12

An introduction to the coordinate-free Maximum Principle

H. J. Sussmann [1]

Department of Mathematics, Rutgers University, New Brunswick, New Jersey 08903

1. Introduction

The purpose of this note is to present an introduction to a version of the finite-dimensional Pontryagin Maximum Principle that stresses its geometric, coordinate-free formulation on manifolds, and at the same time takes into account some recent developments that extend and generalize more classical versions.

The Maximum Principle (abbr. MP) is a "principle" in the true sense of the word. That is, it is a somewhat vague statement that can be made precise in more than one way, in various situations, under different technical assumptions, and with similar but not identical conclusions. So the MP gives rise to a cluster of results about several closely related problems (cf. [1], [2], [3], [4], [7], [8], [11], [13] for various versions). While all these results may look nearly identical to the uninitiated, they differ significantly in the precise technical details of their formulation. These differences—e.g. on

[1] Partially supported by NSF grant DMS92-02554.

technical assumptions on Lipschitz conditions, or on Carathéodory-type requirements on integral bounds for derivatives or Lipschitz constants of time-dependent functions or vector fields—may seem unintuitive to the differential geometer, who is accustomed to dealing with very smooth objects, but are often the crux of the matter for the control theorist. For this reason, there is the danger that a truly general version of the Maximum Principle may appear at first sight to involve intricate lists of unmotivated technical conditions, making it hard for the reader to appreciate the statement's rather natural geometric content.

Our objective here is to provide a formulation that particularly stresses the differential geometric aspect and brings to the fore the crucial role of geometric objects such as Lie brackets, connections along curves, symplectic structures and Poisson brackets, without making unduly restrictive technical assumptions. We insist on providing coordinate-free formulations that make sense on manifolds *directly*, i.e. without having to choose particular coordinate charts, or bases of sections of vector bundles. We will emphasize three features of our version, namely,

F1. the fact that a control system is defined to be a general family of vector fields,

F2. the use of minimal technical assumptions, that are in particular strictly weaker than those of previous versions such as the nonsmooth "Maximum Principle under minimal hypotheses" of [3],

F3. the intrinsic coordinate-free formulation on manifolds.

Our presentation of the MP is preceded by a review of the classical Euler–Lagrange equations, stated so as to make the transition from the classical variational problem to the control problem straightforward. This should facilitate comprehension to those readers who are more familiar with the classical Calculus of Variations (abbr. C. of V.) than with modern Optimal Control Theory. At the same time, our path from the C. of V. to Optimal Control shows the latter to be a very far-reaching generalization of the former: the kinematics (i.e. the set of curves) of a C. of V. problem is trivial, since it consists of the set of all curves, and is generated by the family of all vector fields, so that all the interesting structure of a C. of V. problem arises from its Lagrangian. Optimal Control problems, on the other hand, can have a much more nontrivial kinematics $\mathcal{K}$, generated by fairly general collections $\mathcal{F}$ of vector fields, while the Lagrangian L can still be as complicated as in the C. of V. case. It follows that in the study of these problems a much larger set of issues arises, including many that have no C. of V. counterpart, such as for example various "controllability" questions, i.e. questions about the properties of the set $\mathcal{R}(p)$ of all points reachable from a given point p by means of curves in $\mathcal{K}$, such as whether $\mathcal{R}(p)$ is the whole space, whether it is a neighborhood of p, whether it is a submanifold, whether it has a smooth or piecewise smooth boundary, or whether it is subanalytic. To answer these questions, one has to bring into play a much richer geometric structure, and in particular the Lie-algebraic relations among the vector fields in $\mathcal{F}$. Moreover, those issues that do have a C. of V. counterpart—e.g. the understanding of the structure of optimal trajectories—require an even richer structure, involving the interplay between the Lie-algebraic properties of $\mathcal{F}$ and the Lagrangian L.

Because of limitations of space, many important issues will not be discussed. In particular, (i) *we will only deal with the MP for problems without state space constraints*. Moreover, (ii) we will not touch upon another important aspect of the new version of the MP, namely, the inclusion of *high-order necessary conditions*, and (iii) we will not discuss the Maximum Principle for differential inclusions. Finally, (iv) we will omit the proof. (A rather complete outline of the proof for the case without state space constraints is given in [22] and [23]. These two papers only deal with the case when the state space is an open subset of a Euclidean space $\mathbb{R}^n$, but they allow systems with with jumps. The manifold case is then a trivial corollary, derived by applying the Euclidean space result to a system where the jumps are coordinate changes.)

