(V_{\overline{y}} R_{\overline{x}})$ are cones in, respectively, $V_y\tau$ and $V_{\overline{y}}\overline{\tau}$, where $\overline{x} = (a, q, J_x)$ and $\overline{y} = (b, q', J_y)$, with $J_y - J_x = \mathcal{J}(c^A)$. The fibre $V_y\tau$ of $V\tau$ at y can be identified with $T_{q'}Q$ and, therefore, we shall consider the cone $V_y R_x$ as a cone in $T_{q'}Q$. Similarly, $V_{\overline{y}}\overline{\tau}$ can be identified with $T_{q'}Q \times \mathbb{R}$. Using these identifications, we can reformulate Theorem 6.2 as follows.

Theorem 8.2. *An extremal c^A is abnormal iff $\mathrm{cl}(V_y R_x) \neq T_{q'}Q$. An extremal c^A is strictly abnormal iff the boundary of the cone $\mathrm{cl}(V_{\overline{y}} R_{\overline{x}})$ in $T_{q'}Q \times \mathbb{R}$ does not contain $(0_{q'}, -1)$, where $0_{q'}$ denotes the zero vector in $T_{q'}Q$.*

In the literature the following denomination is often used. A ρ^A-admissible curve c^A is called an *extremal* if there exists a one-form θ along $\widetilde{c}^A$ and a real number λ such that conditions 1, 3 and 4 from Theorem 8.1 are satisfied. In particular, the maximality condition 2 is not required to hold. However, from the implicit Hamiltonian system that $(c^A, \theta(t))$ satisfies, one can deduce that the function h_λ attains a local extremum, i.e.

$$\frac{\partial h_\lambda}{\partial u^a}\big(c^A(t), \theta(t)\big) = 0, \text{ for all } a = 1, \ldots, k \text{ and } t \in I.$$

For that reason, a pair (θ, λ) is called a *local multiplier* of c^A, if the conditions 1, 3 and 4 of Theorem 8.1 are satisfied. Similarly we will then talk about a *local extremal* and *abnormal or normal local extremals*. It is an easy exercise to see that, if (θ_1, λ_1) and (θ_2, λ_2) are local multipliers for a given ρ^A-admissible curve c^A, then the pair $(r_1\theta_1 + r_2\theta_2, r_1\lambda_1 + r_2\lambda_2)$, with $r_1, r_2 \in \mathbb{R}$, is also a local multiplier for c^A. Therefore, the set of local multipliers can be given a linear structure.

In the remainder of this section we shall concentrate on *linear autonomous geometric optimal control structures*, i.e. autonomous geometric optimal control structures satisfying the additional conditions that (ν^A, ρ^A) is a linear anchored bundle, i.e. $\nu^A : C \to Q$ is a linear bundle and $\rho^A : C \to TQ$ is a linear bundle map. Let us denote the distribution generated by ρ^A on the base space Q, by D^A, i.e. $\operatorname{im}\rho^A = D^A$ (cf. Chapter I). We first study the necessary conditions encountered in the maximum principle (Theorem 8.1).

Consider the maximality condition from Theorem 8.1. Fix $\alpha_q \in T^*_q Q$. Then, $s \mapsto h_\lambda(s, \alpha_q)$, regarded as a function on C_q, attains a local extremum at $s = s_0$ iff, when putting $\alpha_q = (q^i, p_i)$ and $s_0 = (q^i, u^a)$, we have

$$\left.\frac{\partial}{\partial u^a}\right|_{s_0} \left({\rho^A}^i_a u^a p_i + \lambda L^A(q^i, u^a, p_i) \right) = 0, \tag{8.8}$$

or equivalently $\rho^{A*}(\alpha_x) = -\lambda \mathbb{F}L^A(s_0)$, where $\mathbb{F}L^A : C \to C^*$ is the so-called fibre derivative of L^A and $\rho^{A*} : T^*Q \to C^*$ is the dual of the linear bundle map ρ^A.

If c^A is an abnormal local extremal (i.e. if there exists a local multiplier (θ, λ) with $\lambda = 0$), then $\rho^{A*}(\theta(t)) = 0$ for any t or, equivalently, $\theta(t) \in D^0_{\widetilde{c}^A(t)}$ (the annihilator space of D at the point $\widetilde{c}^A(t)$). Moreover, in this specific case, the function $s \mapsto h_0(s, \theta(t))$ equals 0 for all $s \in C_{\widetilde{c}^A(t)}$, which implies that c^A is a abnormal extremal. We conclude that for linear autonomous control problems, the abnormal local extremals are abnormal extremals (i.e. they satisfy the maximality condition 2 from Theorem 8.1). This observation further implies that $V_y R_x$ is a linear subspace.

Theorem 8.3. *Let $c^A : [a, b] = I \to C$ denote a ρ^A-admissible curve, taking q to q', and let $c : I \to U$ be the associated control. Then the closure $\mathrm{cl}(V_y R_x)$ of the vertical variational cone, with $x = (a, q)$ and $y = (b, q')$ is a linear subspace of $T_{q'}Q$. Moreover, $\mathrm{cl}(V_y R_x) = V_y R_x$ and the dual cone $V^*_y R_x$ equals the annihilator space $(V_y R_x)^0$ of $V_y R_x$.*

Proof. In order to prove that $\mathrm{cl}(V_y R_x)$ is a linear subspace, it is sufficient to prove that $V^*_y R_x$ is a linear subspace of $T^*_{q'}Q$ (see Remark 3.2). In view of Corollary 5.7 any element in $V^*_y R_x$ corresponds to a multiplier (θ, λ) with $\lambda = 0$. In view of this correspondence, it is sufficient to prove that the set of multipliers of c has a linear structure. This in turn follows from the fact that any local multiplier (θ, λ) with $\lambda = 0$, is a multiplier and that the set of local multipliers is equipped with a linear structure. □

Since $V_y R_x$ is a linear subspace of $T_{q'}Q$, curves in $V_y R_x$ will have their tangent vector in $V_y R_x$. It is precisely this result that we will use to obtain sufficient conditions for an extremal *not to be abnormal.* For notational convenience, we consider a smooth ρ^A-admissible curve c^A, for which there exists a section σ of ν with $\sigma(t, \widetilde{c}^A(t)) = (t, c^A(t))$. Note that Proposition 1.1 precisely says that any ρ^A-admissible curve is a composition of smooth ρ^A-admissible curves of this type. Recall the notations and definitions from Remark 2.8. Using the section σ we can consider the time dependent vector field $X^A : \mathbb{R} \times Q \to TQ : z \mapsto \rho^A(p_C(\sigma(z)))$. Then $\widetilde{c}^A$ is an integral curve of X^A, since $X^A(t, \widetilde{c}^A(t)) = \dot{\widetilde{c}}^A(t)$ holds.

Consider an arbitrary section σ^A of ν^A. We can consider the time dependent vector fields $Y_t^{(i)}$, with $i \in \mathbb{N}_0$, recursively defined by $Y_t^{(0)} = \rho^A \circ \sigma^A$, $Y_t^{(1)} = [X_t^A, \rho^A \circ \sigma^A]$ and $Y_t^{(i)} = \dot{Y}_t^{(i-1)} + [X_t^A, Y_t^{(i-1)}]$, for $i \geq 2$. Let $\mathcal{I}(X^A)$ denote the differentiable distribution on $\mathrm{dom}(X^A) \subset M = \mathbb{R} \times Q$, which is defined as follows: for any $z \in \mathrm{dom}(X^A)$ the linear space $\mathcal{I}(X^A)_z$ is spanned by all tangent vectors of the form $Y^{(i)}(z)$, given any $Y_t^{(0)} = \rho^A \circ \sigma^A$ with $\sigma^A \in \Gamma(\nu^A)$ and $i \in \mathbb{N}_0$. (Note that we used the identification of $V_z\tau$ and $T_{p_Q(z)}Q$.)

Theorem 8.4. *The linear space $\mathcal{I}(X^A)_{(b,q')}$ is a subspace of V_yR_x.*

Proof. Let $\{\psi_t^{t'}\}$ denote the flow of the time dependent vector field X^A. By definition of V_yR_x, any tangent vector of the form

$$T\psi_t^b\Big(\big(\rho^A \circ \sigma^A\big)\big((\psi_t^b)^{-1}(q')\big)\Big), \text{for some } \sigma^A \in \Gamma(\nu^A) \text{ and } t \in]b-\epsilon, b]$$

is contained in V_yR_x. Let us fix any $\sigma^A \in \Gamma(\nu^A)$, then the above tangent vector, considered as a function of t, determines a curve $v(t)$ in V_yR_x through $\rho^A(\sigma^A(q'))$ at $t = b$. The tangent vector through this curve at any t is again contained in V_yR_x, implying that, using Remark 2.8,

$$T\psi_t^b\Big(\big[X_t^A, \rho^A \circ \sigma^A\big]\big((\psi_t^b)^{-1}(q')\big)\Big) = (\psi_b^t)^*(Y_t^{(1)})(q') = -\dot{v}(t) \in V_yR_x,$$

where $Y_t^{(i)}$ denotes the sequence of time dependent vector fields associated with σ^A. If we put $t = b$, then we obtain that $Y_b^{(1)}(q') = [X_b^A, \rho^A \circ \sigma^A](q') \in V_yR_x$. Using the formulae, mentioned in Remark 2.8, we have that:

$$\frac{d^{i-1}}{dt^{i-1}}\Big((\psi_b^t)^*(Y_t^{(1)})(q')\Big) = \Big((\psi_b^t)^*(Y_t^{(i)})(q')\Big) \in V_yR_x, \text{ for any } i \geq 2.$$

This concludes the proof. ∎

Note that, in case X^A can be chosen to be time independent (this is possible if $\dot{\tilde{c}}^A(t) \neq 0$), each vector field $Y_t^{(i)}$ introduced above is time independent and equals:

$$Y^{(i)} = \underbrace{[X^A, [\cdots, [X^A, \rho^A \circ \sigma^A] \cdots]]}_{i}.$$

The distribution $\mathcal{I}(X^A)$ can the be regarded as a distribution on Q. The above theorem says that, given any section σ^A of ν^A, all iterated Lie brackets of X^A with $\rho^A \circ \sigma^A$ evaluated at q', are contained in $V_y R_x$. This result provides us with sufficient conditions for a ρ^A-admissible curve not to be abnormal. For the following theorem we assume that c^A is a ρ^A-admissible curve, for which the control $c(t) = (t, c^A(t))$ is induced by an ordered family $\mathcal{S} = (\sigma_\ell, \ldots, \sigma_1)$ of sections of ν, i.e. $c(t) = \sigma_i(\tilde{c}(t))$ for $t \in [a_{i-1}, a_i]$. Let X_i^A denote the time dependent vector field defined by $\rho^A \circ p_C \circ \sigma_i$.

Theorem 8.5. *The ρ^A-admissible curve $c^A : [a,b] \to C$ can not be an abnormal extremal if there exists a $t \in [a_{i-1}, a_i]$ with $i = 1, \ldots, \ell$, for which the following condition holds:*

$$\mathcal{I}\big(X_i^A\big)_{(t,\tilde{c}^A(t))} = T_{\tilde{c}^A(t)} Q.$$

Proof. Recall Property 2.3 (page 71), saying that the variational cones satisfy the inclusion property. Therefore, if the restriction $c^A|_{[t_0,t_1]}$ to some interval $[t_0, t_1] \subset [a,b]$ is not abnormal (i.e. if $V_{(t_1,\tilde{c}^A(t_1))} R_{(t_0,\tilde{c}^A(t_0))}$ equals $T_{(t_1,\tilde{c}^A(t_1))} Q$), then c^A itself can not be abnormal. Using the previous theorem, we know that the ρ^A-admissible curve $c^A|_{[a_{i-1},t_1]}$, for any $t_1 \in]a_{i-1}, a_i]$, is not abnormal if $\mathcal{I}\big(X_i^A\big)_{(t,\tilde{c}^A(t))} = T_{\tilde{c}^A(t)} Q$. Note that $z \mapsto \mathcal{I}(X_i^A)_z$ is a differentiable distribution, and therefore its rank is a lower semicontinuous function. Thus, if

$$\dim \Big(\mathcal{I}\big(X_i^A\big)_{(a_{i-1},\tilde{c}^A(a_{i-1}))}\Big) = n,$$

then it has maximal rank n at all points $(t, \tilde{c}^A(t))$ for $t \in]a_{i-1}, a_i]$ sufficiently close to a_{i-1}. This concludes the proof. □

We now continue with studying the normal local extremals. If c^A is a normal local extremal i.e. if there exists a multiplier (θ, λ) with $\lambda = -1$, then the equation $\rho^{A*}(\theta(t)) = \mathbb{F}L^A(c^A(t))$ holds. Assume that L^A is a *regular cost*, i.e. if the fibre derivative $\mathbb{F}L^A : C \to C^*$ of L^A is invertible. Then the curve $\theta(t)$ in T^*Q *generates* $c^A(t)$, since $c^A(t) = (\mathbb{F}L^A)^{-1}(\rho^{A*}(\theta(t)))$. In this case, since θ is piecewise smooth, the curve c^A generated by θ is also continuous. Therefore, a normal local extremal is a piecewise smooth curve in C. Moreover, from the following proposition it follows that it is in fact a smooth curve.

Theorem 8.6. *Assume that L^A denotes a regular cost. Every normal local extremal c^A is generated by an integral curve $\theta(t)$ of the Hamiltonian vector field X_G on T^*Q corresponding to the Hamiltonian*

$$G(\alpha_q) = h_\lambda((\mathbb{F}L^A)^{-1}(\rho^{A*}(\alpha_q), \alpha_q),$$

for $\alpha_q \in T_q^*Q$ and with $\lambda = -1$. The converse also holds, i.e. every integral curve of X_G generates a normal local extremal.

Proof. We assume that c^A is a normal local extremal and fix a local multiplier (θ, λ) with $\lambda = -1$. Consider the function $\mathcal{L}^A : T^*Q \to C \times_Q T^*Q$, defined by:

$$\mathcal{L}^A(\alpha_q) = \left((\mathbb{F}L^A)^{-1}(\rho^{A*}(\alpha_q)), \alpha_q\right).$$

Note that $\mathcal{L}^A$ is a section of the bundle $p_{T^*Q} : C \times_Q T^*Q \to T^*Q$, with p_{T^*Q} the projection onto the second factor. Then it is easily seen that $\mathcal{L}^{A*}h_\lambda = G$ and that $\mathcal{L}^{A*}\omega_Q^A = \omega_Q$. Recall the implicit Hamiltonian system (condition 1 in Theorem 8.1): a multiplier has to satisfy: $i_{(\dot{c}^A(t),\dot{\theta}(t))}\omega_Q^A = -dh_\lambda(c^A(t), \theta(t))$, and consider the tangent vectors

$$v = \left(\dot{c}^A(t), \dot{\theta}(t)\right) = T\mathcal{L}^A\left(\dot{\theta}(t)\right),$$

and $w = T\mathcal{L}^A(\overline{w}) \in T(C \times_Q T^*Q)$ for $\overline{w} \in T(T^*Q)$ arbitrary. Then, the following equation holds $\omega_Q^A(v, w) = \omega_Q(\dot{\theta}(t), \overline{w})$. By substituting $\mathcal{L}^{A*}h_\lambda = G$ we obtain that $i_{\dot{\theta}(t)}\omega_Q = -dG(\theta(t))$ for every smooth part of θ. By uniqueness of solutions to differential equations, it follows that θ is smooth (and therefore we have that c^A is smooth).

On the other hand, assume that $i_{\dot{\theta}(t)}\omega_Q = -dG(\theta(t))$ and consider the smooth curve $c^A(t) = (\mathbb{F}L^A)^{-1}(\rho^{A*}(\theta(t)))$ in C. Then, by reversing the above arguments, we obtain that $\omega_Q^A(v, w) = -dh_\lambda(w)$ with

$$v = (\dot{c}^A(t), \dot{\theta}(t)) = T\mathcal{L}^A(\dot{\theta}(t))$$

and $w = T\mathcal{L}^A(\overline{w}) \in T(C \times_Q T^*Q)$ for $\overline{w} \in T(T^*Q)$ arbitrary and $\lambda = -1$. It remains to check that this is also valid for arbitrary $w \in T(C \times_Q T^*Q)$.

Since $\mathcal{L}^A$ is a section of p_{T^*Q}, any element w in $T(C \times_Q T^*Q)$ can be written as $w = T\mathcal{L}^A(\overline{w}) + \hat{w}$ where $\overline{w} = Tp_{T^*Q}(w) \in T(T^*Q)$ and $\hat{w} \in \ker Tp_{T^*Q}$. Since $\omega_Q^A = p_{T^*Q}^*\omega_Q$ it is easily seen that $i_{\hat{w}}\omega_Q^A = 0$. From this we conclude that $i_v\omega_Q^A = -dh_\lambda(c^A(t), \theta(t))$ for arbitrary t. □

9 Some applications of the maximum principle in geometric mechanics

As an illustration of the formalism developed above, we discuss in this section some applications of optimal control theory to mechanics. We first deduce that the solutions to the Euler-Lagrange equations of a given Lagrangian L are precisely the extremals of a trivial optimal control problem with cost L. We then consider the "vakonomic approach" to Lagrangian systems subjected to (regular) nonholonomic constraints and, subsequently, we also consider the formulation of Lagrangian systems on (affine) Lie algebroids.

9.1 Lagrangian systems

A *Lagrangian system* (Q, L) is a pair, where Q denotes an n-dimensional manifold, the configuration space, and where L is a smooth function on TQ called a Lagrangian (for further details, see e.g. [37]). Let us fix a Lagrangian system and consider the following (trivial) linear autonomous geometric control structure $(\tau_Q, \mathrm{id}_{TQ})$, where $\tau_Q : TQ \to Q$ and id_{TQ} is the identity transformation on TQ. The Lagrangian L can be interpreted as a cost function. The following properties are easily verified. Since the anchor map id_{TQ} is the identity, for each $q \in Q$ the set of reachable points R_q equals Q (where we assume that Q is connected). Every piecewise smooth curve c in Q is the base curve of the (id_{TQ})-admissible curve $\dot{c}$ (we deviate from the notations used in the previous sections of this chapter). Since $\mathrm{im}(\mathrm{id}_{TQ}) = TQ$, no abnormal extremals can occur (cf. Theorem 8.5). Assume that $\dot{c}$ is a (normal) extremal. Let $(\theta(t), -1)$ denote a multiplier and consider the local maximum condition (Equation 8.8) from the previous section, with $\rho^A = \mathrm{id}_{TQ}$:

$$\theta(t) = \mathbb{F}L(\dot{\tilde{c}}(t)).$$

We now reformulate Theorem 8.1, for the above specific case considered here. Note that, since an autonomous geometric control structure is an anchored bundle, the following reformulation of the maximum principle holds for any other linear autonomous geometric control structure (ν^A, ρ^A) on Q for which there exists a linear anchored bundle isomorphism between (ν^A, ρ^A) and $(\tau_{TQ}, \mathrm{id}_{TQ})$. Let ω_L denote the closed two-form on TQ, which is defined as the pull-back of the canonical symplectic two-form ω on T^*Q under $\mathbb{F}L : TQ \to T^*Q$, i.e. $\omega_L = \mathbb{F}L^*\omega$. Let E_L denote the energy function associated with L, and which is defined by $E_L = \Delta(L) - L$, where Δ denotes the dilation vector field on the linear bundle τ_Q (cf. Chapter II).

Theorem 9.1. *Let c denote a piecewise smooth curve in Q. Then $\dot{c}$ is an local extremal iff c satisfies the equation*

$$i_{\ddot{c}(t)}\omega_L = -dE_L(\dot{c}(t)),$$

at all points where c is smooth.

Proof. Assume that $\dot{c}$ is a local extremal, with multiplier $(\theta, -1)$. Recall Theorem 8.1, then $\theta(t)$ satisfies, for all t where c is smooth:

$$i_{(\dot{\theta}(t), \ddot{c}(t))}\omega_Q^A = -dh_{-1}(\theta(t), \dot{c}(t)).$$

Now, consider the map

$$TQ \to TQ \times_Q T^*Q : v \mapsto (v, \alpha_q) = (v, \mathbb{F}L(v)),$$

which can be used to pull-back ω_Q^A to a closed two-form on TQ. It is easily seen that this closed two-form is precisely ω_L. The pull-back of h_{-1} under this map equals: $h_{-1}(v, \mathbb{F}L(v)) = \langle \mathbb{F}L(v), v\rangle - L(v) = E_L(v)$, for any $v \in TQ$. Therefore, since $\theta(t) = \mathbb{F}L(\dot{c}(t))$, we easily obtain that

$$i_{(\dot{\theta}(t), \ddot{c}(t))}\omega_Q^A = -dh_{-1}(\theta(t), \dot{c}(t))$$

implies

$$i_{\ddot{c}(t)}\omega_L = -dE_L(\dot{c}(t)).$$

The converse is proven using the same techniques as in the proof of Theorem 8.6. Let p_1 denote the projection of $TQ \times_Q T^*Q$ onto the first factor. It is sufficient to prove that, if $w \in T_{(\theta(t), \dot{c}(t))}(TQ \times_Q T^*Q)$, with $Tp_1(w) = 0$, then $i_w(i_{(\dot{\theta}, \ddot{c})}\omega_Q^A) = -w(h_{-1})$. This easily follows from local coordinate expressions. □

Note that in the case where L is a regular Lagrangian (i.e. $\mathbb{F}L$ is invertible), then ω_L is a symplectic two-form. This, in turn, implies that there exists a unique vector field $\Gamma \in \mathfrak{X}(TQ)$, such that $i_\Gamma \omega_L = -dE_L$. The local extremals are then precisely the integral curves of Γ (which implies that they are smooth).

9.2 Vakonomic approach to Lagrangian systems subjected to nonholonomic constraints

Let (Q, L) again denote a Lagrangian system with $\dim Q = n$, and let C be a k-dimensional (not necessarily linear) subbundle of TQ, representing some nonholonomic constraints to which the Lagrangian system is subjected. Locally, C can be described as the common zero-level set $\Phi^\alpha(q^i, v^i) = 0$ of $n - k$ independent function $\Phi^\alpha : TQ \to \mathbb{R}^{n-k}$. The subbundle condition entails that $\mathrm{rank}(\partial\Phi^\alpha/\partial v^j) = n - k$. Using the implicit function theorem, there exist $n - k$ functions such that, $v = v^i \partial/\partial q^i|_q \in C$ or $\Phi^\alpha(q, v) = 0$ iff $v^\alpha = \Psi^\alpha(q^i, v^a)$ for $a = 1, \ldots, k$ and $\alpha = k+1, \ldots, n$.