Other than the four limitations listed above, our statement will be general, our presentation self-contained, and our technical assumptions minimal. The results will be stated in the general nonsmooth setting under which they were proved in [22] and [23]. In this setting, the reference arc

ξ_* is generated by a control η_* that gives rise to a time-varying vector field $(x,t) \to F_{\eta_*}(x,t)$ that is only assumed to be locally Lipschitz with respect to the state variable x (with an appropriate integral bound of Carathéodory type for the Lipschitz constant as a function of t). On the other hand, for extra clarity, most of our discussion will be confined to the smooth case, in which F_{η_*} is of class C^1 with respect to x.

We incorporate Feature F1 into our version by following the ideas of B. Kaskosz and S. Lojasiewicz Jr. on "generalized control systems" (cf. [6]). We define a control system to be a *parametrized collection of time-varying vector fields*, i.e. a family $F = \{F_\eta\}_{\eta\in\mathcal{U}}$ of time-dependent vector fields—called *control vector fields*[2] , abbr. CVF's—$(x,t) \to F_\eta(x,t)$ on a manifold M. This is more general, and in our view also more natural, than the more traditional approach of considering systems $\dot{x} = f(x,t,u)$, $u \in U$ (or, more generally, systems $\dot{x} = f(x,t,u)$, $u \in U(t)$), where U (or each of the sets $U(t)$) is usually a subset of a Euclidean space, and "open-loop controls" are functions $t \to \eta(t) \in U(t)$ required to satisfy some measurability conditions and some bounds. For us, $\mathcal{U}$ will just be an abstract set, whose elements η will still be called "open-loop controls." Naturally, measurability properties and bounds are still needed, but they now become conditions of the Carathéodory type on the control vector fields F_η, rather than requirements on the parameter values η. *Trajectories* are then locally absolutely continuous curves $t \to \xi(t)$ in M that satisfy the O.D.E. $\dot{\xi}(t) = F_\eta(\xi(t),t)$ for some η. The MP deals with *pairs* (ξ,η), where ξ is an arc and η is a control that generates ξ. (In this paper, the word "arc" always means "absolutely continuous curve defined on a compact interval.") We call such pairs "controlled arcs."[3]

Feature F2 is included by following through on a deep observation due

[2] The reader is warned that this terminology is not standard, since the expression "control vector field" has been used by other authors (for example, F. Albrecht in [1]) with a different meaning.

[3] In our view this conveys the intended meaning much better than expressions such as "admissible pair," that are often encountered in the literature. Naturally, there is no need to distinguish between arcs and controlled arcs for systems such as those of the classical Calculus of Variations, in which a trajectory cannot be generated by more than one control.

to Lojasiewicz[4], who pointed out that the vector fields other than F_{η_*} can be assumed to be just continuous with respect to the state. Thanks to this extra generality, the approach presented here turns out to settle the optimality question even in some situations where other nonsmooth versions of the Maximum Principle fail, as remarked by Kaskosz and Lojasiewicz, one of whose examples is reviewed below (cf. Example 15.4).

The incorporation of Feature F3 is the most important goal of this paper. In order to provide an intrinsic coordinate-free formulation of the MP, we first present a detailed discussion of the precise type of differential-geometric structure that is required for this purpose. As is well known, the conclusion of the MP involves three basic ingredients, namely, (I1) the adjoint equation (abbr. AE), (I2) Hamiltonian minimization, and (I3) transversality conditions. The last two are manifestly intrinsic, but the AE for a controlled arc $\gamma = (\xi, \eta)$ would appear to involve some choices, since it has the form $\dot{\zeta} + \zeta.\frac{\partial F_\eta}{\partial x} = 0$, and the intrinsic meaning of $\frac{\partial F_\eta}{\partial x}$ —the Jacobian with respect to x of the CVF $x \to F_\eta(x,t)$—and of the derivative $\dot{\zeta}$ require further clarification. (It is clear that the correct differential-geometric interpretation of the "adjoint vector" ζ is as a field of covectors along the curve ξ, but then $\dot{\zeta}$ is not well defined, since one does not know how to differentiate fields of vectors or covectors along curves.) The proper answer turns out to be—if F_η is sufficiently smooth—that the expression $\dot{\zeta}+\zeta.\frac{\partial F_\eta}{\partial x}$ must not be broken up into the sum of two separately meaningful terms, but has to be regarded instead as the covariant derivative of ζ relative to a connection $\mathcal{L}^\gamma$ naturally associated with γ. This is explained in detail in Section 4.

The paper is organized as follows. In Section 2 we review the Euler–Lagrange equations and show how to generalize them to get the Control Theory MP. In Section 3 we begin our formal discussion by introducing a number of general definitions, notational conventions and abbreviations that will be used throughout the paper. In Section 4 we discuss the theory of connections along curves and study the adjoint equation along a controlled curve, i.e. Item (I1) in our list of three main ingredients of the

[4] Personal communication.

MP. Section 5 contains the definition of control systems, Lagrangians and reachable sets. Section 6 studies Hamiltonian minimization—i.e. Item (I2) in the list—and introduces the important property of separability, which plays a crucial role in the transition from "weak" to "strong" minimization. Section 7 studies the third item in our list, namely, transversality, and defines the concept of an approximating cone. Section 8 then presents the statement of various forms of the MP, starting with the basic version of the MP as a theorem about set separation, and deducing from it other versions for controllability and optimal control. In Sections 9 and 10 we specialize to systems of the form $\dot{x} = f(x,u)$ and $\dot{x} = f(x,t,u)$ and derive the corresponding versions of the MP, showing in particular that in these cases the strong form of the minimization condition can in fact be used instead of the weak one. Sections 11 to 15 contain illustrations of the application of the MP to various differential-geometric and optimal control questions. In view of the geometric orientation of this paper, we choose as our first example the well known equation of the Riemannian geodesics, and show in Section 11 how to derive it in a simple way from the MP, using only the characterization of the Riemannian connection as the unique torsion-free connection for which the covariant derivative of the metric tensor vanishes, and without having to take coordinates or write formulas for the Christoffel symbols. (It turns out that the MP gives a necessary condition that can be stated in terms of an arbitrary connection $\mathcal{L}$ along the given curve. When $\mathcal{L}$ happens to be the pullback of the Levi-Civita connection, then the condition takes a simpler form, which is precisely the geodesic equation.) In Section 12 we look at a problem lying at the opposite extreme of the C. of V. -Control continuum, namely minimum-time single-input control-affine optimal control. (This means that $L \equiv 1$, so L is trivial, and $\mathcal{F}$ essentially consists of two vector fields, two being the minimum number of vector fields needed in order to have nontrivial control.) We show how the application of the MP—using the geometric form of the AE—directly leads to relationships between properties of the optimal trajectories and Lie algebraic relations among the two vector fields defining the kinematics. In particular, Section 12 provides a very brief introduction to some key concepts and facts in the general regularity theory of optimal trajectories. In Section 13 we let the pendulum swing back to the geometric side, and look at translation-invariant control problems on Lie groups. In particular, we give an example of a problem in the rotation group $SO(3)$ which is not Riemannian —since there are only two independent directions of motion at each point— but where the MP implies smoothness of the optimal arcs. Section 14 discusses sub-Riemannian minimization, and shows a coordinate-free, basis-free derivation of the necessary conditions for optimality, highlighting the importance of the so-called "abnormal extremals." Section 15 presents a number of examples pertaining to various technical issues discussed in the text. These examples show: (i) that the conditions of the MP are not local (Example 15.1); (ii) that our version of the MP, by allowing the vector fields other than the reference vector field to be just continuous, gives in some cases better results than other nonsmooth versions (Example 15.2); (iii) that the conditions of the MP can be satisfied for a controlled trajectory (ξ, η) but be violated for another controlled trajectory $(\xi, \tilde{\eta})$ with the same ξ (Example 15.3), (iv) that by enlarging the original control system—i.e. by adding more vector fields without changing the class of trajectories—one can sometimes rule out arcs that satisfy the conditions of the MP for the original system (Example 15.4), (v) that the requirement that the reference vector field be locally integrably Lipschitz cannot be weakened (Example 15.5). Finally, in Section 16, we briefly point out that the MP is only the beginning of a much richer theory, involving "high-order conditions." We present a simple example of a problem for which the MP in the form of this paper does not suffice, while the addition of a high-order condition involving a Lie bracket leads to a complete solution.