Let $i : C \hookrightarrow TQ$ represent the canonical injection and $\nu : C \to Q$ the restriction of τ_Q to C. The pair (ν, i) is a (not necessarily linear) anchored bundle. If we then put $\widetilde{L} = i^*L \in C^\infty(C)$ and regard $\widetilde{L}$ as a cost function, we obtain an autonomous geometric control structure. The base curves in Q of extremals are called the solutions to the Lagrangian system (Q, L) subjected to the nonholonomic constraint bundle C. Note that these solutions do not necessarily satisfy the equations of the nonholonomic mechanical system with Lagrangian L (cf. Section IV-4). Let c denote a curve in Q, such that $\dot{c} \in C$. The equations from Theorem 8.1, characterising a local normal extremal, correspond to the ones obtained in [10]. Indeed, if we express the equation $i_{(\ddot{c}(t),\dot{\theta}(t)}\omega_Q^A = -dh_{-1}(\dot{c}(t), \theta(t)) = 0$ in a local coordinate neighbourhood, we obtain, putting $c(t) = (q^i(t))$ and $\theta(t) = (q^i(t), p_i(t))$

$$\begin{cases} \dfrac{d}{dt}\left(\dfrac{\partial \widetilde{L}}{\partial v^a}\right) - \dfrac{\partial \widetilde{L}}{\partial q^a} = p_\alpha \left(\dfrac{d}{dt}\left(\dfrac{\partial \Psi^\alpha}{\partial v^a}\right) - \dfrac{\partial \Psi^\alpha}{\partial q^a}\right) + \dot{p}_\alpha \dfrac{\partial \Psi^\alpha}{\partial v^a}, \\ \dot{p}_\alpha(t) = \dfrac{\partial \widetilde{L}}{\partial q^\alpha} - p_\beta \dfrac{\partial \Psi^\beta}{\partial q^\alpha}, \\ \dot{q}^\alpha = \Psi^\alpha(q^i, \dot{q}^a), \end{cases}$$

where we used the local extremum condition: $\partial\widetilde{L}/\partial v^a = p_a + p_\alpha \partial\Psi^\alpha/\partial v^a$, from which it follows that the k momenta p_a can be expressed in terms of the $n-k$ momenta p_α. This implies that the multiplier θ only depends on the $n-k$ momenta p_α. If we want to recover a more familiar formulation of these equations, note that they can be rewritten in an equivalent form as

$$\begin{cases} \dfrac{d}{dt}\left(\dfrac{\partial L}{\partial v^i}\right) - \dfrac{\partial L}{\partial q^i} = \dot{\lambda}_\alpha \dfrac{\partial \widetilde{\Phi}^\alpha}{\partial v^i} + \lambda_\alpha \left(\dfrac{d}{dt}\left(\dfrac{\partial \widetilde{\Phi}^\alpha}{\partial v^i}\right) - \dfrac{\partial \widetilde{\Phi}^\alpha}{\partial q^i}\right), \\ 0 = \widetilde{\Phi}^\alpha(q^i, \dot{q}^a), \end{cases}$$

with $\widetilde{\Phi}^\alpha(q^i, v^i) = \Psi^\alpha(q^i, v^a) - v^\alpha$ and $\lambda_\alpha = p_\alpha - \partial L/\partial v^\alpha$. The functions λ_α are often called *Lagrangian multipliers*. In view of the above correspondence between the p_α and the λ_α, the denomination "multiplier" for $\theta(t) = (q^i(t), p_i(t))$ is justified.

9.3 Regular Lagrangian systems on (affine) Lie algebroids

Assume that $\nu : C \to Q$ is a Lie algebroid with anchor map $\rho : C \to TQ$. In particular, (ν, ρ) is a linear anchored bundle. It is a well known fact that the Lie algebroid structure determines a Poisson structure on the dual bundle C^*, where we shall denote the bracket on $C^\infty(C^*)$ by $\{\cdot,\cdot\}_{C^*}$ and the canonical Poisson bracket on $C^\infty(T^*Q)$ will be denoted by $\{\cdot,\cdot\}$. It is also well-known that both Poisson structures are ρ^* connected, i.e. if $\chi = \rho^*$, then, given arbitrary $f, g \in C^\infty(C^*)$ the following equality holds: $\{\chi^* f, \chi^* g\} = \chi^*\{f, g\}_{C^*}$. This implies that any Hamiltonian vector field X_f on C^* is ρ^*-connected to the Hamiltonian vector field $X_{\chi^* f}$ on T^*Q.

Let $L \in C^\infty(C)$ be a regular cost and consider the function G on T^*Q, introduced in the Section 8 (Theorem 8.6). Define $g \in C^\infty(C^*)$ by $g(\alpha) = \langle \alpha, \mathbb{F}L^{-1}(\alpha)\rangle - L(\mathbb{F}L^{-1}(\alpha))$ with $\alpha \in C^*$. Then, it is easily seen that $\chi^* g = G$. This guarantees that, given any integral curve $\theta(t)$ of X_G, then $\rho^*(\theta(t))$ is an integral curve of X_g and, conversely, any integral curve $\alpha(t)$ of X_g, through a point in the image of ρ^*, is the projection under ρ^* of an integral curve of X_G. From this we conclude that, in the case where C^* is a Lie-algebroid, the integral curves of X_g through a point in the image of ρ^* are projections of normal local extremals. In [40], one can find a detailed treatment of how the above defined "Hamiltonian" system on C^* can be translated into a "Lagrangian" system on C, using the "Legendre" transformation $\mathbb{F}L : C \to C^*$. In particular, one has to generalise the concepts of dilation vector field and vertical endomorphism to the case of a Lie algebroid, which enables one to consider a Poincaré-Cartan two-form associated to the function L.

Here, we shall pay special attention to the case of Lagrangian equations on affine Lie algebroids, which have been studied recently in [41], with a view of deriving a time dependent version of Lagrange equations on Lie algebroids. We follow the discussion of [16, 41]. However, for the sake of simplicity, we only present a local treatment.

Let (ν, ρ) denote an affine anchored bundle on M, i.e. E is an affine bundle over M, with anchor map an affine bundle map from E to TM. The linear bundle, on which E is modeled, is denoted by V. The affine bundle E is called an *affine Lie algebroid* if V is equipped with a Lie algebroid structure (see page 52), with bracket $[\cdot,\cdot]$, and if there exists an action D of $\Gamma(E)$ on $\Gamma(V)$ satisfying the following conditions: for all $f \in C^\infty(M)$, $\sigma \in \Gamma(E)$ and $\zeta, \zeta_1, \zeta_2 \in \Gamma(V)$

1. $D_\sigma f\zeta = \rho(\sigma)(f)\zeta + f D_\sigma \zeta$;
2. $D_\sigma(\zeta_1 + \zeta_2) = D_\sigma \zeta_1 + D_\sigma \zeta_2$;
3. $D_{\sigma+\zeta_1}\zeta_2 = D_\sigma \zeta_2 + [\zeta_1, \zeta_2]$;
4. $D_\sigma[\zeta_1, \zeta_2] = [D_\sigma \zeta_1, \zeta_2] + [\zeta_1, D_\sigma \zeta_2]$.

It should be mentioned that there is a certain redundancy in the above condition, i.e. it is sufficient that the bracket on $\Gamma(V)$ is skew-symmetric. The fact that the bracket on $\Gamma(V)$ determines a Lie algebroid structure ($\mathbb{R}$-linear, the Jacobi identity and the compatibility property) can be deduced from 1 to 4. It should also be noted that, upon replacing ζ_2 by $f\zeta_2$ in 4, one obtains that $\rho(D_\sigma \zeta_1) = [\rho(\sigma), \rho(\zeta_1)]$, where we make no notational distinction between the affine map ρ and its linear part, i.e. the linear bundle map $\overline{\rho} : V \to TM$, induced by ρ, is also denoted by ρ. Take a bundle adapted coordinate system (q^i, u^a) on V and fix one local section of E, say σ_0. Then (q^i, u^a) determines a coordinate system on E, where $(q^i, 0)$ corresponds to the element $\sigma_0(q)$. A ρ-admissible curve $c : [a, b] \to E : t \mapsto (q^i(t), u^a(t))$ satisfies the relations:

$$\dot{q}^i(t) = \gamma_0^i(q(t)) + \gamma_a^i(q(t))u^a(t),$$

The above defined affine Lie algebroid structure is determined by the local structure functions C_{ab}^c of the Lie algebroid structure on $\Gamma(V)$ and, in addition, the local functions C_{0b}^c that are defined by,

$$D_{\sigma_0}\zeta_b(q) = C_{0b}^c(q)\zeta_c(q).$$

Let $L \in C^\infty(U)$ denote a regular cost on U. Then, the Lagrangian equations, for a given ρ-admissible curve c, derived in [41], can be written as:

$$\dot{q}^i = \gamma_0^i + \gamma_a^i u^a,$$
$$\frac{d}{dt}\left(\frac{\partial L}{\partial u^a}\right) = \gamma_a^i \frac{\partial L}{\partial q^i} + \left(C_{0a}^b + C_{ca}^b u^c\right)\frac{\partial L}{\partial u^b}.$$

We now prove prove that these equations can also be obtained with our formalism, from the maximum principle (cf. Theorem 8.1). Indeed, let c be a normal local extremal. Then: $i_{(\dot{c}(t),\dot{\theta}(t))}\omega_M^A = -dh_{-1}(c(t),\theta(t))$ holds, where $(\theta(t),-1)$ is a local multiplier. In coordinates, with $\theta(t) = (q^i(t), p_i(t))$:

$$\dot{q}^i = \gamma_0^i + \gamma_a^i u^a;$$
$$\dot{p}_i = -p_j\left(\frac{\partial \gamma_0^j}{\partial q^i} + \frac{\partial \gamma_a^j}{\partial q^i}u^a\right) + \frac{\partial L}{\partial q^i};$$
$$\frac{\partial L}{\partial u^a} = p_i \gamma_a^i.$$

Multiplying the second equation with γ_a^i and summing over i, we obtain:

$$\frac{d}{dt}\left(p_i\gamma_a^i\right) = p_i \frac{\partial \gamma_a^i}{\partial q^j}\left(\gamma_0^j + \gamma_b^j u^b\right) - p_j\left(\gamma_a^i \frac{\partial \gamma_0^j}{\partial q^i} + \gamma_a^i \frac{\partial \gamma_b^j}{\partial q^i}u^b\right) + \gamma_a^i \frac{\partial L}{\partial q^i}$$

If we consider the coefficient of p_i, and substitute the identities

$$C_{ab}^c \gamma_c^i = \gamma_a^j \frac{\partial \gamma_b^i}{\partial q^j} - \gamma_b^j \frac{\partial \gamma_a^i}{\partial q^j};$$
$$C_{0b}^c \gamma_c^i = \gamma_0^j \frac{\partial \gamma_b^i}{\partial q^j} - \gamma_b^j \frac{\partial \gamma_0^i}{\partial q^j},$$

we obtain, putting $\partial L/\partial u^a = \gamma_a^i p_i$:

$$\frac{d}{dt}\left(\frac{\partial L}{\partial u^a}\right) = \gamma_a^i \frac{\partial L}{\partial q^i} + \left(C_{0a}^b + C_{ca}^b u^c\right)\frac{\partial L}{\partial u^b}.$$

These, together with the relations $\dot{q}^i = \gamma_0^i(q) + \gamma_a^i(q)u^a$, are precisely the Lagrangian equations mentioned above. For further developments on the geometric formulation of Lagrangian mechanics on affine Lie algebroids we refer to [41].

9.4 Connection control systems

We now consider the so called affine connection control systems, introduced by A.D. Lewis in [34, 35]. The number of applications that fit in this geometric formulation seems to be very large, and we will only mention some of them. For further examples, we refer the reader to the work of A.D. Lewis. Let us first specify what we mean by an affine connection control system.

Definition 9.1. An *affine connection control system* consists of a linear anchored bundle (ν, ρ) on Q, and, in addition, a connection ∇ on Q.

To avoid any confusion, we mention explicitly that ∇ is a standard linear connection. The above definition is slightly different from the one given in [35]. Denote the control domain by C and let Γ denote the second order vector field on TQ, representing the geodesic spray of ∇, i.e. if Γ^i_{jk} are the local connection coefficients of ∇, then

$$\Gamma(q, v) = v^i \frac{\partial}{\partial q^i} + \Gamma^i_{jk} v^j v^k \frac{\partial}{\partial v^i}.$$

We always assume that $\operatorname{im}\rho$ is a regular distribution. Consider the autonomous control system (ν', ρ') on TQ, defined by $\nu' : C' = \tau_Q^* C \to TQ$ and $\rho'(v_q, s) = \Gamma(v_q) + (\rho(u))^V_{v_q}$, where $(w_q)^V_{v_q}$ denotes the vertical lift of the tangent vector $w_q \in T_qQ$ to $T_{v_q}TQ$, i.e. $(w_q)^V_{v_q} = d/dt(v_q + tw_q)|_{t=0}$. Let $c' : I \to C'$ denote a ρ'-admissible curve, then by definition, the curve $c'(t) = (v(t), c(t))$, with $v(t) \in TQ$ and $c(t) \in C$, satisfies:

$$\dot{v}(t) = \Gamma(v(t)) + (\rho(c(t)))^V_{v(t)}.$$

First, if we project this equation on Q using $T\tau_Q$, we obtain that, for $\tau_Q(v(t)) = \tilde{c}(t)$, $\dot{\tilde{c}}(t) = v(t)$. Secondly, since $\dot{v}(t) - \Gamma(v(t))$ is vertical to τ_Q and since the second term on the right side is also vertical, we can use the identification of $V_{w_q}\tau_Q$ with T_qQ and, from the definition of the connection map K associated with ∇ (cf. Section II-1), we have that

$$\nabla_{\dot{\tilde{c}}(t)} \dot{\tilde{c}}(t) = \rho(c(t)).$$

Therefore, a curve $(v(t), c(t))$ is ρ'-admissible iff $v(t) = \dot{\tilde{c}}(t)$ and $\nabla_{\dot{\tilde{c}}(t)} \dot{\tilde{c}}(t) = \rho(c(t))$. The above equation is encountered in the study of nonholonomic mechanical systems, which will be discussed in Chapter V. Examples include the upright rolling penny, the snakeboard, the rolling racer, ... (see [34]).

When considering optimal affine connection control problems, the case with a constant Lagrangian function $L \in C^\infty(C')$, i.e. $L = 1$, is often encountered. Thus, $\mathcal{J}(c) = b - a$ if c is defined on $[a,b]$. The strong optimal controls, with starting point $v \in T_qQ$ and endpoint $v' \in T_{q'}Q$, in this autonomous geometric control structure, are often called *time-optimal*, which is clearly justified when considering the cost functional. (One could also consider the problem of finding time-optimal controls for which the velocities at starting point q and endpoint q' are variable, i.e. a control which is $(T_qQ, T_{q'}Q)$-time-optimal.) The maximum principle, developed in Theorem 8.1, can now be translated in the following form. Let the curvature and torsion of ∇ be denoted by R and T, respectively. Consider the following tensors R^* and T^*, associated to R and T:

$$\langle R^*(\alpha, v_2)(v_3), v_1\rangle = \langle \alpha, R(v_1, v_2)(v_3)\rangle;$$
$$\langle T^*(\alpha, v_2), v_1\rangle = \langle \alpha, T(v_1, v_2)\rangle.$$

In order to write down the equations of a normal local extremal, we need to draw attention to the following point. Let θ denote a section of the bundle $(\operatorname{im}\rho)^0$, then we define a map $\Pi(\theta,\cdot) : \Gamma(\operatorname{im}\rho) \to \mathfrak{X}^*(Q)$ according to:

$$\langle \Pi(\theta, X), Y\rangle = \langle \theta, \nabla_Y X\rangle,$$

for all $X \in \Gamma(\operatorname{im}\rho)$ and $Y \in \mathfrak{X}(Q)$. In local coordinates, with $\theta = (q^i, p_i)$, this map reads:

$$\Pi(\theta, \gamma_a^i u^a \frac{\partial}{\partial q^i}) = p_i \left(u^a \frac{\partial \gamma_a^i}{\partial q^j} - \Gamma_{jk}^i \gamma_a^k u^a \right) dq^j,$$

and is clearly tensorial. We no apply Theorem 8.1. A local extremal, in the above defined time-optimal control problem, is a curve $c : [a,b] \to C$ for which there exists a piecewise smooth one-form θ along $\widetilde{c}$, which is contained in $(\operatorname{im}\rho)^0$, and which satisfies the equation (after a tedious computation):

$$\begin{aligned}\nabla^2_{\dot{\widetilde{c}}(t)}\theta(t) + R^*\big(\theta(t), \dot{\widetilde{c}}(t)\big)\big(\dot{\widetilde{c}}(t)\big) - T^*\big(\nabla_{\dot{\widetilde{c}}(t)}\theta(t), \dot{\widetilde{c}}(t)\big)& \\ = \Pi\big(\theta(t), \rho(u(t))\big).&\end{aligned} \tag{9.9}$$

In addition, putting $\eta(t) := 1/2T^*(\theta(t), \dot{\widetilde{c}}(t)) - \nabla_{\dot{\widetilde{c}}(t)}\theta(t)$, the following properties should hold:

1. the triple $(\lambda, \eta(t), \theta(t)) \neq 0$ for all $t \in [a,b]$;
2. $\langle \eta(t), \dot{\widetilde{c}}(t)\rangle + \langle \theta(t), \rho(u(t))\rangle + \lambda = 0$ for all $t \in [a,b]$;

3. if c is $(T_qQ, T_{q'}Q)$-optimal, then $\theta(a) = \theta(b) = 0$.

Remark 9.2. A affine connection control system is an autonomous geometric control structure on TQ, for which multipliers are curves in T^*TQ. Using the decomposition of $T^*TQ = T^*Q \times_{TQ} T^*Q$, determined by the connection ∇, any multiplier will determine two curves in T^*Q, both projecting onto $\widetilde{c}$. These curves are precisely θ and η. In coordinates, this decomposition is given by:

$$T^*TQ \to T^*Q \times_{TQ} T^*Q : p_i dq^i|_v + P_i dv^i|_v \mapsto \\ \Big(\big(p_i + \tfrac{1}{2}(\Gamma^j_{ik} + \Gamma^j_{ki}) P_j v^k\big) dq^i|_q, P_i dq^i|_q\Big),$$

with $v \in T_qQ$.

Another example of an optimal control problem is the case where Q carries a Riemannian metric g. Let the cost function $L \in C^\infty(C)$ be defined by:

$$L(u) = \tfrac{1}{2} g\big(\rho(u), \rho(u)\big).$$

The optimal controls for this cost are called force minimising controls (see [35]). To obtain the differential equations that a multiplier has to satisfy, one has to consider the orthogonal projection π_D of TQ onto $D = \operatorname{im}\rho$, with respect to the metric g. The local extremum condition, is then expressed by $\pi_D(\sharp_g(\theta(t))) = \rho(u)$. Even more, we can define the $(2,0)$-tensor h_Q as

$$h_Q(\alpha_q, \beta_q) = g(\pi_D(\sharp_g(\alpha_q), \pi_D(\sharp_g(\beta_q))).$$

The equation a normal extremal has to satisfy becomes:

$$\nabla^2_{\dot{\widetilde{c}}(t)} \theta(t) + R^*\big(\theta(t), \dot{\widetilde{c}}(t)\big)\big(\dot{\widetilde{c}}(t)\big) - T^*\big(\nabla_{\dot{\widetilde{c}}(t)} \theta(t), \dot{\widetilde{c}}(t)\big) \\ = \tfrac{1}{2}(\nabla h_Q)(\theta(t), \theta(t)) + T^*\big(\theta(t), \sharp_{h_Q}(\theta(t))\big); \tag{9.10}$$

The equation for an abnormal multiplier is precisely Equation 9.9.

IV

Sub-Riemannian geometry

A sub-Riemannian structure on a manifold is a generalisation of a Riemannian structure in that a metric is only defined on a proper vector subbundle of the tangent bundle to the manifold (i.e. on a regular distribution), rather than on the whole tangent bundle. As a result, in sub-Riemannian geometry a notion of length can only be assigned to a certain privileged set of curves, namely curves that are tangent to the given regular distribution on which the metric is defined. The problem then arises to find those curves that minimise length, among all curves connecting two given points. The characterisation of these length minimising curves is one of the main research topics in sub-Riemannian geometry, which has also interesting links to control theory and to vakonomic dynamics (for the latter, see for instance J. Cortés, et al. [10]).

The main goal of this chapter is, on the one hand, to present an application to the theory of generalised connections (cf. Chapter II) in sub-Riemannian geometry and, on the other hand, to illustrate the necessary and sufficient conditions on abnormal and (strictly) abnormal extremals, developed in the previous chapter. Section 1 contains general definitions concerning sub-Riemannian geometry. In particular, we will associate to a sub-Riemannian geometry a geometric optimal control structure, for which the optimal controls are precisely the length minimising curves, tangent to the given distribution. We recall the notion a normal and an abnormal extremal. In Section 2, we will consider some aspects of the theory of generalised connections in the framework of sub-Riemannian geometry. Then, normal extremals will appear as "geodesics" and abnormal extremals as "base curves of parallel transported sections" in the cotangent bundle with respect to a suitable generalised connection associated to the sub-Riemannian structure. Apart from shedding some new light on certain elements of sub-Riemannian geometry, this formulation also allows us to prove some known results in a more elegant way.

The main subtlety in studying length minimising curves of a sub-Riemannian structure lies in the existence of "abnormal minimisers", i.e. length minimising abnormal extremals. R. Montgomery was the first to construct an explicit example of such abnormal curves (see [45]). Since then, many other examples were found, for instance by W. Liu and H.J. Sussmann in [53]. We will deal with this topic in Section 3, where we use the necessary and sufficient conditions for abnormal extremals from Chapter III. In the following section we explain what is meant by a sub-Riemannian structure on a manifold Q, and we define an associated autonomous optimal geometric control structure

1 General definitions

In this section, we first give a brief review of some natural objects associated to a sub-Riemannian structure in order to reformulate the necessary conditions, derived from the maximum principle, for length minimising curves.