2. The Euler–Lagrange equations and the maximum principle

Let $\Omega \subseteq \mathbb{R}^n$ be open, let $a, b \in \mathbb{R}$, $a < b$, and let $\Omega \times \mathbb{R}^n \times [a, b] \ni (x, u, t) \to L(x, u, t) \in \mathbb{R}$ be a function that satisfies technical conditions to be made precise later. In its simplest version, the classical Calculus of

Variations deals with the following minimization problem: given $\bar{x}$ and $\hat{x}$ in Ω, characterize those arcs (i.e. absolutely continuous maps) $\xi_* : [a,b] \to \Omega$ that minimize the integral $J^L(\xi) = \int_a^b L(\xi(t),\dot{\xi}(t),t)\,dt$ among all the arcs $\xi : [a,b] \to \Omega$ such that $\xi(a) = \overline{x}$ and $\xi(b) = \widehat{x}$.

Under suitable technical conditions on L and ξ_*, the implication

(ELI) *a minimizing arc necessarily satisfies* (EL)

is true, where (EL) stands for the *Euler–Lagrange equations*, usually written in the form $\frac{d}{dt}\big(\frac{\partial L}{\partial \dot{x}}\big) = \frac{\partial L}{\partial x}$. For sufficiently smooth curves and Lagrangians (e.g. for ξ_*, $L \in C^2$) the meaning of this equation is clear. Under more general conditions (for example, if ξ_* is Lipschitz but not C^1), some interpretation is required. Define a vector-valued function $\widetilde{\zeta}(\cdot) : [a,b] \to \mathbb{R}^n$ (the *momentum*) by $\widetilde{\zeta}(t) = -\frac{\partial L}{\partial u}(\xi_*(t),\dot{\xi}_*(t),t)$, (i.e. $\widetilde{\zeta}_i(t) = -\frac{\partial L}{\partial u^i}(\xi_*(t),\dot{\xi}_*(t),t)$ for $i = 1,\ldots,n$). Then the Euler-Lagrange equations say that

(EL) $\widetilde{\zeta}$ *is equal a.e. to an absolutely continuous—obviously unique—vector function* ζ, *that satisfies* $\dot{\zeta}_i(t) + \frac{\partial L}{\partial x^i}(\xi_*(t),\dot{\xi}_*(t),t) = 0$ for a.e. $t \in [a,b]$, $i = 1,\ldots,n$.

Then **(ELI)** becomes a rigorous theorem if, for example, (i) L is a function of class C^1, and (ii) ξ_* is Lipschitz. These are not, however, the most general hypotheses that make **(ELI)** true. Much weaker assumptions still suffice, although they are more complicated to state, and the proof of the implication **(ELI)** becomes more delicate. For example, it suffices to require —if we use $\mathcal{T}(\xi_*,\delta)$ to denote the set $\mathcal{T}(\xi_*,\delta) = \{(x,t) : a \le t \le b,\ \|x - \xi_*(t)\| \le \delta\}$—that the following conditions hold on $\mathcal{T}(\xi_*,\delta)$ for some $\delta > 0$:

EL1 $L(x,u,t)$ is continuous with respect to (x,u) for each fixed t,

EL2 $L(x,u,t)$ is Lebesgue measurable with respect to t for each fixed (x,u),

EL3 $L(x,u,t)$ is continuously differentiable with respect to x for each fixed (u,t), and a "Carathéodory type" estimate $\big\|\frac{\partial L}{\partial x}(x,\dot{\xi}_*(t),t)\big\| \le \varphi(t)$ is satisfied for all (x,t) in $\mathcal{T}(\xi_*,\delta)$, where $\varphi : [a,b] \to \mathbb{R}$ is an integrable function,

EL4 $L(x,u,t)$ is everywhere differentiable with respect to u for each fixed (x,t),

MIN $J^L(\xi_*) \leq J^L(\xi)$ for all absolutely continuous curves $\xi : [a,b] \to \Omega$ such that $\xi(a) = \xi_*(a)$, $\xi(b) = \xi_*(b)$, and $(\xi(t),t) \in \mathcal{T}(\xi_*,\delta)$ for all $t \in [a,b]$.