Suppose that Q is a smooth manifold of dimension n, equipped with a regular distribution $D \subset TQ$ (i.e. D is a smooth distribution of constant rank, say of rank k). In view of the regularity, D can alternatively be regarded as a vector subbundle of TQ over D. The natural injection $i : D \hookrightarrow TQ$ is then a linear bundle mapping fibred over the identity. A regular distribution is also completely characterised by its annihilator, i.e. giving D is equivalent to specifying the subbundle D^0 of the cotangent bundle T^*Q whose fibre over $q \in Q$ consists of all co-vectors at q which annihilate all vectors in the subspace D_q of T_qQ. Let U be the domain of a coordinate chart in Q. We will always denote coordinates on U by q^i, $i = 1, \ldots, n$. The coordinates on the corresponding bundle chart of T^*Q are denoted by (q^i, p_i), $i = 1, \ldots, n$.

A smooth Riemannian bundle metric h on D is a smooth section of the tensor bundle $D^* \otimes D^* \to Q$ such that it is symmetric and positive definite, i.e. for all $q \in Q$ and $X_q, Y_q \in D_q$ one has that $h(q)(X_q, Y_q) \in \mathbb{R}$ and

$$\begin{aligned} h(q)(X_q, Y_q) &= h(q)(Y_q, X_q), \\ h(q)(X_q, X_q) &\geq 0, \text{ where the equality holds iff } X_q = 0. \end{aligned}$$

Moreover, h depends in a smooth way on q. With a Riemannian bundle metric on D one can associate a smooth linear bundle isomorphism $\flat_h : D \to D^*$, $X_q \mapsto h(q)(X_q, .)$, fibred over the identity on Q, with inverse denoted by $\sharp_h := \flat_h^{-1} : D^* \to D$.

Definition 1.1. A *sub-Riemannian structure* (Q, D, h) is a triple where Q is a smooth manifold, D a smooth regular distribution on Q, and h a Riemannian bundle metric on D.

Although it is not explicitly mentioned in the definition, it will always be tacitly assumed, as is customary in sub-Riemannian geometry, that D is a *non-integrable distribution* and, therefore, does not induce (directly) a foliation of Q. It should be mentioned that, in the literature one often assumes that Q is connected and that D is bracket generating (i.e. the space spanned by all iterated Lie brackets of vector fields in D equals TQ), which is a sufficient condition for saying that the leaf L_q through $q \in Q$ of the foliation induced by the smallest integrable $\widetilde{D}$ generated by D equals to be Q, i.e. $L_q = Q$ (see also Chapter I page 4). In the remainder of this chapter we always assume that Q is connected and that $\widetilde{D} = TQ$, which is a necessary and sufficient condition for the property that any two points in Q can be connected by a piecewise smooth curve tangent to D, which in turn gives sense to the notion of length minimising curves between two points in a sub-Riemannian structure (see Definition 1.2 later on). Note that, in view of the theory developed in Chapter I, this is equivalent to saying that we are working on the pull-back anchored bundle of the anchored bundle $(\tau_Q|_D, i)$. A manifold Q equipped with a sub-Riemannian structure, will be called a *sub-Riemannian manifold*.

With a sub-Riemannian structure (Q, D, h) one can associate a smooth mapping $g : T^*Q \to TQ$ defined by

$$g(\alpha_q) = i\left(\sharp_h\left(i^*(\alpha_q)\right)\right) \in TQ,$$

where $i^* : T^*Q \to D^*$ is the adjoint mapping of i, i.e. for any $\alpha_q \in T_q^*Q$, $i^*(\alpha_q)$ is determined by $\langle i^*(\alpha_q), X_q\rangle = \langle \alpha_q, i(X_q)\rangle$, for all $X_q \in D_q$. Clearly, g is a linear bundle mapping whose image set is precisely the sub-bundle D of TQ and whose kernel is the annihilator D^0 of D. To simplify notations we shall often identify an arbitrary vector in D with its image in TQ under i and smooth sections of D (i.e. elements of $\Gamma(D)$) will often be regarded as vector fields on Q. It is clearly seen that the pair (π_Q, g) is a linear anchored bundle, where $\pi_Q : T^*Q \to Q$ denotes the natural projection.

With g we can further associate a section $\overline{g}$ of $TQ \otimes TQ \to Q$ according to

$$\overline{g}(q)(\alpha_q, \beta_q) = \langle g(\alpha_q), \beta_q\rangle$$

for all $q \in Q$ and $\alpha_q, \beta_q \in T_q^*Q$. From

$$\begin{aligned}\overline{g}(q)(\alpha_q, \beta_q) := \big\langle g(\alpha_q), \beta_q\big\rangle &= \big\langle \sharp_h(i^*\alpha_q), i^*\beta_q\big\rangle \\ &= h(q)\big(\sharp_h(i^*\alpha_q), \sharp_h(i^*\beta_q)\big) \\ &= h(q)\big(g(\alpha_q), g(\beta_q)\big),\end{aligned}$$

we conclude that $\overline{g}$ is symmetric.

Let G be a Riemannian metric on Q. It is easily seen that, given a regular distribution D on Q, the metric G induces a sub-Riemannian structure (Q, D, h_G) where h_G denotes the restriction of G to the subbundle D, i.e. $h_G(q)(X_q, Y_q) := G(q)(X_q, Y_q)$ for any $q \in Q$ and $X_q, Y_q \in D_q$. Given a sub-Riemannian structure (Q, D, h) and a Riemannian metric G on Q, we say that the Riemannian metric restricts to h if $h_G \equiv h$. Now, every sub-Riemannian structure can be seen as being determined (in a non-unique way) by the restriction of a Riemannian metric. Indeed, let h be a Riemannian bundle metric on a vector sub-bundle D of TQ, and let $\{U_\alpha\}$ be an open covering of Q such that, on each U_α, there exists an orthogonal basis $\{X_1, \ldots, X_k\}$ of local sections of D with respect to h. Extend this to a basis of vector fields $\{X_1, \ldots, X_n\}$ on U_α and define a Riemannian metric on U_α by

$$G_\alpha(q)(X_q, Y_q) = \sum_{i,j=1}^{k} a^i b^j h(q)(X_i(q), X_j(q)) + \sum_{i=k+1}^{n} a^i b^i,$$

where $X_q = a^i X_i(q)$ and $Y_q = b^i X_i(q)$, with $a^i, b^i \in \mathbb{R}$. One can then glue these metrics together, using a partition of unity subordinate to the given covering $\{U_\alpha\}$. This procedure, which is similar to the one adopted for constructing a Riemannian metric on an arbitrary smooth manifold (see for instance [3], Proposition 9.4.1), produces a Riemannian metric on Q which, by construction, restricts to h.

In the sequel we will repeatedly make use of a Riemannian metric G which restricts to a given sub-Riemannian metric h. For that purpose we now introduce some further notations and prove some useful relations between such a G and h. However, it should be stressed that the results on sub-Riemannian structures we obtain in the forthcoming sections do not depend on the choice of the specific Riemannian metric G. The natural bundle isomorphism between TQ and T^*Q induced by G will be denoted by $\sharp_G$, with inverse $\flat_G = \sharp_G^{-1}$. Let $q \in Q$ and $X_q, Y_q \in D_q$, then one has:

$$\langle i^* \flat_G(i(X_q)), Y_q \rangle = \langle \flat_G(i(X_q)), i(Y_q) \rangle = \langle \flat_h(X_q), Y_q \rangle \ ,$$

which implies that $\flat_h = i^* \circ \flat_G \circ i$. Inserting this into $g \circ \flat_G \circ i$ and taking into account the definition of g, we conclude that

$$g \circ \flat_G \circ i = i \text{ or } g \circ \flat_G|_D = \mathrm{id}_D \ ,$$

where id_D is the identity mapping on D. The projections of $TQ = D \oplus D^\perp$ onto D and onto its G-orthogonal complement $D^\perp$ will be denoted by π and $\pi^\perp$, respectively. Likewise, T^*Q can be written as the direct sum of $(D^\perp)^0$ and D^0 and the corresponding projections will be denoted by τ and $\tau^\perp$, respectively. It is easily proven that $(D^\perp)^0 \cong \flat_G(D)$ and that

$$\tau^\perp = \flat_G \circ \pi^\perp \circ \sharp_G \,, \quad \tau = \flat_G \circ \pi \circ \sharp_G \,.$$

Using the fact that $g \circ \flat_G|_D = \mathrm{id}_D$ and $\ker g = D^0$, we also have: $g = g \circ \tau = \pi \circ \sharp_G$.

To any regular distribution D on Q one can associate a natural tensor field acting on $D^0 \otimes D \otimes D$. Indeed, let $\eta \in \Gamma(D^0), X, Y \in \Gamma(D)$ and let $[X,Y]$ denote the Lie bracket of X and Y, regarded as vector fields on Q. Then it is easily proven that the expression $\langle \eta, [X,Y] \rangle$ is $C^\infty(Q)$-linear in all three arguments and, therefore, determines a tensorial object. Now, D is involutive if and only if this tensor is identically zero.

Since $\mathrm{im}\, g = D$, we consider the map g as a bundle map from T^*Q to D, which is a linear anchored bundle morphism between (π_Q, g) and $(\tau_Q|D, i)$, where $\tau_Q|_D : D \to Q$ denotes the restriction of τ_Q to the subbundle D. Therefore, any g-admissible curve is mapped by g onto a i-admissible curve, i.e. the base curve is tangent to D. The following lemma proves the converse. To avoid confusion, we shall always denote a g-admissible curve in T^*Q by a Greek letter, e.g. $\alpha(t), \beta(t)$ and a curve in Q is denoted by a Latin letter, e.g. $c(t)$.

Lemma 1.1. *Given a sub-Riemannian structure (Q, D, h) and any piecewise smooth curve c in Q, tangent to D. Then, there always exists a g-admissible curve in T^*Q which projects onto c.*

Proof. Take an arbitrary Riemannian metric G which restricts to h on Q. If $c : [a,b] \to Q$ is a smooth curve tangent to D, one can simply put $\alpha(t) = \flat_G(\dot{c}(t))$ for all $t \in [a,b]$. Clearly, α then defines a g-admissible curve in T^*Q with base curve c. Next, assume $c : [a,b] \to Q$ is a piecewise smooth curve, tangent to D, i.e. c is the concatenation of ℓ smooth curves $c^i : [a_{i-1}, a_i] \to Q$, such that $c = c_\ell \cdot \ldots \cdot c_1$. We can then define ℓ curves α^i in T^*Q as follows: put $\alpha^i(t) = \flat_G(\dot{c}_i(t))$ for $t \in [a_{i-1}, a_i]$. It is easily seen that α^i is g-admissible, with base curve c_i and, hence, $\alpha^\ell \cdot \ldots \cdot \alpha^1$ is also g-admissible with projection c. □

We will now introduce the notion of length of piecewise smooth curves, tangent to D (i.e. base curves of i-admissible curves).

Definition 1.2. Given a sub-Riemannian structure (Q, D, h), then the *length* $l(c)$ of a curve $c : [a, b] \to M$, tangent to D, is given by

$$l(c) := \int_a^b \sqrt{h(c(t))(\dot{c}(t), \dot{c}(t))}dt.$$

Given any g-admissible curve α in T^*Q with base curve c, and a Riemannian metric G which restricts to h, then the length of c still equals

$$l(c) = \int_a^b \sqrt{\overline{g}(c(t))(\alpha(t), \alpha(t))}dt = \int_a^b \sqrt{G(c(t))(\dot{c}(t), \dot{c}(t))}dt.$$

These integrals do not depend on the choice of α, resp. G. Using standard arguments, it is easily seen that the length of c does not depend on a specific parametrisation. Therefore, we may always assume that c is defined on $[0, 1]$ (see page 11). We say that c, taking q_0 to q_1, is *length minimising* if, given any other base curve of an i-admissible curve $c' : [0, 1] \to Q$, with $c'(0) = q_0$ and $c'(1) = q_1$, then $l(c) \leq l(c')$. From the fact that i-admissible curves are precisely curves tangent to D, it is easily seen that the functional l is a cost functional in the sense of Chapter III, associated with the cost $L(v) = \sqrt{h(v, v)}$, for $v \in D$. To avoid confusion, we shall denote the cost functional associated with L by $\mathcal{J}^L$, i.e. for c a curve tangent to D, we have

$$\mathcal{J}^L(\dot{c}) = \int_0^1 L(\dot{c}(t))dt.$$

Using these notations, we have $\mathcal{J}^L(\dot{c}) = l(c)$. Consider the energy cost function $E \in C^\infty(D)$, defined by $E(v) = \frac{1}{2}h(v, v)$. It should be noted that whereas L is not a regular cost, E is regular. Indeed, the fibre derivative of E is invertible: $\mathbb{F}E = \flat_h : D \to D^*$ with inverse $\sharp_h$. We now prove that, if the base curve c of the i-admissible curve $\dot{c}$ is length minimising, then $\dot{c}$ is optimal with respect to the cost E, and vice versa.[1] We shall make use of Hölder's inequality (see [20, p 55]), which says that,

$$l(c) = \mathcal{J}^L(\dot{c}) \leq \sqrt{2\mathcal{J}^E(\dot{c})}, \tag{1.1}$$

when adapted to this particular situation. The above inequality reduces to an equality if $E(\dot{c}(t))$ is constant on every smooth part of c. If this is

[1] The proof of this correspondence is taken from W. Respondek.

the case we say that c is *parameterised by constant arc length*. The above inequality implies that, given a length minimising curve c, parameterised by constant arc length, then $\dot{c}$ is optimal with respect to the energy E. On the other hand, assume that $\dot{c}$ is energy optimal and that there exists a curve c' such that $l(c') \leq l(c)$. We may assume that c' is parameterised by constant arc length since this doesn't affect the value of $l(c')$. Then the inequality $\sqrt{2\mathcal{J}^E(\dot{c}')} = l(c') \leq l(c) \leq \sqrt{2\mathcal{J}^E(\dot{c})}$ holds, leading to a contradiction. Moreover, it also follows that c is necessarily parameterised by constant arc length. We conclude that the length minimising problem and the optimal control problem with respect to the cost E are equivalent. Therefore, we study 'energy optimal' i-admissible curves in the linear autonomous geometric control structure $(\tau_Q|_D, i)$. In view of these remarks we shall consider in the following abnormal and normal (local) extremals without always mentioning explicitly that they are defined with respect to the anchored bundle $(\tau_Q|_D, i)$ with cost function E (cf. Section III-8).

Since E is a regular cost, we can now apply Theorem III-8.6 (page 117), saying that every normal local extremal $\dot{c}(t)$ is generated by an integral curve $\alpha(t)$ of the Hamiltonian vector field on T^*Q determined by the function H on T^*Q:

$$\begin{aligned}\alpha_q \mapsto h_\lambda\Big((\mathbb{F}E)^{-1}\big(i^*(\alpha_q)\big), \alpha_q\Big) = \\ \big\langle \sharp_h\big(i^*(\alpha_q)\big), \alpha_q\big\rangle - \tfrac{1}{2} h\big(\sharp_h\big(i^*(\alpha_q)\big), \sharp_h\big(i^*(\alpha_q)\big)\big) = \tfrac{1}{2}\overline{g}(\alpha_q, \alpha_q),\end{aligned}$$

where $\lambda = -1$ and for any $\alpha_q \in T_q^*Q$. The fact that $\alpha(t)$ generates $\dot{c}(t)$ is expressed by $\sharp_h\big(i^*(\alpha(t))\big) = g(\alpha(t)) = \dot{c}(t)$, i.e. $\alpha(t)$ is g-admissible.

Let us now consider the equations for an abnormal extremal $\dot{c}(t)$. Let α denote an arbitrary g-admissible curve with base curve c. If we apply the condition for $\dot{c}$ to be an abnormal extremal (cf. Theorem III-8.1), then there exists a piecewise smooth section $\theta(t)$ of D^0 such that the following (local) equations are satisfied:

$$\dot{p}_i(t) = -\frac{\partial g^{jk}}{\partial q^i}(q(t))\alpha_j(t)p_k(t),$$

where $\alpha(t) = (q^i(t), \alpha_i(t))$, $\theta(t) = (q^i(t), p_i(t))$. Note that these equations are derived from the implicit Hamiltonian system in Theorem III-8.1 where we used the fact that $\theta(t) \in D^0$ and that $\dot{c}(t) = g(\alpha(t))$. Let ∇ denote an arbitrary g-connection on T^*M adapted to D, i.e. $\nabla_\beta\eta = i_\beta d\eta = \delta_{g(\beta)}\eta$, for $\eta \in \Gamma(D^0)$ (cf. Section II-7). Then, if we express $\nabla_{dq^i}\eta$ in coordinates, we obtain:

$$
\begin{array}{rcl}
\nabla_{dq^i}\eta & = & \left(g^{ik}\dfrac{\partial \eta_j}{\partial q^k} - \dfrac{\partial \eta_k}{\partial q^j} g^{ik}\right) dq^j \\
\left(g^{ik}\dfrac{\partial \eta_j}{\partial q^k} + \Gamma^{ik}_j \eta_k\right) dq^j & = & \left(g^{ik}\dfrac{\partial \eta_j}{\partial q^k} + \dfrac{\partial g^{ik}}{\partial q^j}\eta_k\right) dq^j \\
& \Updownarrow & \\
\Gamma^{ik}_j \eta_k & = & \dfrac{\partial g^{ik}}{\partial q^j}\eta_k
\end{array}
$$

From these local expressions, we conclude that the i-admissible curve $\dot{c}(t) = g(\alpha(t))$ is an abnormal extremal iff there exists a piecewise smooth section θ of D^0 such that $\nabla_\alpha \theta(t) = 0$, given any g-connection adapted to the bundle D. The following corollary summarises the above obtained results on normal and abnormal local extremals in sub-Riemannian geometry.

Corollary 1.2. *Let $c : [0,1] \to Q$ denote a piecewise smooth curve, tangent to D. If c is length minimising and parameterised by constant arc length, then, at least one of the following two conditions is satisfied:*

1. *there exists an integral curve $\alpha(t)$ of X_H, which projects onto c and for which $g(\alpha(t)) = \dot{c}(t)$;*
2. *there exists a section $\theta(t)$ of D^0 along c, such that $\nabla_\alpha \theta(t) = 0$, where α is g-admissible with base curve c and where ∇ is a g-connection adapted to the bundle D.*

As the above corollary suggests, it will be natural to study g-connections and their relation to extremal curves. This is the topic of the following section. Strictly speaking an extremal $\dot{c}$ refers to a curve in $D \subset TQ$, where c is a piecewise smooth curve in Q. But, for our convenience, we will call here, in the sub-Riemannian context, the curve c itself an extremal.

2 Connections in sub-Riemannian geometry

Fix a sub-Riemannian structure (Q, D, h) and consider the associated bundle map $g : T^*Q \to TQ$. In sub-Riemannian geometry we only consider generalised connections on T^*Q over g. Our main goal is the characterisation of normal and abnormal extremals of the sub-Riemannian structure in terms of such generalised connections. We first define a notion of "geodesic" for g-connections.

Definition 2.1. A g-admissible curve $\alpha : I \to T^*Q$ is said to be an *autoparallel curve* with respect to a g-connection ∇ if it satisfies $\nabla_\alpha \alpha(t) = 0$ for all $t \in I$. Its base curve $c = \pi_Q \circ \alpha$ is then called a *geodesic* of ∇.

In canonical coordinates on T^*Q, an autoparallel curve $\alpha(t) = (q^i(t), p_i(t))$ satisfies the equations

$$\dot{q}^i(t) = g^{ij}(q(t))p_j(t) \ , \quad \dot{p}_j(t) = -\Gamma_j^{ik}(q(t))p_i(t)p_k(t),$$

where g^{ij} and $\Gamma_j^{ik} \in C^\infty(U)$ are the local components of the contravariant tensor field $\overline{g}$, associated to the sub-Riemannian structure, and the connection coefficients of ∇, respectively. In fact, given a g-connection ∇ one can always define a smooth vector field Γ^∇ on T^*Q whose integral curves are autoparallel curves with respect to ∇. In canonical coordinates, this vector field (also called the *geodesic spray* of ∇) reads:

$$\Gamma^\nabla(q,p) = g^{ij}(q)p_j \frac{\partial}{\partial q^i} - \Gamma_j^{ik}(q)p_i p_k \frac{\partial}{\partial p_j}.$$

(A proof of this property follows by standard arguments, and is left to the reader.) This implies, in particular, that given any $\alpha_0 \in T^*Q$, there exists an autoparallel curve $\alpha(t)$ passing through α_0. Note that it may happen that two different autoparallel curves correspond to the same base curve (i.e. may project onto the same geodesic). We can also define a notion of *Killing vector field* for a g-connection ∇.

Definition 2.2. A vector field $\xi \in \mathfrak{X}(Q)$ is called a *Killing vector field* for a given g-connection ∇ if, given any autoparallel curve $\alpha(t)$ with base curve $c(t)$, $\langle \alpha(t), \xi(c(t)) \rangle$ is constant for all t.

Using standard techniques, it is easily seen that ξ is Killing iff $\langle \alpha, \nabla_\alpha \xi \rangle = 0$ for any one-form $\alpha \in \mathfrak{X}(Q)$. Note that, if we define a *Killing one-form* $\kappa \in \mathfrak{X}^*(Q)$ by the requirement that $g(\kappa(c(t)), \alpha(t))$ is constant for any autoparallel curve α with base curve c, then the vector field $g \circ \kappa$ is a Killing vector field. Therefore, we shall only consider Killing vector fields, which turns out to be a more general concept then Killing one-forms.

Now, we would like to find a g-connection on a sub-Riemannian manifold whose geodesics are precisely the normal extremals. Recalling the characterisation of a normal extremal in Corollary 1.2, it follows that we have to look for a g-connection ∇ for which $\Gamma^\nabla = X_H$, where X_H denotes the Hamiltonian vector field corresponding to $H(\alpha_q) = \frac{1}{2}\overline{g}(\alpha_q, \alpha_q) \in C^\infty(T^*Q)$. A first step in that direction is the construction of a symmetric product associated with a given g-connection, which fully characterises the geodesics of the g-connection under consideration.

Two g-connections ∇ and $\overline{\nabla}$ have the same autoparallel curves if and only if the tensor field $T : \mathfrak{X}^*(Q) \otimes \mathfrak{X}^*(Q) \to \mathfrak{X}^*(Q), (\alpha, \beta) \mapsto \nabla_\alpha \beta - \overline{\nabla}_\alpha \beta$ is skew-symmetric or, equivalently, $T(\alpha, \alpha) \equiv 0$ for all $\alpha \in \mathfrak{X}^*(Q)$. In local coordinates, the components of T are given by $T^{ij}_k = \Gamma^{ij}_k - \overline{\Gamma}^{ij}_k$, where Γ^{ij}_k and $\overline{\Gamma}^{ij}_k$ are the connection coefficients of ∇ and $\overline{\nabla}$, respectively. We immediately see that T is skew-symmetric iff $\Gamma^\nabla = \Gamma^{\overline{\nabla}}$, proving the previous statement. Define the *symmetric product* of a connection ∇ as

$$\langle \alpha : \beta \rangle_\nabla := \nabla_\alpha \beta + \nabla_\beta \alpha \ , \quad \text{for } \alpha, \beta \in \mathfrak{X}^*(Q) \ .$$

(Observe that this is not a tensorial quantity, i.e. $\langle \alpha : \beta \rangle_\nabla$ is not $C^\infty(Q)$-linear in its arguments). By replacing α by $\alpha + \beta$ in $T(\alpha, \alpha)$ the following lemma is easily proven.

Lemma 2.1. *The geodesics of a g-connection ∇ are completely determined by the symmetric product $\langle \alpha : \beta \rangle_\nabla$ in the sense that, given two g-connections ∇ and $\overline{\nabla}$, then both have the same geodesics if and only if $\langle \alpha : \beta \rangle_\nabla = \langle \alpha : \beta \rangle_{\overline{\nabla}}$, for all $\alpha, \beta \in \mathfrak{X}^*(Q)$.*

In the following we shall construct a symmetric bracket of one-forms, associated to a sub-Riemannian structure (Q, D, h), which coincides with the symmetric product of a g-connection ∇ on T^*Q iff $\Gamma^\nabla = X_H$.

Before proceeding, we first recall that the Levi-Civita connection ∇^G associated to an arbitrary Riemannian metric G is completely determined by the relation:

$$\begin{aligned} 2G(\nabla^G_X Y, Z) &= X(G(Y,Z)) + Y(G(X,Z)) - Z(G(X,Y)) \\ &\quad + G([X,Y],Z) - G([X,Z],Y) - G(X,[Y,Z]) \ , \end{aligned}$$

for all $X, Y, Z \in \mathfrak{X}(Q)$ (see [3]). This can still be rewritten as

$$2\flat_G(\nabla^G_X Y) = \mathcal{L}_X \flat_G(Y) + \mathcal{L}_Y \flat_G(X) + \flat_G([X,Y]) - d(G(X,Y)),$$

and the symmetric product of two vector fields X, Y, defined by $\langle X : Y \rangle_{\nabla^G} = \nabla^G_X Y + \nabla^G_Y X$, then satisfies

$$\flat_G(\langle X : Y \rangle_{\nabla^G}) = \mathcal{L}_X \flat_G(Y) + \mathcal{L}_Y \flat_G(X) - d(G(X,Y)).$$

The right-hand side of this equation now inspires us to propose the following definition for a symmetric bracket of one-forms on a sub-Riemannian manifold.

Definition 2.3. The *symmetric bracket* associated to a sub-Riemannian structure (Q, D, h) is the mapping $\{\cdot,\cdot\} : \mathfrak{X}^*(Q) \times \mathfrak{X}^*(Q) \to \mathfrak{X}^*(Q)$ defined by:

$$\{\alpha, \beta\} = \mathcal{L}_{g(\alpha)}\beta + \mathcal{L}_{g(\beta)}\alpha - d\left(\overline{g}(\alpha, \beta)\right).$$

In the following proposition we list some properties of this bracket, the first of which justifies the denomination "symmetric bracket". The proofs of these properties are straightforward and immediately follow from the above definition.

Proposition 2.2. *The symmetric bracket satisfies the following properties: for any $\alpha, \beta \in \mathfrak{X}^*(Q)$*

1. $\{\alpha, \beta\} = \{\beta, \alpha\}$*;*
2. *the bracket is* $\mathbb{R}$*-bilinear;*
3. $\{f\alpha, \beta\} = g(\beta)(f)\alpha + f\{\alpha, \beta\}$*, with* $f \in C^\infty(Q)$*,*
4. $\{\alpha, \eta\} = \delta_{g(\alpha)}\eta$*, for any* $\eta \in \Gamma(D^0)$*, and* $\{\alpha, \eta\} = 0$ *if both* α *and* η *belong to* $\Gamma(D^0)$*.*

The first three properties justify the following definition.

Definition 2.4. A g-connection ∇ is said to be *normal* if the associated symmetric product equals the symmetric bracket, i.e. if $\langle \alpha : \beta \rangle_\nabla = \{\alpha, \beta\}$ holds for all $\alpha, \beta \in \mathfrak{X}^*(Q)$.

The connection coefficients of a normal g-connection satisfy the relations

$$\Gamma^{ij}_k + \Gamma^{ji}_k = \frac{\partial g^{ij}}{\partial q^k}, \text{ for all } i, j, k = 1, \ldots, n.$$

Theorem 2.3. *Let ∇ be a g-connection, then the following statements are equivalent:*

1. ∇ *is a normal g-connection;*
2. *for all* $\alpha \in \mathfrak{X}^*(Q) : \nabla_\alpha \alpha = \frac{1}{2}\{\alpha, \alpha\}$*;*
3. $\langle \nabla_\alpha X, \beta \rangle + \langle \nabla_\beta X, \alpha \rangle = \langle [g(\alpha), X], \beta \rangle + \langle [g(\beta), X], \alpha \rangle + X(\overline{g}(\alpha, \beta))$ *for all* $\alpha, \beta \in \mathfrak{X}^*(Q)$ *and* $X \in \mathfrak{X}(Q)$*;*

4. *$\Gamma^\nabla = X_H$ or, equivalently, every autoparallel curve of ∇ is a normal multiplier and vice versa;*

5. *let G be a Riemannian metric restricting to h and let ∇^G be the corresponding Levi-Civita connection, then for all $\alpha \in \mathfrak{X}^*(Q)$, ∇ satisfies:*

$$\nabla_\alpha \alpha = \nabla^G_{g(\alpha)} \tau(\alpha) + \delta_{g(\alpha)} \tau^\perp(\alpha).$$

(Note that the right-hand side of the third property above agrees with the definition of the symmetrised covariant derivative considered in [50].)

Proof. The equivalence of (1) and (2) follows directly from the definition of a normal g-connection, and the equivalence of (1) and (3) follows from $\langle \nabla_\alpha \beta, X\rangle = g(\alpha)(\langle \beta, X\rangle) - \langle \beta, \nabla_\alpha X\rangle$ after some tedious but straightforward calculations.

(2) $\Leftrightarrow$ (4). Choose an arbitrary $\alpha_0 \in T^*Q$. Let U be a coordinate neighbourhood of $q_0 = \pi_Q(\alpha_0)$ and put $\alpha_0 = (q^i_0, p^0_j)$. Then, $\nabla_\alpha \alpha = \frac{1}{2}\{\alpha, \alpha\}$ implies, in particular, that the connection coefficients Γ^{ij}_k of ∇ on U satisfy

$$\Gamma^{ij}_k(q_0) p^0_i p^0_j = \frac{1}{2} \frac{\partial g^{ij}}{\partial q^k}(q_0) p^0_i p^0_j.$$

The coordinate expression for the Hamiltonian vector field X_H at α_0 reads:

$$X_H(\alpha_0) = g^{ij}(q_0) p^0_j \left.\frac{\partial}{\partial q^i}\right|_{\alpha_0} - \frac{1}{2} \frac{\partial g^{ij}}{\partial q^k} p^0_i p^0_j \left.\frac{\partial}{\partial p_k}\right|_{\alpha_0}.$$

Recalling the definition of Γ^∇ it is easy to see that $\Gamma^\nabla(\alpha_0) = X_H(\alpha_0)$ for any $\alpha_0 \in T^*Q$ if and only if $\nabla_\alpha \alpha = \frac{1}{2}\{\alpha, \alpha\}$ for each $\alpha \in \mathfrak{X}^*(Q)$.

(2) $\Leftrightarrow$ (5). Let G be a Riemannian metric restricting to h. Recall the following property of the Levi-Civita connection ∇^G:

$$\flat_G(\langle X : Y\rangle_{\nabla^G}) = \mathcal{L}_X \flat_G(Y) + \mathcal{L}_Y \flat_G(X) - d(G(X,Y)).$$

Putting $X = Y = g(\alpha)$, this equation becomes

$$\flat_G(\nabla^G_{g(\alpha)} g(\alpha)) = \mathcal{L}_{g(\alpha)} \flat_G(g(\alpha)) - \tfrac{1}{2} d(\overline{g}(\alpha, \alpha)).$$

Using the identity $\flat_G(g(\alpha)) = \tau(\alpha)$ derived in Section 1 (page 131) of this chapter, and taking into account that ∇^G preserves the metric G, i.e. $\nabla^G \circ \flat_G = \flat_G \circ \nabla^G$, we obtain

$$\begin{aligned} \nabla^G_{g(\alpha)} \tau(\alpha) &= \mathcal{L}_{g(\alpha)} \tau(\alpha) - \tfrac{1}{2} d(\overline{g}(\alpha, \alpha)), \\ &= \tfrac{1}{2}\{\alpha, \alpha\} - \mathcal{L}_{g(\alpha)} \tau^\perp(\alpha). \end{aligned}$$

Since $\tau^\perp(\alpha) \in \Gamma(D^0)$ and $g(\alpha) \in \Gamma(D)$, the last term on the right-hand side reduces to $\delta_{g(\alpha)}\tau^\perp(\alpha)$, which completes the proof. □

Theorem 2.3 implies, in particular, that normal g-connections exist. For instance, the operator ∇ defined by

$$\nabla_\alpha \beta = \nabla^G_{g(\alpha)}\tau(\beta) + \delta_{g(\alpha)}\tau^\perp(\beta)$$

is a g-connection and it is normal in view of the equivalence of (1) and (5). Moreover, for $\beta \in \Gamma(D^0)$ we find that $\nabla_\alpha\beta = \delta_{g(\alpha)}\beta$, i.e. the connection under consideration is also D-adapted. Summarising, we have shown the following result.

Proposition 2.4. *Given a sub-Riemannian structure (Q, D, h), one can always construct a g-connection which is both normal and D-adapted.*

Furthermore, Theorem 2.3 (in particular, the equivalence between (1) and (5)) provides us with a relation between a normal g-connection, the Levi-Civita connection ∇^G of any Riemannian metric restricting to h, and the operator δ. This relation will be very useful when we study the relation between vakonomic dynamics and nonholonomic mechanics (see Section 4). The following property gives a characterisation of Killing vector fields in sub-Riemannian geometry:

Proposition 2.5. *Let ∇ denote a normal g-connection. A vector field $\xi \in \mathfrak{X}(Q)$ is a Killing vector field for ∇ iff $\mathcal{L}_\xi \overline{g} = 0$.*

Proof. Let ξ denote an arbitrary vector field, and consider the following equalities, for arbitrary $\alpha \in \mathfrak{X}^*(Q)$:

$$\begin{aligned} -\tfrac{1}{2}\mathcal{L}_\xi\overline{g}(\alpha,\alpha) &= -\tfrac{1}{2}\xi\big(\overline{g}(\alpha,\alpha)\big) + \langle \mathcal{L}_\xi\alpha, g(\alpha)\rangle \\ &= \tfrac{1}{2}\xi\big(\overline{g}(\alpha,\alpha)\big) - \langle \alpha, [\xi, g(\alpha)]\rangle \\ &= \tfrac{1}{2}\xi\big(\overline{g}(\alpha,\alpha)\big) + \langle \alpha, [g(\alpha), \xi]\rangle. \end{aligned}$$

Now, since ∇ is normal, i.e. it satisfies (3) from Theorem 2.3, it follows that the right-hand side of the last equality equals $\langle \alpha, \nabla_\alpha \xi\rangle$. By definition, this vanishes if ξ is a Killing vector field and, therefore, $\mathcal{L}_\xi\overline{g} = 0$. The proof in the other direction simply follows by reversing the previous arguments. □

To conclude this section we make some further remarks on normal and D-adapted g-connections. It is well known that the Levi-Civita connection

∇^G, associated with a Riemannian metric G, is uniquely determined by the properties that it preserves the metric, i.e. $\nabla^G G = 0$, and that its torsion is zero. We would like to consider g-connections ∇ on a sub-Riemannian manifold which are *metric*, i.e. $\nabla\overline{g} = 0$ (where $\overline{g}$ is the symmetric contravariant 2-tensor field defined in Section 1). From above we know that normal extremals of a sub-Riemannian structure, resp. abnormal extremals, can be characterised as geodesics of a normal g-connection, resp. as parallel transported sections of D^0 for a D-adapted g-connection (see Theorem 2.3, resp. Corollary 1.2). Therefore it is natural to look for g-connections that are simultaneously normal and D-adapted. It has been shown above that such a g-connection always exists, namely

$$\nabla_\alpha\beta = \nabla^G_{g(\alpha)}\tau(\beta) + \delta_{g(\alpha)}\tau^\perp(\beta),$$

with G any Riemannian metric restricting to h. We will prove, however, that no metric g-connection can be found that is also D-adapted, provided that D is integrable. First we prove an interesting result relating the notion of partial g-connection (cf. Section II-7) with that of a D-adapted normal g-connection.

Proposition 2.6. *Let ∇ be a normal g-connection. Then ∇ is partial if and only if ∇ is D-adapted.*

Proof. Let ∇ be a normal g-connection, i.e. $\nabla_\alpha\beta + \nabla_\beta\alpha = \{\alpha,\beta\}$, for all $\alpha,\beta \in \mathfrak{X}^*(Q)$. Suppose ∇ is partial, then for $\beta \in \Gamma(D^0)$ the previous relation becomes:

$$\nabla_\alpha\beta = \{\alpha,\beta\} = \mathcal{L}_{g(\alpha)}\beta = \delta_{g(\alpha)}\beta,$$

i.e. ∇ is D-adapted. Conversely, suppose ∇ is normal and D-adapted, then $\nabla_\alpha\beta = \{\alpha,\beta\} - \nabla_\beta\alpha$. Let $\alpha \in \Gamma(D^0)$, then the right hand side of this equation is zero, and thus $\nabla_\alpha\beta = 0$ for all $\alpha \in \Gamma(D^0)$ and $\beta \in \mathfrak{X}^*(Q)$. This proves the proposition. □

We will now describe a general method for constructing normal g-connections.

Let $[\cdot\,,\cdot] : \mathfrak{X}^*(Q) \times \mathfrak{X}^*(Q) \to \mathfrak{X}^*(Q)$ denote a skew-symmetric bracket that is $\mathbb{R}$-linear in both arguments and satisfies, for any $f \in \mathcal{F}(Q)$, $[\alpha, f\beta] = g(\alpha)(f)\beta + f[\alpha,\beta]$. Given such a bracket on $\mathfrak{X}^*(Q)$, one can define a unique normal g-connection ∇ for which $[\alpha,\beta] = \nabla_\alpha\beta - \nabla_\beta\alpha$, namely:

$$\nabla_\alpha\beta = \tfrac{1}{2}\left([\alpha,\beta] + \{\alpha,\beta\}\right).$$

Conversely, given a normal g-connection ∇, one can define a skew-symmetric bracket with the desired properties by putting $[\alpha,\beta] = \nabla_\alpha\beta - \nabla_\beta\alpha$. Henceforth, we shall denote the bracket associated with a normal g-connection ∇ by $[\alpha,\beta]_\nabla$.

As can be easily verified, for a g-connection ∇ which is both normal and D-adapted, the skew-symmetric bracket satisfies: $[\alpha,\eta]_\nabla = \delta_{g(\alpha)}\eta$, for all $\eta \in \Gamma(D^0)$ and $\alpha \in \mathfrak{X}^*(Q)$. Therefore, if a Riemannian metric G is chosen, with projections τ and $\tau^\perp$ on $\flat_G(D)$ and D^0 respectively, and which restricts to h, this bracket takes the form:

$$[\alpha,\beta]_\nabla = [\tau(\alpha),\tau(\beta)]_\nabla + \delta_{g(\alpha)}\tau^\perp(\beta) - \delta_{g(\beta)}\tau^\perp(\alpha).$$

We only have to know the value of the bracket acting on sections of $\flat_G(D) \cong D$. For example, for the g-connection defined by

$$\nabla_\alpha\beta = \nabla^G_{g(\alpha)}\tau(\beta) + \delta_{g(\alpha)}\tau^\perp(\beta),$$

the associated bracket becomes:

$$[\alpha,\beta]_\nabla = \flat_G\left([g(\alpha),g(\beta)]\right) + \delta_{g(\alpha)}\tau^\perp(\beta) - \delta_{g(\beta)}\tau^\perp(\alpha),$$

where $[g(\alpha),g(\beta)] = \mathcal{L}_{g(\alpha)}g(\beta)$ is the usual Lie bracket on vector fields. Note, however, that there does not seem to exist a 'natural' skew-symmetric bracket on $\mathfrak{X}^*(Q)$, independent of the chosen Riemannian extension G of h, which could be used to identify a 'standard' g-connection which is both normal and D-adapted. One might think of imposing a metric condition in order to completely determine such a ∇, but the following result shows that this is impossible. Note that we have assumed that D is not integrable.

Proposition 2.7. *A D-adapted g-connection can not be metric.*

Proof. Let ∇ be D-adapted g-connection. Suppose that ∇ leaves $\overline{g}$ invariant. This can be equivalently expressed by $g(\nabla_\alpha\beta) = \nabla_\alpha(g(\beta))$ for all $\alpha,\beta \in \mathfrak{X}^*(Q)$. Let $\eta \in \Gamma(D^0)$, then, since ∇ is D-adapted this equation becomes $g(\delta_{g(\alpha)}\eta) = 0$ for all $\alpha \in \mathfrak{X}^*(Q)$ and $\eta \in \Gamma(D^0)$. However, this is equivalent to saying that D is involutive. Indeed, from $g(\delta_{g(\alpha)}\eta) = 0$ we have

$$0 = \langle\beta, g(\delta_{g(\alpha)}\eta)\rangle = \langle\delta_{g(\alpha)}\eta, g(\beta)\rangle = -\langle\eta,[g(\alpha),g(\beta)]\rangle,$$

for arbitrary $\alpha,\beta \in \mathfrak{X}^*(Q)$ and $\eta \in \Gamma(D^0)$, hence $[g(\alpha),g(\beta)] \in \Gamma(D)$. □

3 Normal and abnormal extremals revisited

In the first part of this section we will make use of Theorem 2.3 to recover some known results about normal extremals. In the second part we further investigate abnormal extremals. Consider a sub-Riemannian structure (Q, D, h) and let G be an arbitrary Riemannian metric on Q restricting to h. Theorem 2.3 then says $\nabla_\alpha \alpha(t) = \nabla^G_{\dot{c}} \tau(\alpha)(t) + \delta_{\dot{c}} \tau^\perp(\alpha)(t)$, where α is a g-admissible curve with base curve c, and ∇ is any normal g-connection. This immediately leads to the following result.

Proposition 3.1. *Let $c : I \to Q$ be a curve tangent to D which is geodesic with respect to a Riemannian metric G on Q restricting to h, then c is a normal extremal of the sub-Riemannian structure.*

Proof. The curve c is a normal extremal if there exists a g-admissible curve α with base curve c, which is autoparallel with respect to a normal g-connection ∇. Since $c : I \to Q$ is a geodesic with respect to G, i.e. $\nabla^G_{\dot{c}} \dot{c}(t) = 0$ for all $t \in I$, we know from Section 1 that $\alpha = \flat_G(c)$ is a g-admissible curve with base curve c for which $\tau(\alpha) = \alpha$ or $\tau^\perp(\alpha) = 0$. It then follows that $\nabla_\alpha \alpha(t) = 0$ since $\nabla_\alpha \alpha(t) = \nabla^G_{\dot{c}} \tau(\alpha)(t) = \flat_G(\nabla^G_{\dot{c}} \dot{c}(t))$ =0. □

We now try to prove the converse of the above proposition. Let $c : I = [a, b] \to Q$ be a normal extremal in a sub-Riemannian structure (Q, D, h). Then there exists a g-admissible curve α which is autoparallel with respect to a normal g-connection. Given any $t_0 \in I$, then one can always find a one-form $\overline{\alpha}$ and a compact subinterval J of I containing t_0, such that $\overline{\alpha}(c(t)) = \alpha(t)$ for all $t \in J$ and $c(J)$ is contained in a coordinate neighbourhood U. We will now construct a local Riemannian metric G restricting to h on D such that $c|_J$ is a geodesic with respect to this Riemannian metric.

Since $g(\overline{\alpha}) \neq 0$, one can construct a local basis of $\mathfrak{X}^*(U)$, namely $\{\overline{\alpha} = \beta^1, \beta^2, \dots, \beta^n\}$, such that $\beta^{k+1}, \dots, \beta^n$ determine a local basis for $\Gamma(D^0)$, defined on U. Let $\{X_1, \dots, X_n\}$ denote the dual basis. Then the vector fields X_j, for $j = 1, \dots, k$, form a local basis for $\Gamma(D)$, since $\langle \beta^i, X_j \rangle \equiv 0$ for $i = k+1, \dots, n$. We can now define a Riemannian metric G on U, restricting to h, as in Section 1, i.e. for arbitrary vector fields Y and Z on U,

$$G(q)(Y, Z) = \sum_{r,s=1}^{k} Y^r Z^s h(q)(X_r(q), X_s(q)) + \sum_{r=k+1}^{n} Y^r Z^r,$$

where we have put $Y(q) = Y^r X_r(q)$ and $Z(q) = Z^r X_r(q)$ for some $Y^r, Z^r \in \mathbb{R}$ $(r = 1, \ldots, n)$. From the definition of G we can derive that $D^\perp$ is spanned by $\{X_{k+1}, \ldots, X_n\}$ such that, in particular, $\tau^\perp(\overline{\alpha}) = 0$. This implies that $\tau^\perp(\alpha(t)) = 0$ and, hence, $\flat_G(\dot{c}(t)) = \alpha(t)$. From $\nabla_\alpha \alpha(t) = 0$ and $\tau^\perp(\alpha(t)) = 0$ we obtain that $\nabla^G_{\dot{c}} \dot{c}(t) = 0$ holds for any $t \in J$. Herewith, we have shown that the following result holds.