The Euler–Lagrange equations can also be written in "Hamiltonian form" by introducing the *Hamiltonian* $\mathcal{H}^L : \Omega \times \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}$ given by

$$\mathcal{H}^L(x,z,t,u) = \langle z,u\rangle + L(x,u,t). \tag{2.1}$$

If we plug in $x = \xi_*(t)$, $u = \dot{\xi}_*(t)$, $z = \zeta(t)$, and let $\Gamma(t) = (\xi_*(t),\zeta(t),t,\dot{\xi}_*(t))$, then the O.D.E. $\dot{\zeta}(t) = -\frac{\partial L}{\partial u}(\xi_*(t),\dot{\xi}_*(t),t)$—which holds a.e. on $[a,b]$—just says that, along Γ, the identity

$$\frac{\partial \mathcal{H}^L}{\partial u^i} = 0 \tag{2.2}$$

holds for a.e. $t \in [a,b]$. Clearly, $\dot{x} = u$ also holds a.e., and this, together with the Euler-Lagrange equations $\dot{\zeta}_i + \frac{\partial L}{\partial x^i} = 0$, says that *Hamilton's equations*

$$\dot{x}^i = \frac{\partial \mathcal{H}^L}{\partial z_i}, \qquad \dot{z}_i = -\frac{\partial \mathcal{H}^L}{\partial x^i}. \tag{2.3}$$

hold for a.e. $t \in [a,b]$. The Hamiltonian version of the classical necessary conditions then becomes:

Theorem 2.1. *Let Ω be an open subset of $\mathbb{R}^n$, and let $\xi_* : [a,b] \to \Omega$ be absolutely continuous. Let $L : \Omega \times \mathbb{R}^n \times [a,b] \to \mathbb{R}$ be a function such that Conditions* **EL1**, **EL2**, **EL3**, **EL4** *and* **MIN** *are satisfied on some set $\mathcal{T}(\xi_*,\delta)$. Then there is an absolutely continuous function $\zeta : [a,b] \to \mathbb{R}^n$ such that, if we define $\mathcal{H}^L$ by* (2.1), *then* (2.2) *and* (2.3) *hold for a.e. $t \in [a,b]$, $i = 1,\ldots,n$, along the curve Γ given by $t \to (x,z,t,u) = (\xi_*(t),\zeta(t),t,\dot{\xi}_*(t))$.*

This can be strengthened by observing that (2.2) is a consequence of the following assertion

$$\mathcal{H}^L(\xi_*(t),\zeta(t),t,\dot{\xi}_*(t)) = \min\{\mathcal{H}^L(\xi(t)_*,\zeta(t),t,v) : v \in \mathbb{R}^n\} \qquad \text{for a.e. } t \in [a,b], \tag{2.4}$$

known as the *Weierstrass side condition*. It turns out that (2.4) is also necessary for optimality. So we may reformulate our necessary conditions using (2.4) in place of (2.2). This has two deep consequences. First, the new result is better even for a Lagrangian L and arc ξ_* that satisfy the technical conditions of Theorem 2.1, since (2.4) is in general strictly stronger than (2.2). A second and even more important consequence of the new formulation is that, since (2.2) was the only place where our necessary conditions involved the partial derivatives $\frac{\partial L}{\partial u_i}$, these derivatives have now completely disappeared. It is then to be expected that the differentiability condition **EL4** might no longer be needed. Indeed, one can prove:

Theorem 2.2. *Let Ω be an open subset of $\mathbb{R}^n$, and let $\xi_* : [a,b] \to \Omega$ be absolutely continuous. Let $L : \Omega \times \mathbb{R}^n \times [a,b] \to \mathbb{R}$ be a function such that* **EL1**, **EL2**, **EL3** *and* **MIN** *hold on some set $\mathcal{T}(\xi_*, \delta)$. Then there is an absolutely continuous function $\zeta : [a,b] \to \mathbb{R}^n$ such that, if we define $\mathcal{H}^L$ by (2.1), then (2.3) and (2.4) hold for a.e. $t \in [a,b]$, $i = 1, \ldots, n$, along the curve Γ given by $t \to (x,z,t,u) = (\xi_*(t), \zeta(t), t, \dot{\xi}_*(t))$.*

Theorem 2.2 is a particular case of the general Maximum Principle, stated below.