Proposition 3.2. *Let $c : I \to Q$ be a normal extremal. Then for any $t \in I$ there exists a compact neighbourhood $J \subset I$ of t such that c, when restricted to J, is a geodesic with respect to some Riemannian metric restricting to h on D.*

As a corollary this proves that a normal extremal is locally length minimising.

Let c be a normal extremal and let ∇ be a normal and D-adapted g-connection (recall that such a ∇ always exists). Suppose that c is *degenerate* in the following sense: there exist two g-admissible curves α, β with base curve c, such that $\nabla_\alpha \alpha(t) = \nabla_\beta \beta(t) = 0$. We will now see that c is then also an abnormal extremal. We have proven before that a normal and D-adapted connection is partial, i.e. $\nabla_\alpha = \nabla_\beta$ if $g(\alpha) = g(\beta)$. Therefore one obtains that $\nabla_\alpha(\alpha - \beta)(t) = 0$. Since $g(\alpha(t) - \beta(t)) = 0$, or $\theta(t) = (\alpha - \beta)(t) \in D^0$ for all t, θ is a parallel transported section along α, lying entirely in D^0 and, hence, c is an abnormal extremal. Conversely, assume that c is a normal extremal, i.e. c is the base curve of an autoparallel curve α with respect to ∇, and that c is also an abnormal extremal. Let θ denote a parallel transported section along α lying in D^0. Then, using the same arguments as before, $\alpha + \theta$ is also an autoparallel curve with base curve c. We can conclude that curves that are both normal and abnormal are degenerate in the sense that they admit more than one g-admissible curve that is autoparallel.

Let us now reconsider the equivalent characterisation of an abnormal extremal as described in Corollary 1.2. Let $\alpha : [0,1] \to T^*Q$ denote an arbitrary g-admissible curve taking q_0 to q_1, with base curve c. Using the theory of linear autonomous optimal control problems from Section III-8, we have the following correspondence: let $\theta(t)$ denote a section of D^0 along c, then $\nabla_\alpha \theta(t) = 0$ iff $\theta(1) \in V^*_{(1,q_1)} R_{(0,q_0)}$. This observation was one of the key propositions in [28], which is now more general since we also admit curves c with singular points (i.e. curves for which there exists a $t \in [0,1]$ such that $\dot{c}(t) = 0$). The sufficient conditions for c not to be abnormal, obtained in Theorem III-8.5, can be used in sub-Riemannian geometry without modification. In the remainder of this section we consider two examples of a sub-Riemannian structure on $\mathbb{R}^3$ and use the necessary and sufficient conditions for the existence of (strictly) abnormal extremals to prove that for the structures under consideration, strictly abnormal extremals do exist. The first example is the sub-Riemannian structure that was used in [45] to prove the existence of a strictly abnormal minimiser (see also Example I-1.2 page 6). The second example is taken from [38].

Example 3.3. Take $Q = \mathbb{R}^3$ (we use cylindrical coordinates (r, θ, z) on $\mathbb{R}^3$), and $D = \operatorname{span}\{X_1, X_2\}$ with $X_1 = \partial/\partial r$ and $X_2 = \partial/\partial\theta - p(r)\partial/\partial z$, where $p(r)$ is a function on $\mathbb{R}$ with a single non degenerate maximum at $r = 1$, i.e. p satisfies:

$$\frac{d}{dr}p(r)\bigg|_{r=1} = 0 \quad \text{and} \quad \frac{d^2}{dr^2}p(r)\bigg|_{r=1} < 0.$$

The distribution thus defined is everywhere of rank two, and is differentiable by construction. The flows of X_1, X_2 are denoted by $\{\phi_s\}$, $\{\psi_s\}$, respectively. In particular, we have $\phi_t(r, \theta, z) = (t+r, \theta, z)$, $\psi_t(r, \theta, z) = (r, \theta+t, z-p(r)t)$. Let $c : [0, 1] \to Q$ be an integral curve of X_1 through $q_0 = (r_0, \theta_0, z_0)$ at $t = 0$, with endpoint q_1. In the following we will prove that the integral curves of X_1 are not abnormal. We use the necessary and sufficient conditions on abnormal extremals developed in Theorem III-8.3. Note that these conditions do not depend on the cost function, i.e. the Riemannian metric on D, which allows us to say that the results below hold for any metric h on D. The vertical variational cone equals the subspace of $T_{q_1}Q$, given by

$$V^*_{(1,q_1)}R_{(0,q_0)} = \operatorname{span}\left\{ X_1(q_1), X_2(q_1), \frac{\partial}{\partial\theta}\bigg|_{q_1} - p(r_0+t)\,\frac{\partial}{\partial z}\bigg|_{q_1} \bigg| \forall t \in [0,1] \right\}.$$

This subspace coincides with the whole tangent space, i.e. $V^*_{(1,q_1)}R_{(0,q_0)} = T_{q_1}Q$, by observing that for an arbitrary $(v_r, v_\theta, v_z) \in T_qQ$:

$$\begin{aligned} v_r\,\frac{\partial}{\partial r}\bigg|_{q_1} + v_\theta\,\frac{\partial}{\partial\theta}\bigg|_{q_1} + v_z\,\frac{\partial}{\partial z}\bigg|_{q_1} &= v_r X_1(q_1) + v_\theta X_2(q_1) \\ &+ \frac{v_z + v_\theta p(r_0)}{p(r_0+t) - p(r_0)}(X_2 - \phi^*_{-t}X_2)(q_1), \end{aligned}$$

where t is chosen such that $p(r_0+t) \neq p(r_0)$. Consequently, in view of Theorem III-8.3, one can conclude that an integral curve of X_1 is not abnormal with respect to any metric on $D \to Q$. Let us obtain the same result using the sufficient conditions from Theorem III-8.5. Consider the following iteration of Lie brackets of X_1 and X_2:

$$[X_1, X_2] = -p'(r)\partial/\partial z;$$
$$[X_1, [X_1, X_2]] = -p''(r)\partial/\partial z;$$
$$\vdots$$

By assumption, we know that $p'(r) = 0$ iff $r = 1$ and that, for $r = 1$, $p''(1) < 0$ holds. Therefore, the vector fields $X_1, X_2, [X_1, X_2]$ and $[X_1, [X_1, X_2]]$ span the total tangent space at every point in Q. It is clear that the sufficient conditions from Theorem III-8.5 are more easily computed then the necessary and sufficient conditions from Theorem III-8.3.

We now repeat the above computations for the integral curves of X_2. Let $c' : [0, 1] \to Q$ be an integral curve of X_2, with $c'(0) = q_0 = (r_0, \theta_0, z_0)$ and endpoint q_1. The vertical variational cone now becomes

$$V^*_{(1,q_1)}R_{(0,q_0)} = \text{span}\left\{ X_1(q_1), X_2(q_1), \left.\frac{\partial}{\partial r}\right|_{q_1} + p'(r_0)t \left.\frac{\partial}{\partial z}\right|_{q_1} \;\middle|\; \forall t \in [0, 1] \right\}.$$

If q_0 is a point on the cylinder defined by $r = 1$, then one easily sees that $V^*_{(1,q_1)}R_{(0,q_0)} \neq T_{q_1}Q$ since $p'(1) = 0$. Therefore, every helix $c' : [0, 1] \to Q : t \mapsto (1, \theta + t, z - p(1)t)$ is an abnormal extremal, i.e. there exists a section $\theta(t)$ of D^0 along the curve c' through $q_0 = (1, 0, 0)$ such that

$$\theta(t) := T^*\psi_{-t}(p(1)\; d\theta|_x + dz|_x) = p(1)\; d\theta|_{(1,t,-p(1)t)} + dz|_{(1,t,-p(1)t)}\,,$$

or, equivalently $\nabla_\alpha \theta(t) = 0$, with ∇ a D-adapted g-connection.

The following step in our treatment consists of proving that c' is strictly abnormal. For that purpose, we need to work in the extended setting (i.e. on $Q \times \mathbb{R}$) and compute the extended vertical cone of variations. It is now necessary to fix a sub-Riemannian metric h, determining the energy cost E. In the example constructed by R. Montgomery [45], the metric h on D is given by $h_{11} = 1$, $h_{12} = h_{21} = 0$ and $h_{22} = r^2$, when expressed with respect to the basis $\{X_1, X_2\}$ of D. It is easily seen with this choice for h, that the sub-Riemannian length of a curve tangent to D is precisely the length of its projection on the (x, y)-plane with respect to the standard Riemannian metric on $\mathbb{R}^2$. In order to compute the extended vertical variational cone, we first recall the definition of the extended geometric control structure. The autonomous optimal control structure associated with this sub-Riemannian structure equals $(\tau_Q|_D, i)$, where i denotes the natural injection of D into TQ, and with cost the energy E. The anchor map in the extended setting $\overline{\rho} : (Q \times \mathbb{R}) \times \mathbb{R}^2 \to T(Q \times \mathbb{R})$ is defined by:

$$\overline{\rho}\big((q, J), (u^1, u^2)\big) = u^1 \frac{\partial}{\partial r} + u^2 \left(\frac{\partial}{\partial \theta} - p(r) \frac{\partial}{\partial z} \right) + \frac{1}{2}\big((u^1)^2 + (ru^2)^2\big) \frac{\partial}{\partial J},$$

where we have used the basis $\{X_1, X_2\}$ of D to fix a coordinate system (u^1, u^2) on the anchored bundle $\tau_Q|_D$. Note that, as is already explained in Theorem III-8.2, we have omitted explicit reference to the time coordinate. The vector field on $Q \times \mathbb{R}$ defined by $\overline{X}_2(q, J) = \partial/\partial\theta - p(r)\partial/\partial z + \frac{1}{2} r^2 \partial/\partial J$, with flow $\{\overline{\psi}_t\}$ given by $\overline{\psi}_t(r, \theta, z, J) = (r, \theta + t, z - p(r)t, J + \frac{1}{2} r^2 t)$, satisfies the property that its integral curve through $(q_0, 0)$, with $q_0 = (1, 0, 0)$, is precisely the curve $t \mapsto \big(c'(t), (\mathcal{J}^E_{c'})^{(0,t)}\big)$. The extended vertical variational cone is generated by tangent vectors of the form:

$$T\overline{\psi}_{1-t}\Big(\overline{\rho}\big(\overline{\psi}_t(q_0, 0), (u^1, u^2)\big)\Big) - \overline{X}_2(q_1) = u^1 \left(\frac{\partial}{\partial r} + (1-t) \frac{\partial}{\partial J} \right) + \\ (u^2 - 1) \left(\frac{\partial}{\partial \theta} - p(1) \frac{\partial}{\partial z} \right) + \frac{1}{2} \left((u^1)^2 + (u^2)^2 - 1 \right) \frac{\partial}{\partial J},$$

where $u^1, u^2 \in \mathbb{R}$ and $t \in]0, 1]$ are arbitrary. In order to prove that $-\partial/\partial J$ is contained in the boundary of extended variational cone, we introduce new control coordinates (s, ϕ): $u^1 = s \cos\phi$ and $u^2 = s \sin\phi$. By replacing u^1 and u^2 in the above family of tangent vectors by these new coordinates, we now construct two circles in the extended variational cone. Let $t = 1$, then it is easily seen that the circle in $T_{\overline{\psi}_1(q_0,0)}(Q \times \mathbb{R})$ with centre at the point corresponding to the tangent vector $-(\partial/\partial\theta - p(1)\partial/\partial z)$ and determined by the tangent vectors parameterised by $s = 1$ and $\phi \in [0, 2\pi]$, is entirely contained in the extended vertical variational cone. The tangent line to this circle at the origin is spanned by $\partial/\partial r$. Therefore, both vectors $\partial/\partial r$ and $-\partial/\partial r$ are contained in the closure of this cone (see Remark 3.4 and the picture below).

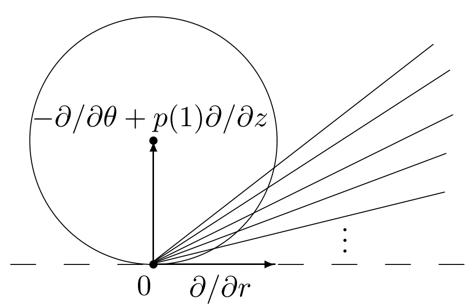

By taking $t = 0$, $s = 1$ and $\phi = [0, 2\pi]$ and repeating the above reasoning, we obtain that the straight line spanned by $-\partial/\partial r - \partial/\partial J$ is contained in the closure of the cone. Adding $\partial/\partial r$ to this vector, shows that $-\partial/\partial J$ is contained in the closure of the extended vertical variational cone, which proves that c' is strictly abnormal.

Before proceeding, we first give a physical interpretation of the above system. A curve $c(t) = (r(t), \theta(t), z(t))$ taking $(x_0, y_0, 0)$ to (x_0, y_0, z_1) and tangent to the distribution D, is a curve such that

$$\dot{z}(t) = -p(r)\dot{\theta}(t).$$

Therefore, we have that, using Stokes theorem:

$$z_1 = \int_0^1 -p(r(t))\dot{\theta}(t)dt = -\iint_{\Omega} B(r)rdrd\theta,$$

with $B(r) = (1/r)p'(r)$, which can be regarded as the strength of a magnetic field, normal to the (x, y)-plane, with vector potential

$$\Big(A_x = -\frac{p(r)}{r}\sin\theta, A_y = \frac{p(r)}{r}\cos\theta, A_z = 0\Big).$$

The surface Ω is the oriented surface enclosed by the projection of the curve c on the (x, y)-plane. The value $-z(1)$ therefore measures the flux of the magnetic field through Ω. Let us now consider the Hamiltonian H, generating normal extremals: $H = \frac{1}{2}(p_r^2 + \frac{1}{r^2}(p_\theta^2 - p_z p(r))^2$. If one expresses this function in the momenta with respect to the coordinates (x, y, z), then one obtains the Hamiltonian for a particle in the plane (x, y) with charge $q = p_z$ in a magnetic field with vector potential $(A_x, A_y, 0)$ orthogonal to the plane (note that $\dot{p}_z = 0$). The above connection between a normal extremal of the sub-Riemannian structure under consideration and the motion of a charged particle in a magnetic field was obtained by R. Montgomery [44, 45].

Remark 3.4. Let C denote a cone in a finite dimensional vector space $\mathcal{V}$, with vertex at the origin. Let $c : [0, 1] \to \mathcal{V}$ denote a curve through 0 at $t = 0$ such that $c([0, 1]) \subset C$. The tangent vector at 0 is defined by:

$$\dot{c}(0) = \lim_{h\to 0^+} \frac{1}{h}c(h).$$

Since $h > 0$ the argument of the above limit is contained in C for any h, which in turn implies that the limit itself is contained in the closure of C. We thus conclude that $\dot{c}(0) \in \mathrm{cl}\, C$. A similar argument can be applied if we consider a curve $c : [-1,0] \to \mathcal{V}$ such that $c([-1,0]) \subset C$ and $c(0) = 0$. Consider again the tangent vector at $t = 0$:

$$\lim_{h\to 0^-} \frac{1}{h} c(h) = - \lim_{(-h)\to 0^+} \frac{1}{(-h)} c(h).$$

We obtain that $-\dot{c} \in \mathrm{cl}\, C$.

We consider an example, constructed by W. Liu and H.J. Sussmann in [38], that contains strictly abnormal extremals that are minimising.

Example 3.5. Let $M = \mathbb{R}^3$ and D spanned by $X_1 = \partial/\partial x, X_2 = (1-x)\partial/\partial y + x^2\partial/\partial z$, where we use cartesian coordinates (x,y,z). The set $\{X_1, X_2\}$ forms a basis for D and determines a coordinate system (u^1,u^2) on the control space. This allows us to define a metric h on D and the cost function on D then becomes

$$E((x,y,z),u^1,u^2) = \frac{1}{2}\left((u^1)^2 + (u^2)^2\right),$$

Let us calculate the iterated brackets of X_1 and X_2:

$$\begin{aligned}[X_1, X_2] &= -\frac{\partial}{\partial y} + 2x\frac{\partial}{\partial z},\\ [X_1,[X_1,X_2]] &= 2\frac{\partial}{\partial z},\end{aligned}$$

from which follows easily that the integral curves of X_1 are not abnormal. It should be noted that this result is independent of the metric h on D. We investigate the abnormality of the integral curves of X_2. The flows $\{\phi_t\}$ of X_1 and $\{\psi_t\}$ of X_2 are given by $\phi_t(x,y,z) = (x+t,y,z)$ and $\psi_t(x,y,z) = (x,(1-x)t+y,x^2t+z)$. The pull-back of X_1 under ψ_t equals $\psi_t^*X_1 = \partial/\partial x + t\partial/\partial y - 2xt\partial/\partial z$, and this vector field can be written as a linear combination of X_1, X_2 for any value of t and at all points for which $x = 0$ or $x = 2$. Indeed, if $x = 0$, then $\psi_t^*X_1(0,y,z) = X_1(0,y,z) + tX_2(0,y,z)$. If $x = 2$, then $\psi_t^*X_1(2,y,z) = X_1(2,y,z) - tX_2(2,y,z)$. Therefore, each curve defined by $c : I \to M : t \mapsto (x,(1-x)t+y,x^2t+z)$ for any given point (x,y,z) with $x = 0$ or $x = 2$, is an abnormal extremal (i.e. the vertical variational cone equals D). We now compute that the integral curves of X_2 are strictly abnormal. The anchor map $\overline{\rho}$ in the extended setting $\mathbb{R}^4$ becomes (we leave out the time coordinate) in coordinates with $\overline{q} = (x,y,z,J)$:

$$\begin{aligned}\overline{\rho}(q,u^1,u^2) = u^1 \left.\frac{\partial}{\partial x}\right|_{\overline{q}} + u^2\left((1-x)\left.\frac{\partial}{\partial y}\right|_{\overline{q}} + x^2\left.\frac{\partial}{\partial z}\right|_{\overline{q}}\right)\\ +\frac{1}{2}\left((u^1)^2+(u^2)^2\right)\left.\frac{\partial}{\partial J}\right|_{\overline{q}}.\end{aligned}$$

Consider the vector field $\overline{X}_2$, defined on $\mathbb{R}^4$:

$$\overline{X}_2(\overline{q}) = \overline{\rho}(\overline{q}, 0, 1) = (1-x)\left.\frac{\partial}{\partial y}\right|_{\overline{q}} + x^2 \left.\frac{\partial}{\partial z}\right|_{\overline{q}} + \frac{1}{2}\left.\frac{\partial}{\partial J}\right|_{\overline{q}},$$

with flow $\{\overline{\psi}_t\}$ defined by $\overline{\psi}_t(x,y,z,J) = (\psi_t(x,y,z), J + \frac{1}{2}t)$. By definition, the extended vertical variational cone is generated by tangent vectors of the form:

$$\begin{aligned} T\overline{\psi}_{b-t}\Big(\overline{\rho}(\overline{\psi}_{t-a}(\overline{q}), u^1, u^2)\Big) - \overline{X}_2(\overline{\psi}_b(\overline{q})) = & \\ + \quad u^1 \left(\frac{\partial}{\partial x} - (b-t)\frac{\partial}{\partial y} + 2(b-t)x\frac{\partial}{\partial z}\right) & \\ + \quad (u^2 - 1)\left((1-x)\frac{\partial}{\partial y} + x^2\frac{\partial}{\partial z}\right) + \tfrac{1}{2}\Big((u^1)^2 + (u^2)^2 - 1\Big)\frac{\partial}{\partial J}, & \end{aligned} \tag{3.2}$$

with $u^1, u^2 \in \mathbb{R}$ and $t \in]a,b]$. Consider the following parametrisation: let $u^1 = s\cos\theta$ and $u^2 = s\sin\theta$, for $s > 0$ and $\theta \in [0, 2\pi[$. If we assume that $s = 1$, the coefficient of $\partial/\partial J$ becomes zero. If θ varies, we obtain a curve in the cone trough the origin. We know from Remark 3.4 that the tangent ray at the origin, i.e. the straight line spanned by

$$\frac{\partial}{\partial x} - (b-t)\frac{\partial}{\partial y} + 2(b-t)x\frac{\partial}{\partial z}$$

lies in the *closure* of the cone. If we substitute $b = t$, then we have that $\pm Y_1 = \pm\partial/\partial x$ is contained in the closure of the cone. If, on the other hand, $b > t$, then we obtain that $\pm Y_2 = \pm(\partial/\partial y - 2x\partial/\partial z)$ is contained in the closure of the cone. Now assume that $x = 0$ or $x = 2$. It is now easily seen that, given u^1, u^2 such that $(u^1)^2 + (u^2)^2 < 1$, then a linear combination (with positive coefficients) of the tangent vector in (3.2) and $\pm Y_1, \pm Y_2$ can be found which is proportional to $-\partial/\partial J$, up to a positive multiple.

4 Vakonomic dynamics and nonholonomic mechanics

As an application of our approach to sub-Riemannian structures in terms of generalised connections, we will see how to establish coordinate independent conditions for the motions of a free mechanical system subjected to linear nonholonomic constraints to be normal extremals with respect to an associated sub-Riemannian structure, and vice versa. We first give a definition of what is meant here by a free mechanical systems subjected to linear nonholonomic constraints, shortly called “a free nonholonomic mechanical system” and its associated sub-Riemannian structure. The definition below will be more detailed in the next chapter. We would like to emphasise that the equations of motion are closely related to the control equations for a control affine connection system (cf. Section III-9). However, the solutions to the free nonholonomic mechanical system are not in general optimal.

Assume that a manifold Q is equipped with a non-integrable regular distribution D and a Riemannian metric G. A free mechanical system with linear nonholonomic constraint D consists of a free particle described by the Lagrangian $L(v) = \frac{1}{2}G(v,v) \in C^\infty(TQ)$ and subjected to the constraint that the integral curves should live in D. (“Free” refers here to the absence of external forces.) The problem of determining the *dynamics of the free nonholonomic mechanical system* then consists in finding the solutions of the following equation (see [1, 4])

$$\pi(\nabla^G_{\dot{c}}\dot{c}(t)) = 0 \quad \text{and} \quad \dot{c}(t) \in D\ , \forall t,$$

where π is the orthogonal projection of $TQ = D \oplus D^\perp$ onto D with respect to G and, as before, ∇^G is the Levi-Civita connection corresponding to G. The associated sub-Riemannian structure is then given by (Q, D, h_G), with h_G the restriction of G to D.

In Chapter V we will construct a unique generalised connection ∇^{nh} on D over the linear vector bundle map $i : D \hookrightarrow TQ$, i.e. the natural inclusion map, namely: $\nabla^{nh}_X Y = \pi(\nabla^G_X Y)$. Here and in the sequel, we identify $X \in \Gamma(D)$ with $i \circ X \in \mathfrak{X}(Q)$. The i-connection ∇^{nh} preserves the sub-Riemannian metric h_G on D and satisfies $\nabla^{nh}_X Y - \nabla^{nh}_Y X - \pi[X,Y] = 0$ for all $X, Y \in \Gamma(D)$. In Chapter V it is proven that ∇^{nh} is completely determined by these two properties. A motion c of the free nonholonomic mechanical system is then characterised by the condition that $\nabla^{nh}_{\dot{c}}\dot{c}(t) = 0$, for all t.

The *vakonomic dynamical problem*, associated with the free particle with linear nonholonomic constraints, consists in finding normal extremals with respect to the associated sub-Riemannian structure (Q, D, h_G). It is interesting to compare the solutions of the nonholonomic mechanical problem with the solutions of the vakonomic dynamical problem, because the equations of motion for the mechanical problem are derived by means of (a generalisation of) d’Alembert’s principle, whereas the normal extremals are derived from a variational principle (see for instance Section III-9). This has been discussed for more general Lagrangian systems by J. Cortés et al. [10]. For the free particle case, we shall present here an alternative (coordinate free) approach.

Given a Riemannian metric G and a regular distribution D on a manifold Q, we can define the following two tensorial operators:

$$\begin{aligned}
&\Pi^G : \Gamma(D) \otimes \Gamma(D) \to \Gamma(D^\perp), (X, Y) \mapsto \pi^\perp(\nabla^G_X Y),\\
&\Pi^B : \Gamma(D) \otimes \Gamma(D^0) \to \Gamma((D^\perp)^0), (X, \eta) \mapsto \tau(\delta_X \eta).
\end{aligned}$$

It is indeed easily seen that both Π^G and Π^B are $C^\infty(Q)$-bilinear in their arguments and, hence, their action can be defined point-wise, with expressions like $\Pi^G(X_q, Y_q)$ and $\Pi^B(X_q, \eta_q)$, for $X_q, Y_q \in D_q$ and $\eta_q \in D^0$, having an obvious and unambiguous meaning.

The operator Π^B is related to the 'curvature' of the distribution D as follows: let $X, Y \in \Gamma(D)$, then one has:

$$\langle \Pi^B(X, \eta), Y \rangle = \langle \delta_X \eta, Y \rangle = -\langle \eta, [X, Y] \rangle, \text{ for any } \eta \in \Gamma(D^0).$$

Thus $\Pi^B \equiv 0$ if and only if D is involutive. The following lemma illustrates the importance of both tensors Π^G and Π^B. First, define a connection $\widetilde{\nabla}^B$ over $i : D \hookrightarrow TQ$ on the bundle D^0 by the prescription $\widetilde{\nabla}^B_X \eta = \tau^\perp(\delta_X \eta)$ with $X \in \Gamma(D)$ and $\eta \in \Gamma(D^0)$.

Lemma 4.1. *Consider a Riemannian metric G and a regular distribution D on a manifold Q. Assume that $c : I = [a, b] \to Q$ is a curve tangent to D and let ∇ be a D-adapted g-connection with respect to the associated sub-Riemannian structure (Q, D, h_G). Then, the following properties hold:*

1. *Given $Y_a \in D_{c(a)}$, denote the parallel transported curves along c, with initial point Y_a, with respect to ∇^{nh}, resp. ∇^G, by $\widetilde{Y}(t)$, resp. $Y(t)$. Then $\widetilde{Y}(t) = Y(t)$ for all $t \in I$, if and only if $\Pi^G(\dot{c}(t), \widetilde{Y}(t)) = 0$ for all $t \in I$.*
2. *Given $\eta_a \in D^0_{c(a)}$, denote the parallel transported curves along c, with initial point η_a, with respect to $\widetilde{\nabla}^B$, resp. ∇, by $\widetilde{\eta}(t)$, resp. $\eta(t)$. Then $\widetilde{\eta}(t) = \eta(t)$ if and only if $\Pi^B(\dot{c}(t), \widetilde{\eta}(t)) = 0$ for all $t \in I$.*

Proof. (1) From the definition of Π^G it follows that, given any section $\widetilde{Z}(t)$ of D along c, the following equation holds: $\nabla^{nh}_{\dot{c}} \widetilde{Z}(t) = \nabla^G_{\dot{c}} \widetilde{Z}(t) -$

$\Pi^G(\dot{c}(t), \widetilde{Z}(t))$. Assume that $\widetilde{Z}(t) = \widetilde{Y}(t) = Y(t)$, then we have

$$\Pi^G(\dot{c}(t), \widetilde{Y}(t)) = 0.$$

This already proves the statement in one direction. The converse follows from the fact that parallel transported curves with respect to any connection are uniquely determined by their initial conditions.

The proof of (2) follows from similar arguments. □

Note that property (2) of the previous lemma gives necessary and sufficient conditions for the existence of curves that are abnormal extremals, i.e.: c is an abnormal extremal if and only if there exists a parallel transported section $\widetilde{\eta}$ of D^0 along c with respect to $\widetilde{\nabla}^B$ such that, in addition, $\Pi^B(\dot{c}(t), \widetilde{\eta}(t)) = 0$ for all t. We shall now investigate some further properties of the operators Π^B and Π^G.

For $q \in Q$, let X_q be a non-zero element of D_q and consider the following subspace of T_qQ:

$$D_q + [X_q, D_q] \quad := \quad \mathrm{span}\{Y(q) + [\widetilde{X}, Y'](q) \mid Y, Y' \in \Gamma(D); \tag{4.3}$$
$$\widetilde{X} \in \Gamma(D) \text{ with } \widetilde{X}(q) = X\}. \tag{4.4}$$

As a side result of the next lemma it will be seen that the space $D_q + [X_q, D_q]$ is independent of the extension $\widetilde{X}$ of X_q used in its definition and, hence this also justifies the notation.

Lemma 4.2. *Let $\eta_q \in D_q^0$ and $X_q \in D_q$ for some $q \in Q$. The following equivalence holds:*

$$\Pi^B(X_q, \eta_q) = 0 \textit{ if and only if } \eta_q \in (D_q + [X_q, D_q])^0.$$

Proof. Let $\Pi^B(X, \eta) = 0$. This is equivalent to $\langle \eta, [\widetilde{X}, Y'](q) \rangle = 0$ for any $\widetilde{X}, Y' \in \Gamma(D)$ with $\widetilde{X}(q) = X_q$. Since $\eta_q \in D_q^0$, we may conclude that $\eta_q \in (D_q + [X, D_q])^0$. The converse follows by reversing the previous arguments. □

Another useful property is given by the following lemma.

Lemma 4.3. *Let Q be a manifold with a Riemannian metric G and a regular non-integrable distribution D, and consider the associated sub-Riemannian structure (Q, D, h_G). Let ∇ be a normal g-connection. We then have for $\alpha \in \mathfrak{X}^*(Q)$ that $\nabla_\alpha \alpha = 0$ if and only if*

$$\begin{aligned}&\flat_G(\nabla^{nh}_{g(\alpha)} g(\alpha)) = -\Pi^B(g(\alpha), \tau^\perp(\alpha)) \text{ and}\\ &\widetilde{\nabla}^B_{g(\alpha)} \tau^\perp(\alpha) = -\flat_G(\Pi^G(g(\alpha), g(\alpha))).\end{aligned}$$

Proof. From Theorem 2.3 one has that $\nabla_\alpha \alpha = 0$ if and only if $\nabla^G_{g(\alpha)} \tau(\alpha) + \delta_{g(\alpha)} \tau^\perp(\alpha) = 0$. Using the following relations

$$\begin{aligned}&\tau(\alpha) = \flat_G(g(\alpha)),\\ &\nabla^G \circ \flat_G = \flat_G \circ \nabla^G,\\ &\nabla^G_{g(\alpha)} g(\alpha) = \nabla^{nh}_{g(\alpha)} g(\alpha) + \Pi^G(g(\alpha), g(\alpha)),\\ &\delta_{g(\alpha)} \tau^0(\alpha) = \widetilde{\nabla}^B_{g(\alpha)} \tau^0(\alpha) + \Pi^B(g(\alpha), \tau^0(\alpha)),\end{aligned}$$

together with the fact that $T^*Q = \flat_G(D) \oplus D^0$ and $D^0 \cong \flat_G(D^\perp)$, the equivalence is immediately proven. □

The previous lemmas can now be used to derive necessary and sufficient conditions for a motion of a free nonholonomic mechanical system to be normal extremals (i.e. solution to the corresponding vakonomic problem) and vice versa. Let Q again be a manifold with a Riemannian metric G and a regular non-integrable distribution D.

Proposition 4.4. *A solution $c : [a, b] \to Q$ of the free nonholonomic system determined by the triple (Q, D, G) is a solution of the corresponding vakonomic problem, and vice versa, if and only if there exists a section η of D^0 along c such that*

$$\widetilde{\nabla}^B_{\dot{c}} \eta(t) = -\flat_G(\Pi^G(\dot{c}(t), \dot{c}(t))) \tag{4.5}$$

and such that, in addition, $\eta(t) \in (D_{c(t)} + [\dot{c}(t), D_{c(t)}])^0$ for all t.

Proof. The condition for any g-admissible curve $\alpha(t) = \flat_G(\dot{c}(t)) + \eta(t)$ with base curve c (where $\eta(t)$ is any section of D^0 along c) to be parallel transported with respect to a normal g-connection is that $\nabla_\alpha \alpha(t) = 0$. This can equivalently be written as:

$$\begin{aligned}&\flat_G(\nabla^{nh}_{\dot{c}} \dot{c}(t)) = -\Pi^B(\dot{c}(t), \eta(t)) \text{ and}\\ &\widetilde{\nabla}^B_{\dot{c}} \eta(t) = -\flat_G(\Pi^G(\dot{c}(t), \dot{c}(t))).\end{aligned}$$

Thus $\nabla^{nh}_{\dot{c}}\dot{c}(t) = 0$ if and only if $\Pi^B(\dot{c}(t), \eta(t)) = 0$, where $\eta(t)$ is a solution of $\widetilde{\nabla}^B_{\dot{c}}\eta(t) = -\flat_G(\Pi^G(\dot{c}(t), \dot{c}(t)))$. □

Remark 4.5. Given any $\eta_0 \in (D_{c(a)} + [\dot{c}(a), D_{c(a)}])^0$, then (4.5) always admits a solution, $\eta(t)$ with initial condition $\eta(a) = \eta_0$. The obstruction for c to be simultaneously a motion of the nonholonomic mechanical system and a solution of the vakonomic dynamical problem, lies in the fact that $\eta(t)$ should belong to $(D_{c(t)} + [\dot{c}(t), D_{c(t)}])^0$ for all t, and this is not guaranteed by the fact that $\eta(t)$ is a solution of (4.5). The search for geometric conditions for solutions $\eta(t)$ of equation (4.5) to remain in $(D_{c(t)} + [\dot{c}(t), D_{c(t)}])^0$ for all t, is left for future work.

V

Nonholonomic Mechanics

In this chapter we present an alternative approach to the treatment of nonholonomic systems with symmetry from a differential geometric perspective. Several authors have contributed to the field of nonholonomic mechanics (see [1, 4, 7, 8, 9, 10] and references therein). This has resulted in several quite different, but equivalent approaches to the subject. In this chapter we follow the affine connection approach, i.e. we formulate the equations of motion of a nonholonomic system using the notion of a (affine) connection [1].

1 General setting

It is well-known that the equations for a nonholonomic mechanical systems are derived from the Lagrange-d'Alembert principle, saying that the virtual work of the reaction forces should be zero. The geometric framework in which a nonholonomic system is formulated is basically the same as the geometric framework introduced in Section III-9.2, i.e. we assume that a Lagrangian system (Q, L) is given as well as a k-dimensional regular distribution D on Q, which is locally determined by $n-k$ independent functions $\Phi^\alpha : TQ \to \mathbb{R}$, $\alpha = 1, \ldots, n-k$, called constraint functions. The equations determining the dynamics of the nonholonomic system, then read

$$\frac{d}{dt}\left(\frac{\partial L}{\partial v^i}\right)\big(q^i(t), v^i(t)\big) - \frac{\partial L}{\partial q^i}\big(q^i(t), v^i(t)\big) = \lambda_\alpha \frac{\partial \Phi^\alpha}{\partial v^i} \text{ and } \Phi(\dot{c}(t)) = 0;$$

for some real-valued functions $\lambda_\alpha(t) \in \mathbb{R}$ with $\alpha = 1, \ldots, n-k$ and where we have put $\Phi = (\Phi^1, \ldots, \Phi^{n-k})$ and $c(t) = (q^i(t), v^i(t))$. Note that these equations are different from the ones obtained in Section III-9.2. In the remainder of this chapter we assume that the Lagrangian L takes the form

$$L(v) = \tfrac{1}{2}G(v, v) - V(\tau_Q(v)),$$

where G is a Riemannian metric on Q and V represents the potential energy function. Using standard techniques, and with $\{\eta^1, \ldots, \eta^{n-k}\}$ representing a basis of D^0, the equations of motion of the nonholonomic system can be put in the form:

$$\nabla^G_{\dot c}\dot c(t) = -\operatorname{grad} V\big(c(t)\big) + \sharp_G(\lambda_A\eta^A),$$

together with the constraint condition $\dot c(t) \in D$. The gradient is defined here using the Riemannian metric G as follows: $\operatorname{grad} V = \sharp_G(dV)$. One can verify (see [1]) that the above equation is equivalent to

$$\pi\big(\nabla^G_{\dot c}\dot c(t) + \operatorname{grad} V\big(c(t)\big)\big) = 0,$$

where $\pi : TQ \to D$ is the orthogonal projection defined using the metric G: we have a direct sum decomposition $TQ = D \oplus D^\perp$, where $D^\perp$ is the orthogonal complement of D (in this chapter we follow the same notations as in the previous chapter). The projection of TQ onto $D^\perp$ is denoted by $\pi^\perp$. If $V = 0$ (or if V is constant), i.e. if there are no external forces, then we say that we are studying *the nonholonomic free particle*. The equations of motion then reduce to $\pi(\nabla^G_{\dot c}\dot c(t)) = 0$, together with the constraint condition $\dot c(t) \in D$ for all t. Furthermore, one can define a connection $\bar\nabla$ on Q according to $\bar\nabla_X Y = \nabla^G_X Y + (\nabla^G_X \pi_{D^\perp})(Y)$ for $X, Y \in \mathfrak{X}(Q)$. This connection leaves D invariant and the equation of motion of the nonholonomic free particle can be rewritten as $\bar\nabla_{\dot c}\dot c(t) = 0$, with initial velocity taken in D (see [4, 33]).

There are many examples of nonholonomic systems that fit into the geometric framework proposed above. As an example, we consider the *Snakeboard*, which is a variant of the skateboard in which the passive wheel assemblies can pivot freely about a vertical axis (this example is taken from [7, 8, 36]). A peculiar characteristic of the Snakeboard is the fact that the rider can generate a snake-like locomotion without having to kick off the ground. The picture below sketches a simplified model. The human torso is simulated by a momentum wheel, rotating about the vertical axis through the centre of mass.

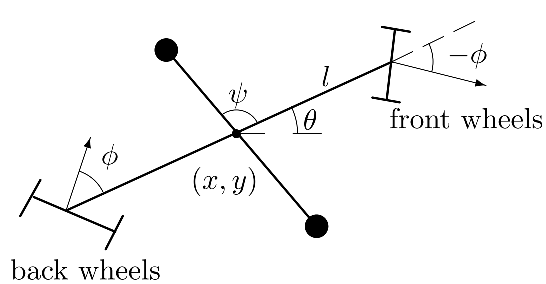

The configuration space can be identified with $Q = SE(2) \times S^1 \times S^1$, with local coordinates denoted by $(x, y, \theta, \psi, \phi)$. The requirement that the wheels do not slip in the direction of their axis imposes the two nonholonomic constraints:

$$\begin{aligned} &-\sin(\theta+\phi)\dot{x} + \cos(\theta+\phi)\dot{y} - l\cos\phi\,\dot{\theta} = 0;\\ &-\sin(\theta-\phi)\dot{x} + \cos(\theta-\phi)\dot{y} + l\cos\phi\,\dot{\theta} = 0; \end{aligned}$$

determining the distribution D. It follows that D is spanned by:

$$\frac{\partial}{\partial\psi}, \frac{\partial}{\partial\phi} \text{ and } -l\cos\phi\cos\theta\frac{\partial}{\partial x} - l\cos\phi\sin\theta\frac{\partial}{\partial y} + \sin\phi\frac{\partial}{\partial\theta}. \tag{1.1}$$

The kinetic energy Lagrangian determining the motion of the Snakeboard takes the form

$$L = \tfrac{1}{2}m(\dot{x}^2 + \dot{y}^2) + \tfrac{1}{2}(J + J_r + 2J_w)\dot{\theta}^2 + J_r\dot{\theta}\dot{\psi} + \tfrac{1}{2}J_r\dot{\psi}^2 + J_w\dot{\phi}^2,$$

where m is the mass of the Snakeboard, J is the moment of inertia of the board, J_w the moment of inertia of the wheels about the vertical axis and J_r the moment of inertia of the rotor. Following [8] we make the additional simplifying assumption $ml^2 = J + J_r + 2J_w$, which keeps the inertias on similar scales. The metric G on Q has components:

$$\begin{aligned} &G_{xx} = m && G_{yy} = m\\ &G_{\theta\theta} = J + J_r + 2J_w && G_{\theta\psi} = J_r = G_{\psi\theta}\\ &G_{\psi\psi} = J_r && G_{\phi\phi} = 2J_w; \end{aligned}$$

where all other components are zero.

We now reconsider the nonholonomic free particle from the point of view of connections over a vector bundle map. It is our goal to define a generalised connection for which the "geodesic equation" is equivalent to the equations of motion of the nonholonomic free particle. For that purpose, let $i : D \to TQ$ denote the natural embedding of D into TQ. In the sequel we will again identify $X \in \Gamma(D)$ with $Ti \circ X$, regarded as a vector field on Q. In terms of the notations used in Chapter II, we consider the following situation: $N = E = D$, $\nu = \pi = (\tau_Q)_{|D}$ and $\rho = i$. We may now define a connection $\nabla^{nh} : \Gamma(D) \times \Gamma(D) \to \Gamma(D)$ over i on the vector bundle $\pi : D \to Q$ by the prescription

$$\nabla^{nh}_X Y = \pi(\nabla^G_X Y)\,,$$

where the superscript nh stand for "nonholonomic". It is easily seen that this determines indeed a i-connection and that, moreover, $\nabla^{nh}_X Y = \bar{\nabla}_X Y$ for $X, Y \in \Gamma(D)$. Admissible curves in this setting are curves $\widetilde{c}$ in D that are prolongations of curves in Q, i.e. $\widetilde{c}(t) = \dot{c}(t)$ for some curve c in Q. Note that for any base curve c, $\dot{c}$ may be regarded here both as an admissible curve in D and as a section of π defined along c. It follows that the equation of motion of the given nonholonomic problem can be written as $\nabla^{nh}_{\dot{c}}\dot{c}(t) = 0$, with c a curve in Q tangent to D. Admissible curves satisfying this equation are also called autoparallel curves (see also Chapter IV page 135). We now continue to investigate some properties of the nonholonomic connection ∇^{nh}.

The restriction of the given Riemannian metric G on Q to sections of D defines a bundle metric on D which we denote by G^o (to avoid notational confusion with the notion of a lift over an anchor map introduced in the following section, we use a different notion for the restriction of G to the distribution D in comparison with the previous section). The i-connection ∇^{nh} considered above now admits the following characterisation.

Proposition 1.1. *∇^{nh} is uniquely determined by the conditions that it is 'metric', i.e. for all $X, Y, Z \in \Gamma(D)$ one has*

$$X(G^o(Y, Z)) = G^o(\nabla^{nh}_X Y, Z) + G^o(Y, \nabla^{nh}_X Z),$$

and that it satisfies

$$\nabla^{nh}_X Y - \nabla^{nh}_Y X = \pi[X, Y],$$

for all $X, Y \in \Gamma(D)$.

Proof. First we prove that ∇^{nh} satisfies both conditions. Using the fact that ∇^G is metric for G, and regarding sections of D as vector fields on Q, we find:

$$\begin{aligned} X(G^o(Y, Z)) &= X(G(Y, Z)) \\ &= G(\nabla^G_X Y, Z) + G(Y, \nabla^G_X Z) \\ &= G^o(\nabla^{nh}_X Y, Z) + G^o(Y, \nabla^{nh}_X Z), \end{aligned}$$

where the last equality follows from the fact that $G(X, Y) = 0$ whenever $X \in \Gamma(D)$ and $Y \in \Gamma(D^\perp)$. The second condition follows from the symmetry property of ∇^G (i.e. ∇^G has zero torsion).