What we have done so far can be described in the following alternative way: we started with the dynamical equation $\dot{x} = \eta(t)$, in which x belongs to an open subset Ω of $\mathbb{R}^n$ and the "control" η is an arbitrary integrable $\mathbb{R}^n$-valued function, and with a function $(x,u,t) \to L(x,u,t)$, the "Lagrangian." For each control function η, let $L_\eta(x,t) = L(x,\eta(t),t)$. Letting $\mathcal{U}$ be the set of all control functions η, we can define a family $\mathcal{H}^L = \{H^L_\eta\}_{\eta\in\mathcal{U}}$ of time-varying *Hamiltonian functions*, parametrized by $\mathcal{U}$, by letting

$$\mathcal{H}^L_\eta(x,z,t) = \langle z, \eta(t)\rangle + L_\eta(x,t). \tag{2.5}$$

Associated with each Hamiltonian function $H : \Omega \times \mathbb{R}^n \times I \to \mathbb{R}$ —where I is a subinterval of $\mathbb{R}$—there is a *time-varying vector field* $\vec{H}$ on $\Omega \times \mathbb{R}^n$, whose components are $\left(\frac{\partial H}{\partial z_1}, \ldots, \frac{\partial H}{\partial z_n}, -\frac{\partial H}{\partial x^1}, \ldots, -\frac{\partial H}{\partial x^n}\right)$. So to the family $\mathcal{H}^L$ of Hamiltonians there corresponds a family $\vec{\mathcal{H}}^L = \{\vec{\mathcal{H}}^L_\eta\}_{\eta\in\mathcal{U}}$ of time-varying vector fields on $\Omega \times \mathbb{R}^n$.

The necessary conditions for optimality then say, simply, that an optimal arc ξ_* and its corresponding control function η_* must satisfy the following: there is an arc ζ in $\mathbb{R}^n$ such that (i) $\Xi = (\xi_*, \zeta)$ is an integral curve of the Hamiltonian system on $\Omega \times \mathbb{R}^n$ whose dynamical equations are $\dot{X}(t) = \vec{\mathcal{H}}^L_\eta(X(t), t)$, (ii) for almost every t, η_* minimizes the function $\mathcal{U} \ni \eta \to \mathcal{H}^L_\eta(\Xi(t), t)$. More succinctly:

(a) *to each Lagrangian $L : \Omega \times \mathbb{R}^n \times I \to \mathbb{R}$ there corresponds a Hamiltonian system—i.e. a collection $\vec{\mathcal{H}}^L$ of time-varying Hamiltonian vector fields—on $\Omega \times \mathbb{R}^n$, and*

(b) *every optimal arc ξ_* corresponding to a control η_* can be lifted to an integral curve Ξ of the Hamiltonian system in such a way that the control η_* minimizes the Hamiltonian.*

To understand the control theory generalization, notice that the situation considered above really has two ingredients, namely (a) a dynamics (or kinematics) Σ, which singles out the class of curves over which the minimization is to be performed, and (b) a Lagrangian L. The dynamics Σ happens to be trivial, in the sense that Σ is the class of all absolutely continuous curves in Ω. Let us use U_Ω to denote the class of all constant vector fields on Ω, and let $\mathcal{U}_\Omega$ be the set of all absolutely continuous maps $t \to \eta(t) \in U_\Omega$. Now let us identify our set $\mathcal{U}$ of controls with $\mathcal{U}_\Omega$ in an obvious way, by thinking of our control functions η as U_Ω-valued functions (so each value $\eta(t) \in \mathbb{R}^n$ is now thought of as a constant vector field, and then η itself becomes a time-varying vector field). Then Σ is the set of all arcs that are integral curves of time-varying vector fields $\eta \in \mathcal{U}_\Omega$.

The control theory generalization consists of replacing this special class Σ by much more general classes of curves, arising from more general families of time-varying vector fields.