Conversely, let ∇ be an arbitrary i-connection that satisfies both conditions. One then easily derives that for any chosen $X, Y \in \Gamma(D)$ and for all $Z \in \Gamma(D)$

$$\begin{aligned}2G^o(\nabla_X Y, Z) &= X(G(Y,Z)) + Y(G(X,Z)) - Z(G(X,Y)) \\ &\quad + G(\pi[X,Y],Z) - G(\pi[X,Z],Y) - G(X,\pi[Y,Z]) \\ &= 2G(\nabla^G_X Y, Z)\,,\end{aligned}$$

from which one readily deduces that $\nabla_X Y = \pi(\nabla^G_X Y)$, i.e. $\nabla \equiv \nabla^{nh}$. □

It is easily proven that if D is an integrable distribution, it induces a foliation of Q (i.e. the given constraints are holonomic), then the connection ∇^{nh} induces the Levi-Civita connection on the leaves of this foliation with respect to the induced metric.

By definition, the nonholonomic connection ∇^{nh} is metric, i.e. for any $X, Y \in \Gamma(D)$ we have that

$$X(G^o(X,Y)) = G^o(\nabla^{nh}_X X, Y) + G^o(X, \nabla^{nh}_X Y)$$

holds. The second term on the right-hand side can be rewritten as

$$\begin{aligned}G^o(X, \nabla^{nh}_X Y) &= G^o(X, \nabla^{nh}_Y X) + G(X,[X,Y]) \\ &= \tfrac{1}{2}\mathcal{L}_Y(G(X,X)) + G(X,[X,Y]) \\ &= \tfrac{1}{2}(\mathcal{L}_Y G)(X,X)\,,\end{aligned}$$

(with $\mathcal{L}$ denoting the Lie derivative operator). With any given $Y \in \Gamma(D)$ one can associate a function J_Y on D, given by $J_Y(X_q) := G^o(X_q, Y(q))$, for all $q \in Q$ and $X_q \in D_q$. Using the preceding identities, and considering a base curve c in Q which is "geodesic" with respect to ∇^{nh} (i.e. a solution of the nonholonomic equations), one easily derives that

$$\frac{d}{dt}(J_Y(\dot{c}))(t) = \tfrac{1}{2}(\mathcal{L}_Y G)(\dot{c}(t), \dot{c}(t)).$$

This equation implies that the condition that a section Y of D which, regarded as a vector field on Q, leaves the metric G invariant (i.e. is a Killing vector field) is a sufficient condition for the function J_Y to determine a conserved quantity for the given nonholonomic system. However, a necessary and sufficient condition for J_Y to be conserved is that the restriction of the tensor $(\mathcal{L}_Y G)$ to $D \otimes D$ is zero, i.e. $(\mathcal{L}_Y G)|_{D\otimes D} = 0$.

2 Reduction of the nonholonomic free particle with symmetry.

Consider again a nonholonomic free particle with configuration space Q, kinetic energy Lagrangian $L(v) = \frac{1}{2}G(v,v)$, and linear nonholonomic constraints represented by a regular distribution D on Q. We now investigate the role of symmetries in the description of such systems.

Let H be a Lie group defining a free and proper right action on Q, denoted by $R_a : Q \to Q, q \mapsto R_a(q) = qa$, for all $a \in H$, such that we have a principal fibre bundle $Q \xrightarrow{\mu} \hat{Q} := Q/H$. Assume this action leaves invariant both the Riemannian metric G and the constraint distribution D, i.e. $R_a^*G = G$ and $TR_a(D) \subset D$ for all $a \in H$. We already know from above that the equations of motion of the nonholonomic free particle are given by the "geodesic" equations: $\nabla^{nh}_{\dot{c}}\dot{c}(t) = 0$. Using the symmetry assumption (i.e. the H-invariance of G and D), it is easily proven that if $c(t)$ is a solution, so is $c(t)a$ for all $a \in H$. Therefore, one obtains equivalence classes of solutions, where two solutions c_1 and c_2 are called equivalent iff $c_1 = c_2a$ for some $a \in H$. In the reduction procedure described below, it is our intention to construct a reduced connection over a suitable vector bundle map, such that the corresponding "geodesics" are precisely these equivalence classes.

First of all, we note that the set D/H, the quotient space of D under the lifted action of H on D, admits a vector bundle structure over $\hat{Q}$, with projection $\tau : D/H \to \hat{Q}$ defined by $\tau([X_q]) = \mu(q)$. Here, $[X_q]$ represents the H-orbit of $X_q \in D$ under the lifted right action. Using the fact that this action on D is fibre linear, and relying on the local triviality of the principal bundle $Q \to \hat{Q}$, one can verify that τ indeed determines a vector bundle structure (see e.g. [49, p 29]). Next, we define a map $\rho : D/H \to T\hat{Q}$ according to $\rho([X_q]) := T\mu(X_q)$. Once again one can easily see that this map is well defined (i.e. does not depend on the chosen representative X_q of $[X_q]$) and is fibred over the identity on $\hat{Q}$. We now first construct a principal ρ-connection on Q which, subsequently, will be used to define a ρ-connection on D/H.

Let $h : \mu^*(D/H) \to TQ : (q,[X_q]) \to X_q$, i.e. we take the image $h(q,[X_q])$ to be the unique tangent vector at q belonging to the equivalence class $[X_q]$. Since the action of H is free, it immediately follows that h is well defined and, moreover, $\operatorname{im} h = D$. We can also verify that $h(qa,[X_q]) = TR_a(X_q) = TR_a(h(q,[X_q]))$ and $T\mu(h(q,[X_q])) = \rho([X_q])$. Consequently, h determines a principal ρ-connection on Q.

Note that sections of the bundle $\tau : D/H \to \hat{Q}$ can be put into one to one correspondence with the set of right invariant vector fields on Q taking values in D (i.e. the right equivariant sections of $D \to Q$). Indeed, for $\psi \in \Gamma(D/H)$ and $q \in Q$ such that $\mu(q) \in \mathrm{dom}\,\psi$, put

$$\psi^h(q) := h(q, \psi(\mu(q)))\,.$$

Then ψ^h is a H-equivariant section of D. On the other hand, let X be a right invariant vector field on Q with values in D. Then, define an element X_h of $\Gamma(D/H)$ by

$$X_h(\hat{q}) = [X_q],$$

with $q \in \mu^{-1}(\hat{q})$. Clearly, this does not depend on the choice of q in the fibre over $\hat{q}$. Thus, by means of h we have established a bijective correspondence between $\Gamma(D/H)$ and the set of H-equivariant sections of $D \to Q$. For the following derivation of a reduced ρ-connection on D/H, we may refer to F. Cantrijn et al. [4] where, at least for the so-called Chaplygin-case, a similar construction has been made in terms of 'ordinary' connections and, therefore, we will not enter here into details. For completeness, however, we recall the following useful properties. First, from the H-invariance of G one can deduce that the vector field $\nabla^G_X Y$ is right invariant whenever $X, Y \in \mathfrak{X}(Q)$ are right invariant, and that $\pi : TQ \to D$ commutes with TR_a for any $a \in H$. Secondly, the symmetry assumptions also imply that the induced bundle metric G^o on D is H-invariant and, hence, determines a reduced bundle metric $\hat{G}^o$ on D/H. Using h we can construct $\hat{G}^o$ as follows: for any $\phi, \psi \in \Gamma(D/H)$ put

$$\hat{G}^o(\hat{q})(\phi(\hat{q}), \psi(\hat{q})) := G^o(q)(\phi^h(q), \psi^h(q)),$$

with $q \in \mu^{-1}(\hat{q})$. Let $a \in H$, then

$$\begin{aligned} G^o(qa)(\phi^h(qa), \psi^h(qa)) &= G(qa)(TR_a\phi^h(q), TR_a\psi^h(q)) \\ &= G^o(q)(\phi^h(q), \psi^h(q)), \end{aligned}$$

where, again, we have relied on the H-invariance of G. From this we may conclude that $\hat{G}^o$ is indeed well defined.

Let ∇^{nh} be the nonholonomic connection over i introduced in the previous section. We now construct a ρ-connection on the bundle D/H, as follows: for any $\psi, \phi \in \Gamma(D/H)$ put

$$\hat{\nabla}^{nh}_{\psi}\phi = (\nabla^{nh}_{\psi^h}\phi^h)_h\,.$$

Again, one may check that this is well defined and verifies the conditions of a ρ-connection.

Proposition 2.1. *The ρ-connection $\hat{\nabla}^{nh}$ is metric with respect to the reduced bundle metric $\hat{G}^o$ on D/H, and satisfies the property:*

$$\hat{\nabla}^{nh}_{\psi}\phi - \hat{\nabla}^{nh}_{\phi}\psi - [\psi, \phi] = 0,$$

where, by definition, $[\psi, \phi] := (\pi[\psi^h, \phi^h])_h$.

Proof. For any $\psi \in \Gamma(D/H)$, we have that ψ^h is μ-related to $\rho \circ \psi$ as vector fields on Q and $\hat{Q}$, respectively. Using this, together with the properties of ∇^{nh}, we can prove that $\hat{\nabla}^{nh}$ is metric with respect to $\hat{G}^o$. Indeed, let $\psi, \phi, \eta \in \Gamma(D/H)$, then

$$\begin{aligned}
(\rho \circ \psi)(\hat{G}^o(\phi, \eta)) \circ \mu &= \psi^h(\hat{G}^o(\phi, \eta) \circ \mu) \\
&= \psi^h(G^o(\phi^h, \eta^h)) \\
&= G^o(\nabla^{nh}_{\psi^h}\phi^h, \eta^h) + G^o(\phi^h, \nabla^{nh}_{\psi^h}\eta^h) \\
&= \Big(\hat{G}^o(\hat{\nabla}^{nh}_{\psi}\phi, \eta) + \hat{G}^o(\phi, \hat{\nabla}^{nh}_{\psi}\eta)\Big) \circ \mu,
\end{aligned}$$

from which the result readily follows.

The second property can also be proven in a straightforward manner. □

It is also not difficult to verify that $\hat{\nabla}^{nh}$ is uniquely determined by the two properties mentioned in the proposition.

To complete the reduction picture, it can be proved that every solution of the geodesic equation for ∇^{nh} projects onto a solution of the "geodesic problem" for the reduced nonholonomic connection $\hat{\nabla}^{nh}$ in the following sense. Assume that c is a solution of the nonholonomic equations, i.e. $\nabla^{nh}_{\dot{c}}\dot{c}(t) = 0$. Consider the curve $\hat{c} = \mu \circ c$ in $\hat{Q}$. Then the section $[\dot{c}](t) = [\dot{c}(t)]$ of D/H along $\hat{c}$ is autoparallel with respect to the ρ-connection $\hat{\nabla}^{nh}$, i.e. $\hat{\nabla}^{nh}_{[\dot{c}]}[\dot{c}](t) = 0$. This follows from the fact that for each $q \in Q$, $h(q, .) : \tau^{-1}(\mu(q)) \to T_qQ$ is injective and that for any base curve c in Q,

$$h(c(t), \hat{\nabla}^{nh}_{[\dot{c}]}[\dot{c}](t)) = \nabla^{nh}_{\dot{c}}\dot{c}(t), \ \forall t.$$

On the other hand, any solution $[\dot{c}]$ of the equation $\hat{\nabla}^{nh}_{[\dot{c}]}[\dot{c}](t) = 0$ determines an equivalence class of solutions of the initial nonholonomic problem on Q. Given any point c_0 in $\mu^{-1}(\tau([\dot{c}](0)))$, a unique curve c in Q can be constructed which is horizontal with respect to the principal ρ-connection

h on Q, i.e. c satisfies for all t: $\dot{c}(t) = h(c(t), [\dot{c}](t))$ with initial condition $c(0) = c_0$ (note that $[\dot{c}(t)] = [\dot{c}](t)$). It is easily seen that $\mu(c) = \tau([\dot{c}])$ and from this we can deduce $\nabla^{nh}_{\dot{c}}\dot{c}(t) = 0$.

We conclude that the set of equivalence classes of solutions of the free nonholonomic mechanical problem in Q is in a one-to-one correspondence with the set of solutions of autoparallel admissible curves with respect to the reduced nonholonomic connection (i.e. using the principal ρ-connection h).

To close this section, we note that much of the preceding discussion can be easily extended to more general nonholonomic systems with symmetry, admitting forces derivable from a H-invariant potential energy function.

3 The Snakeboard revisited

Let us reconsider the example of the Snakeboard from the beginning of this chapter with configuration space $Q = SE(2) \times S^1 \times S^1$. The constraints as well as the metric G are invariant under the right action of $SE(2)$. Denoting the elements of $SE(2)$ by $g = (a, b, \alpha)$, this action is given by

$$R_g(x, y, \theta, \psi, \phi) = (x \cos\alpha - y \sin\alpha + a, x \sin\alpha + y \cos\alpha + b, \theta + \alpha, \psi, \phi).$$

Also the action of S^1 on Q, defined by $R_\beta(x, y, \theta, \psi, \phi) \mapsto (x, y, \theta, \psi + \beta, \phi)$ leaves the constraints and the metric invariant. Putting $H = SE(2) \times S^1$, the direct product of $SE(2)$ and S^1, the configuration space Q thus inherits the structure of a principal fibre bundle with structure group H over the base space S^1. The three vector fields from (1.1) form a basis for D which is invariant under this action and therefore, since these vector fields correspond to sections of $D/H \to S^1$, they determine a basis for the bundle D/H, which will be denoted by $\{e_1, e_2, e_3\}$.

Since, eventually, we have to find the coordinate expression for the equation $\hat{\nabla}^{nh}_{[\dot{c}]}[\dot{c}] = 0$, it will be profitable to work with the following basis of TQ

$$\Big\{ X_1 = \frac{\partial}{\partial\psi},\ X_2 = \frac{\partial}{\partial\phi},$$

$$X_3 = -l\cos\phi\cos\theta\frac{\partial}{\partial x} - l\cos\phi\sin\theta\frac{\partial}{\partial y} + \sin\phi\frac{\partial}{\partial\theta},$$

$$X_4 = \sin\theta\frac{\partial}{\partial x} - \cos\theta\frac{\partial}{\partial y},$$

$$X_5 = l\sin\phi\cos\theta\frac{\partial}{\partial x} + l\sin\phi\sin\theta\frac{\partial}{\partial y} + \cos\phi\frac{ml^2}{ml^2 - J_r}\left(\frac{\partial}{\partial\theta} - \frac{\partial}{\partial\psi}\right)\Big\},$$

where X_4, X_5 are a basis for $D^\perp$. Recall the definition of a lift of a section of $D/H \to S^1$ to a vector field on Q, then we can write $e_i^h = X_i$ for $i = 1, 2, 3$. In order to obtain the local equations for $\hat{\nabla}^{nh}$ one might compute the connection coefficients for this connection. These are non-trivial functions and it would require a long and tedious calculation to derive them. Therefore, we shall follow a different route. Let $[\dot{c}](t) = \omega^i(t)e_i(\psi(t), \phi(t))$, with $\dot{\phi} = \omega_2$ (this is the admissibility condition). We now derive the coordinate expression of the equation $\pi(\nabla^G_{\dot{c}}\dot{c}) = 0$, where $\dot{c} = [\dot{c}]^h$, i.e. in coordinates:

$$\dot{c} = \omega^1 X_1 + \omega^2 X_2 + \omega^3 X_3.$$

Since G has only constant coefficients, we have that

$$\nabla^G_{\dot{c}}\dot{c}(t) = \dot{\omega}^1(t)X_1(c(t)) + \dot{\omega}^2(t)X_2(c(t)) + \dot{\omega}^3(t)X_3(c(t)) + \omega^3(t)\dot{X}_3(t),$$

where $\dot{X}_3(t)$ is the tangent vector, defined by:

$$\frac{d}{dt}\left(-l\cos\phi\cos\theta\right)\frac{\partial}{\partial x} - \frac{d}{dt}\left(l\cos\phi\sin\theta\right)\frac{\partial}{\partial y} + \frac{d}{dt}\left(\sin\phi\right)\frac{\partial}{\partial\theta},$$

where d/dt represents the time derivation along $c(t)$ at t. The orthogonal projection of this tangent vector on D, gives us the coefficients of $\pi(\nabla^G_{\dot{c}}\dot{c}(t))$ with respect to the basis $\{X_1, X_2, X_3\}$ of D and in turn the coefficients of $\hat{\nabla}^{nh}_{[\dot{c}]}[\dot{c}]$ with respect to $\{e_1, e_2, e_3\}$, where $a = ml^2/(ml^2 - J_r)$:

$$\begin{aligned}\hat{\nabla}^{nh}_{[\dot{c}]}[\dot{c}] \quad = \quad & \left(\dot{\omega}^1 + \frac{a\cos\phi}{a\cos^2\phi + \sin^2\phi}\dot{\phi}\omega^3\right)e_1 + \dot{\omega}^2 e_2 + \\ & \left(\dot{\omega}^3 + \frac{(1-a)\cos\phi\sin\phi}{a\cos^2\phi + \sin^2\phi}\dot{\phi}\omega^3\right)e_3.\end{aligned}$$

If we substitute $\omega^2 = \dot{\phi}$, then the equations for $\omega(t) = (\omega^1(t), \omega^2(t), \omega^3(t))$ to be autoparallel with respect to $\hat{\nabla}^{nh}$ are precisely:

$$\begin{aligned}\dot{\omega}^1 &= -\frac{a\cos\phi}{a\cos^2\phi + \sin^2\phi}\omega^2\omega^3;\\ \dot{\omega}^2 &= 0;\\ \dot{\omega}^3 &= -\frac{(1-a)\cos\phi\sin\phi}{a\cos^2\phi + \sin^2\phi}\omega^2\omega^3.\end{aligned}$$

It is easily seen that the connection coefficients only depend on the coordinate ϕ. The third equation is easily integrated, giving

$$\omega^3 = K^3\left(a\cos^2\phi + \sin^2\phi\right)^{-\frac{1}{2}},$$

with K^3 an integration constant. If we substitute this expression for ω^3 into the differential equation for ω^1, we obtain, after integration

$$\omega^1 = -\sin\phi\,\omega^3 + K^1,$$

(K^1 =constant). Since $\dot\phi = \omega^2 = K^2$ is constant, we have that $\phi(t) = K^2 t + \phi_0$. The expressions for x, y, θ, ψ then follow from reconstruction, which involves integrating the system of differential equations:

$$\begin{aligned}
&\dot x = -l\cos\phi\cos\theta\omega^3;\\
&\dot y = -l\cos\phi\sin\theta\omega^3;\\
&\dot\theta = \sin\phi\omega^3;\\
&\dot\psi = \omega^1,
\end{aligned}$$

where on the right hand side ω^1, ω^2 and ϕ are given by the expressions determined above. It should be noted that the expression for ω^3 is closely related to the modified "Lagrange-d'Alembert-Poincare" function from [7]. Let us consider the vector field $\partial/\partial\psi \in \Gamma(D)$. It is easily seen that it is a fundamental vertical vector field of the principal fibre bundle $Q \to S^1$, which leaves the metric G invariant (since G is right invariant). Therefore, the function $v \mapsto G(v, \partial/\partial\psi)$ on D is conserved along solutions of the nonholonomic mechanical system. Locally this function reads $\dot\theta + \dot\psi$. Using the above expressions for $\dot\theta$ and $\dot\psi$, we immediately obtain that $\dot\psi + \dot\theta = K^1$.

4 Some remarks

Our approach to the reduction problem of a nonholonomic free particle with symmetry, using the notion of generalised connections over a bundle map, differs from other approaches in that we do not have to make any additional assumption regarding the constraint distribution D. In treatments of the so-called Chaplygin case, for instance, the assumption is that D is the horizontal distribution of a principal H-connection (see e.g. [4, 9, 23]), i.e. besides being H-invariant D also satisfies $TQ = D \oplus \ker T\mu$. In the more general case treated e.g. by H. Cendra et al. [7], it is assumed that $TQ = D + \ker T\mu$ (but one may have $D \cap \ker T\mu \neq \{0\}$). In our treatment we only require H-invariance of D.

Bibliography

[1] A.M. Bloch and P.E. Crouch. Newton's law and integrability of nonholonomic systems. *SIAM J. Control Optim.*, 36:2020–2039, 1998.

[2] R. Bott, S. Gitler, and I.M. James. *Lectures on Algebraic and Differential topology*, volume 279 of *Lecture Notes in Mathematics*. Springer, 1972.

[3] F. Brickell and R.S. Clark. *Differentiable Manifolds. An Introduction*. Van Nostrand Reinhold, London, 1970.

[4] F. Cantrijn, J. Cortés, M. de Léon, and D. Martín de Diego. On the geometry of generalized Chaplygin system. *Math. Proc. Camb. Phil. Soc.*, 132:323–351, 2002.

[5] F. Cantrijn and B. Langerock. Generalised connections over a vector bundle map. *Diff. Geom. Appl.*, 2002. to appear (math.DG/0201274).

[6] J.F. Cariñena and E. Martínez. Lie algebroid generalization of geometric mechanics. In J. Kubarski, P. Urbański, and R. Wolak, editors, *Lie algebroids and related topics in differential geometry*, volume 54 of *Banach Center Publications*, pages 201–215. Warzawa, 2001.

[7] H. Cendra, J.E. Marsden, and T.S. Ratiu. Geometric mechanics and lagrangian reduction and nonholonomic mechanics. In B. Enguist and W. Schmid, editors, *Mathematics Unlimited-2001 and Beyond*, pages 221–273. Springer-Verlag, New York, 2000.

[8] J. Cortés. *Geometric, Control and Numerical aspects of Nonholonomic Systems*, volume 1793 of *Lecture Notes in Mathematics*. Springer, Berlin, 2002.

[9] J. Cortés and M. de León. Reduction and reconstruction of the dynamics of nonholonomic systems. *J. Phys. A: Math. Gen.*, 32:8615–8645, 1999.

[10] J. Cortés, M de León, D. Martín de Diego, and S. Martínez. Geometric description of vakonomic and nonholonomic dynamics, comparison of solutions. *SIAM J. Control Optim.*, 41(5):1389–1412, 2003.

[11] M. Crainic and R.L. Fernandez. Integrability of lie brackets. *Annals of Mathematics.* to appear.

[12] F. Etayo. A coordinate-free survey on pseudo-connections. *Rev. Acad. Canar. Cienc.*, 5:125–137, 1993.

[13] E. Falbel, C. Gorodski, and M. Rumin. Holonomy of sub-riemannian manifolds. *Int. J. Math.*, 8:317–344, 1997.

[14] R.L. Fernandes. Connections in poisson geometry. I: Holonomy and invariants. *J. Diff. Geom.*, 54:303–366, 2000.

[15] R.L. Fernandes. Lie algebroids, holonomy and characterisitic classes. *Adv. Math.*, 170:119–179, 2002.

[16] J. Grabowski, K. Grabowska, and P. Urbanski. Lie brackets on affine bundles. 2002. preprint (math.DG/0203112).

[17] J. Grabowski and P. Urbanski. Lie algebroids and Poisson-Nijenhuis structures. *Rep. Math. Phys.*, 40:195–208, 1997.

[18] J. Grabowski and P. Urbanski. Algebroids - general differential calculi on vector bundles. *J. Geom. Phys.*, 31:111–141, 1999.

[19] W. Greub, S. Halperin, and R. Vanstone. *Connections, Curvature and Cohomology, Vol. I, II and III*, volume 47 of *Pure and Applied Mathematics.* Academic Press, New York, 1972.

[20] V. Hutson and J.S. Pym. *Applications of functional analysis and operator theory*, volume 146 of *Mathematics in science and engineering.* Academic Press, London, 1980.

[21] F. Kamber and P. Tondeur. *Foliated bundles and characteristic classes*, volume 493 of *Lecture Notes in Math.* Springer, Berlin, 1975.

[22] S. Kobayashi and K. Nomizu. *Foundations of differential geometry*, volume I and II. Intersience Publishers, 1963.

[23] J. Koiller. Reduction of some classical non-holonomic systems with symmetry. *Arch. Ration Mech. Anal.*, 118:113–148, 1992.

[24] I. Kolář, P. Michor, and J. Slovák. *Natural Operations in Differential Geometry*. Springer, Berlin, 1993.

[25] Y. Kosmann-Schwarzbach and F. Magri. Poisson-nijenhuis structures. *Ann. Inst. Henri Poincaré, Phys. Théor.*, 53:35–81, 1990.

[26] G. Köthe. *Topological vector spaces I*. Springer-Verlag, Berlin, 1969.

[27] B. Langerock. Nonholonomic mechanics and connections over a bundle map. *J. Phys. A: Math. Gen.*, 34:L609–L615, 2001.

[28] B. Langerock. A connection theoretic approach to sub-Riemannian geometry. *J. Geom. Phys.*, 2002. to appear (math.DG/0210004).

[29] B. Langerock. Connections in sub-Riemannian geometry. In *Proceedings of the 8th International Conference on Differential Geometry and Its Applications*, volume 3 of *Mathematical Publications*. Silesian University at Opava, 2002.

[30] B. Langerock. Autonomous optimal control problems. *Rep. Math. Phys.*, 2003. to appear.

[31] B. Langerock. Geometric aspects of the maximum principle and lifts over a bundle map. *Acta Appl. Math.*, 2003. to appear (math.DG/0212055).

[32] A. D. Lewis and R. M. Murray. Configuration controllability of simple mechanical control systems. *SIAM Review*, 41(3):555–574, 1999.

[33] A.D. Lewis. Affine connections and distributions with applications to nonholonomic mechanics. *Rep. Math. Phys.*, 42(1/2):135–164, 1998.

[34] A.D. Lewis. Simple mechanical control systems with constraints. *IEEE Transactions on Automatic Control*, 45(8):1420–1436, 2000.

[35] A.D. Lewis. The geometry of the maximum principle for affine connection control systems. 2001.

[36] A.D. Lewis, J.P. Ostrowski, R.M. Murray, and J.W. Burdick. Nonholonomic mechanics and locomotion: the snakeboard example. In *Proc. IEEE Conf. Robotics & Automation*, pages 2391–2397. San Diego, CA, 1994.

[37] P. Libermann and C.-M. Marle. *Symplectic Geometry and Analytical Mechanics*. Reidel, Dortrecht, 1987.

[38] W. Liu and H.J. Sussmann. Shortest paths for sub riemannian metrics on rank two distributions. *Memoirs AMS*, 118, 1995.

[39] L. Mangiarotti and G. Sardanashvily. *Connections in Classical and Quantum Field Theory.* World Scientific, Singapore, 2000.

[40] E. Martínez. Lagrangian Mechanics on Lie algebroids. *Acta. Appl. Math.*, 67:295–320, 2001.

[41] E. Martínez, T. Mestdag, and W. Sarlet. Lie algebroid structures and lagrangian systems on affine bundles. *J. Geom. Phys.*, 44:70–95, 2002.

[42] T. Mestdag. Generalised connections on affine Lie algebroids. 2002. preprint.

[43] T. Mestdag, W. Sarlet, and E. Martínez. Note on generalised connections and affine bundles. *J. Phys. A: Math. Gen.*, 35:9843–9856, 2002.

[44] R. Montgomery. The isoholonomic problem and some applications. *Comm. Math. Phys.*, 128:565–592, 1990.

[45] R. Montgomery. Abnormal minimizers. *SIAM J. Control Optim.*, 32:1605–1620, 1994.

[46] R. Montgomery. *A tour of Subriemannian geometries, their geodesics and applications*, volume 91 of *Mathematical Surveys and Monographs.* American mathematical society, 2002.

[47] L.S. Pontryagin, V.G. Boltyanskii, R.V. Gamklelidze, and E.F. Mishchenko. *The Mathematical Theory of Optimal Processes.* Wiley, Interscience, 1962.

[48] P. Popescu. On the geometry of relative tangent spaces. *Rev. roum. math. pures appl.*, 37:727–733, 1992.

[49] D.J. Saunders. *The geometry of jet bundles.* Cambridge University Press, Cambrige, 1979.

[50] R.S. Strichartz. Sub-Riemannian geometry. *J. Diff. Geom.*, 24:221–263, 1986.

[51] R.S. Strichartz. Corrections to “sub-Riemannian geometry”. *J. Diff. Geom.*, 30:595–596, 1989.

[52] H.J. Sussmann. Orbits of families of vector fields and integrability of distributions. *Trans. Amer. Math. Soc.*, 180:171–188, 1973.

[53] H.J. Sussmann. A cornucopia of four-dimensional abnormal subriemannian minimizers. In A. Bellaïche and J.-J. Risler, editors, *Sub-Riemannian geometry*, volume 144 of *Progr. Math.*, pages 341–364. Birkhäuser, Basel, 1996.

[54] H.J. Sussmann. An introduction to the coordinate-free maximum principle. In B. Jakubczyk and W. Respondek, editors, *Geometry of Feedback and Optimal Control*, pages 463–557. Marcel Dekker, New York, 1997.

[55] I. Vaisman. On the geometric quantization of poisson manifolds. *J. Math. Phys.*, 32:3339–3345, 1991.

[56] I. Vaisman. *Lectures on the Geometry of Poisson Manifolds*, volume 118 of *Progress in Math.* Birkhäuser, Basel, 1994.

[57] J. Vilms. Connections on tangent bundles. *J. Diff. Geom.*, 1:235–243, 1967.

[58] A. Weinstein. Lagrangian mechanics and groupoids. In W.F. Shadwick, P.S. Krishnaprasad, and T.S. Ratiu, editors, *Mechanics Day*, volume 7 of *Fields Institute Communications*, pages 207–232. AMS, Providence, Rhode Island, 1995.

[59] Y.C. Wong. Linear connections and quasi connections on a differentiable manifold. *Tôhuku Math. J*, 14:48–63, 1962.

Index

Samenvatting

De resultaten, voorgesteld in deze verhandeling, vinden hun oorsprong in recent werk van R.L. Fernandes [11, 14, 15] omtrent integrabiliteit van Lie algebroids en de constructie van nieuwe karakteristieke klassen die kunnen geassocieerd worden met geometrische structuur van een Lie algebroid. Hiertoe werd een veralgemening van het concept connectie ingevoerd, waarvan de essentie kan teruggevonden worden in eerdere werken van o.a. I. Vaisman [56], Y.C. Wong [59] en F. Kamber en P. Tondeur [21] over, respectievelijk, contravariante connecties, pseudo-connecties en partiële connecties.

Het belang van het creëren van een algemeen geometrisch model voor een connectie waartoe de hierboven vermelde verschillende types van connecties behoren, ligt in het feit dat we hiermee een nieuw licht kunnen werpen op differentiaalmeetkundige structuren die tot nu toe niet werden bestudeerd vanuit het standpunt van connectietheorie. Enkele toepassingsgebieden van de veralgemeende connectietheorie die we zullen ontwikkelen, zijn o.a. controletheorie (Hoofdstuk III) en daarmee geassocieerd sub-Riemannmeetkunde (Hoofdstuk IV), en niet-holonome mechanica (Hoofdstuk V).

In Hoofdstuk I bestuderen we de differentiaalmeetkundige structuren waarop we, in Hoofdstuk II, het begrip van een veralgemeende connectie zullen definiëren. Deze structuren worden *verankerde bundels* genoemd en bestaan uit een vezelbundel $\nu : N \to M$ over de basisvariëteit M, ook wel configuratieruimte genoemd, en een bundelafbeelding $\rho : N \to TM$, gevezeld over de identiteit op M. De idee die aan de grondslag ligt om dergelijke structuren te onderzoeken, is het feit dat de dynamische systemen die voorkomen in controletheorie, sub-Riemannmeetkunde en niet-holonome mechanica, kunnen geformuleerd worden op verankerde bundels. De integraalkrommen van dergelijke dynamische systemen kunnen geassocieerd worden met speciale krommen in N die *toelaatbare krommen* worden genoemd en die als volgt worden gedefinieerd. Een kromme $c : I = [a, b] \to N$ is toelaatbaar als ze voldoet aan de relatie

$$\left.\frac{d}{dt}\right|_t \nu(c(t)) = \rho(c(t)).$$

De geprojecteerde kromme $\widetilde{c} = \nu \circ c$ in de configuratieruimte M wordt de basiskromme van de toelaatbare kromme c genoemd en we zeggen dat c het punt $x = \widetilde{c}(a) \in M$ naar het punt $y = \widetilde{c}(b) \in M$ brengt, of nog dat het punt *y bereikbaar is vanuit x* onder de toelaatbare kromme c.

In Hoofdstuk I onderzoeken we eigenschappen van deze toelaatbare krommen. In het bijzonder zullen we nagaan hoe de verzameling van alle *bereikbare punten vanuit een gegeven vast punt x* in M kan gekarakteriseerd worden. Het blijkt dat deze verzameling van bereikbare punten vanuit x bevat is in de maximale integraalvariëteit door x (in het Engels de 'leaf' door x genoemd) van de kleinste integreerbare veralgemeende distributie op M die gegenereerd wordt door alle raakvectoren van de vorm $\rho(s)$, met $s \in N$. In het bijzondere geval dat de bundel ν de structuur heeft van een vectorbundel en dat ρ een lineair morfisme is, blijkt dat elk punt in de 'leaf' door x bereikbaar is vanuit x. Deze resultaten zijn van essentieel belang omdat zij toelaten een notie van holonomie te ontwikkelen voor veralgemeende connecties. Verder zullen we *toelaatbare lussen* definiëren en tonen we aan dat deze krommen een deelgroep bepalen van de eerste fundamentaalgroep van M.

Voordat we de veralgemeende connecties bespreken, herhalen we eerst bondig wat klassiek onder een connectie wordt verstaan. Er bestaan vele equivalente definities van het begrip connectie op een vectorbundel $\pi : E \to M$. Eén ervan karakteriseert een connectie op π als een afbeelding h die raakvectoren in TM '*lift*' naar raakvectoren in TE. Meer specifiek is h een afbeelding van de bundel $E \times_M TM$ (het vezelproduct van E en TM) naar TE die aan de volgende eigenschappen voldoet: $\tau_E(h(e,\cdot)) = e$ voor alle $e \in E$, h is linear in zijn tweede argument, en h commuteert met $T\pi$, i.e. als $(e,v) \in E \times_M TM$, dan is $T\pi(h(e,v)) = v$. Het beeld van de lift h bepaalt een distributie $H\pi$ op E, complementair aan de distributie van verticale raakvectoren $V\pi = \ker T\pi$ ($H\pi$ wordt daarom ook wel de *horizontale distributie* genoemd). Een connectie op een vectorbundel $\pi : E \to M$ laat toe vezels van π over verschillende punten met elkaar in verband te brengen, op voorwaarde dat een kromme in M is gegeven die deze punten met elkaar verbindt. Wiskundig wordt dit verband uitgedrukt met behulp van een *parallel transportoperator* die geassocieerd is aan de gegeven kromme.

Veronderstel nu dat een verankerde bundel $\nu : N \to M$ gegeven is met ankerafbeelding $\rho : N \to TM$. De veralgemening van het begrip connectie, ingevoerd in Hoofdstuk II, bestaat erin dat we nu elementen van N naar raakvectoren in TE liften. Meer specifiek: een *veralgemeende connectie* over de afbeelding ρ wordt gedefinieerd als een afbeelding h van $E \times_M N$

naar TE, zodat voldaan is aan $\tau_E(h(e,s)) = e$ en $T\pi(h(e,s)) = \rho(s)$ voor willekeurige (e,s) in $E \times_M N$.

De distributie Q op E, opgespannen door alle raakvectoren van de vorm $h(e,s)$ met $(e,s) \in E \times_M N$, is niet noodzakelijk complementair aan de distributie $V\pi$ van verticale raakvectoren. In het algemeen is $Q \cap V\pi \neq \{0\}$ en geldt er dat $T\pi(Q)$ niet samenvalt met de volledige raakbundel van M. Deze eigenschappen zorgen ervoor dat we hier niet kunnen spreken over Q als een 'horizontale' distributie. Omwille van het feit dat de distributie Q niet complementair is aan $V\pi$, is een zinvolle veralgemening van de notie van kromming van een veralgemeende connectie niet haalbaar. Echter, de noties van parallel transport en holonomie kunnen wel worden veralgemeend, en dit is het belangrijkste resultaat uit Hoofdstuk II. We merken op dat de transportoperator van een veralgemeende connectie niet langer gedefinieerd is langs een willekeurige kromme in M, maar slechts langs toelaatbare krommen in de verankerde bundel $\nu : N \to M$. De holonomiegroepen worden dan gegenereerd door toelaatbare lussen. Omdat de toelaatbare krommen (of lussen) bevat zijn in de leaves van de foliatie die geïnduceerd wordt door de ankerafbeelding, zullen we spreken over 'leafwise holonomy'.

Zoals reeds eerder vermeld, worden in de Hoofdstukken III, IV en V enkele mogelijke toepassingsgebieden van de veralgemeende connectietheorie die we nu kort overlopen.

Controletheorie behelst de studie van dynamische systemen die kunnen voorgesteld worden met behulp van een stelsel differentiaalvergelijkingen van de vorm

$$\dot{q}^i(t) = \gamma^i(q^1(t), \ldots, q^n(t), u^1(t), \ldots, u^k(t)),$$

waarbij $q(t) = (q^1(t), \ldots, q^n(t))$ en $u(t) = (u^1(t), \ldots, u^k(t))$ krommen zijn in, respectievelijk, $\mathbb{R}^n$ en $\mathbb{R}^k$, en waarbij γ een gladde afbeelding is van $\mathbb{R}^n \times \mathbb{R}^k$ naar $\mathbb{R}^n$. Dergelijke dynamische systemen kunnen teruggevonden worden bij de modellering van fysische systemen die voorkomen bij tal van technologische toepassingen. De ruimte $\mathbb{R}^n$ wordt opgevat als de configuratieruimte van het systeem en de kromme $u(t)$ stelt een externe input voor in het systeem. De afbeelding γ geeft aan hoe de input $u(t)$ de configuratie van het systeem beïnvloedt. Meer nog, de input (of *controle*) $u(t)$ zal de evolutie van het systeem volledig bepalen, op voorwaarde dat een beginconfiguratie is gegeven (dit volgt uit de uniciteit van oplossingen van differentiaalvergelijkingen).

Veronderstel nu dat de afbeelding γ de lokale voorstelling is van de ankerafbeelding ρ van een verankerde bundel $\nu : N \to M$ in een aangepast coordinatensysteem (q^i, u^a) op N. Het is eenvoudig in te zien dat het koppel $(q(t), u(t))$ dat voldoet aan $\dot{q}^i(t) = \gamma^i(q(t), u(t))$, een toelaatbare kromme is. Dit legt het verband tussen de theorie omtrent verankerde bundels en controlesystemen.

We wensen nu te zoeken naar een toelaatbare kromme $(q(t), u(t))$ die twee vooropgestelde punten q_0 en q_1 in $\mathbb{R}^n$ met elkaar verbindt, zodanig dat deze toelaatbare kromme de 'beste keuze' is met betrekking tot een vooropgestelde grootheid L, die soms ook de *kostfunctie* wordt genoemd. Hiermee bedoelen we dat deze toelaatbare kromme *optimaal* is ten opzichte van andere toelaatbare krommen die de punten q_0 en q_1 verbinden, in de zin dat ze een extremum oplevert van de functionaal

$$\int_{t_0}^{t_1} L(q(t), u(t))dt.$$

De kostfunctie kan bijvoorbeeld de 'tijd' zijn. Men zoekt dan naar die toelaatbare kromme die beide punten met elkaar verbindt in een zo kort mogelijke tijdspanne. Een andere mogelijkheid is de 'energie', men wil een toelaatbare kromme vinden die zo weinig mogelijke 'energie' verbruikt. De zoektocht naar deze *optimale* toelaatbare kromme werd gedeeltelijk opgelost door het *maximumprincipe*. Het maximumprincipe geeft nodige voorwaarden opdat een toelaatbare kromme optimaal zou zijn. Toelaatbare krommen die voldoen aan deze nodige voorwaarden worden *extremale* krommen genoemd. Onder bepaalde voorwaarden gebeurt het dat een extremale kromme onafhankelijk is van de beschouwde kostfunctie. Deze vaststelling leidde tot het invoeren van de benaming *abnormale extremalen*. In het verleden werd het bestaan van deze krommen vaak over het hoofd gezien. R. Montgomery toonde in 1994 aan dat er abnormale extremale krommen bestaan die optimaal zijn.

In Hoofdstuk III leiden we een differentiaalmeetkundige versie af van het maximumprincipe. Hiertoe voeren we een notie van variatie van een toelaatbare kromme in. De raakvectoren aan dergelijke variaties genereren een kegel in de raakruimte aan de configuratieruimte in het eindpunt q_1 van de gegeven toelaatbare kromme. We tonen aan dat deze kegel in een omgeving van het eindpunt de verzameling van bereikbare punten vanuit het beginpunt q_0 genereert. Dit laat toe om nodige voorwaarden op te stellen voor optimale toelaatbare krommen. Tevens onderzoeken we het geval waarbij de eindpunten van een optimale kromme kunnen variëren op een vooropgegeven begin- en eindoppervlak.

De rol die veralgemeende connecties spelen in deze theorie, bestaat erin dat de *variatiekegel* wordt gegenereerd door de transportoperator van een canonisch bepaalde veralgemeende connectie. Nodige en voldoende voorwaarden opdat een toelaatbare kromme een abnormale extremale kromme is, kunnen worden geformuleerd in termen van deze variatiekegel. Deze karakterisering van abnormale krommen kan samengevat worden als volgt: als een toelaatbare kromme abnormaal is, dan, en slechts dan, is de variatiekegel niet gelijk aan de volledige raakruimte in het eindpunt van de kromme. Met andere woorden: als de te onderzoeken toelaatbare kromme niet voldoende variaties toelaat opdat alle punten in een open omgeving van het eindpunt bereikbaar zouden zijn, dan, en slechts dan, is de toelaatbare kromme abnormaal. Het meetkundig beeld achter deze karakterisering is intuïtief duidelijk. Het blijkt echter dat analytische berekeningen om deze voorwaarden te verifiëren meestal bijzonder omslachtig zijn. Daarom hebben we, in het geval dat ρ een lineaire afbeelding is, nieuwe voldoende voorwaarden geformuleerd voor abnormale extremale krommen, die veel gemakkelijker hanteerbaar zijn. Op het einde van Hoofdstuk III passen we de differentiaalmeetkundige versie van het maximumprincipe die we hebben opgesteld toe op enkele gekende voorbeelden uit de variatierekening: o.a. Lagrangiaanse systemen (al dan niet met niet-holonome bindingen) en Lagrangiaanse systemen op (affiene) Lie algebroids.

In Hoofdstuk IV bestuderen we sub-Riemannmeetkunde. Sub-Riemannmeetkunde onderscheidt zich van Riemannmeetkunde in het feit dat de metriek slechts gedefinieerd is op een deelbundel van de raakbundel van de onderliggende variëteit. De lengte van een kromme is dan ook enkel gedefinieerd voor een bepaalde klasse van krommen, namelijk die krommen die overal raken aan die deelbundel. Het is algemeen bekend dat deze 'bevoorrechte' krommen kunnen opgevat worden als toelaatbare krommen voor een geassocieerd controlesysteem. Een lengte-minimaliserende kromme in een sub-Riemannstructuur komt overeen met een optimale toelaatbare kromme in het geassocieerd controlesysteem. Dit laat toe om de bekomen resultaten omtrent het maximumprincipe te vertalen naar nodige voorwaarden voor lengte-minimaliserende krommen van een sub-Riemannstructuur. Het blijkt dat we deze nodige voorwaarden op een elegante wijze kunnen herformuleren aan de hand van veralgemeende connecties. De resultaten uit Hoofdstuk III omtrent de karakterisering van abnormale extremalen worden getoetst aan de hand van een tweetal concrete voorbeelden.

We eindigen in Hoofdstuk V met een toepassing van veralgemeende connecties op de reductie van mechanische systemen met niet-holonome bindingen waarvoor de bewegingsvergelijkingen en de bindingen invariant zijn onder de actie van een symmetriegroep. Onze behandeling verschilt van andere behandelingen van dit vraagstuk in het feit dat we geen bijkomende voorwaarden opleggen aan de niet-holonome bindingen, behalve de voorwaarde dat ze een reguliere invariante niet-integreerbare distributie induceren op de configuratieruimte. Om de gedachten te vestigen, illustreren we onze theorie aan de hand van een concreet voorbeeld, namelijk het 'Snakeboard'. Dit is een skateboard waarvan de wielen vrij kunnen roteren om een verticale as. We veronderstellen dat de wielen rollen zonder glijden. Een interessant fenomeen hierbij is dat de bestuurder een slangachtige beweging kan genereren zonder zich van de grond te moeten afduwen.