More precisely, rather than take $\mathcal{U}$ to be the special class $\mathcal{U}_\Omega$, we now let $\mathcal{U}$ be an arbitrary parameter set, and consider families $F = \{F_\eta\}_{\eta \in \mathcal{U}}$ of time-varying vector fields, parametrized by η. The dynamical equations then become $\dot{x} = F_\eta(x, t)$, and the Lagrangian will be, as in the classical case, a family of functions of x, t parametrized by η. (The classical case corresponds to letting $F_\eta(x, t) = \eta(t)$.)

It is clear that once we have decided to operate on this level of generality there is no longer any need to work on open subsets of Euclidean spaces, so *the natural spaces for our optimization problems are manifolds*.

In the Calculus of Variations problem, the O.D.E. $\dot{x}(t) = \eta(t)$ determines the control η given the curve $x(\cdot)$, so the cost can be viewed as functional on the space of curves. In the more general setting, it may happen that a curve ξ is generated by more than one control η, and the cost J could depend on η as well as ξ. For this reason, it is now natural to introduce the concept of a "controlled trajectory," i.e. of a pair (ξ, η) such η is a control and ξ is a trajectory generated by η.

Example 2.1. Consider "the control system $\dot{x}_1 = \eta_1(t)$, $\dot{x}_2 = \eta_2(t)x_1$, $(x_1, x_2) \in \mathbb{R}^2$, η integrable." That is, we take

$$\mathcal{U} = \bigcup_{-\infty < a \leq b < \infty} L^1([a,b], \mathbb{R}^2)$$

and associate with each $\eta \in \mathcal{U}$ with domain I the time-varying vector field $F_\eta : (x_1, x_2, t) \to (\eta_1(t), \eta_2(t)x_1)$, i.e. $F_\eta(\cdot, t) = \eta_1(t)\frac{\partial}{\partial x_1} + \eta_2(t)x_1\frac{\partial}{\partial x_2}$. Let the Lagrangian L be given by $L(x_1, x_2, u_1, u_2) = u_1^2 + u_2^2$, i.e. let $L_\eta(x_1, x_2, t) = \eta_1(t)^2 + \eta_2(t)^2$. Let $\xi : [0,1] \to M$ be the arc given by $\xi(t) = (0,0)$. Then ξ is generated by the control $\eta(t) \equiv (0,0)$ and also by the control $\widetilde{\eta}(t) = (0,1)$. The costs of the controlled arcs (ξ, η) and $(\xi, \widetilde{\eta})$ are different, since $J(\xi, \eta) = 0$ and $J(\xi, \widetilde{\eta}) = 1$. ∎

Pursuing the analogy with the classical case, it is natural to assign to each control system $\Sigma = (M, \mathcal{U}, F)$—where $F = \{F_\eta\}_{\eta \in \mathcal{U}}$—and Lagrangian $L = \{L_\eta\}_{\eta \in \mathcal{U}}$, a family of *Hamiltonian functions* $\mathcal{H}_\eta^{\Sigma,L}$, defined by

$$\mathcal{H}_\eta^{\Sigma,L}(x, z, t) = \langle z, F_\eta(x,t) \rangle + L_\eta(x,t), \tag{2.6}$$

(This is formally analogous to (2.5), since both formulas can be written, somewhat less precisely, as $\mathcal{H} = \langle z, \dot{x} \rangle + L$.) Since $F_\eta(x,t)$ is a tangent vector at x, it is clear that, for (2.6) to make sense, z has to be a *covector* at x. So each $\mathcal{H}_\eta^{\Sigma,L}$ should be regarded as a function on the cotangent bundle T^*M. We recall that, for a Hamiltonian function H on T^*M, there is a well-defined Hamiltonian vector field $\vec{H}$, whose integral curves obey, in coordinates, the Hamilton equations $\dot{x} = \frac{\partial H}{\partial z}$, $\dot{z} = -\frac{\partial H}{\partial x}$.

Having defined the Hamiltonian functions, it is natural to expect that, under suitable technical assumptions (such as continuous differentiability with respect to x of each F_η and each L_η, which is needed to guarantee the existence of the partial derivatives of the Hamiltonians with respect to x), the obvious analogue of Theorem 2.2 may be true. Precisely, suppose $\Sigma = (M, \mathcal{U}, F)$ is a control system and L is a Lagrangian for Σ. Use $\mathrm{Carc}(\Sigma)$ to denote the set of all controlled arcs of Σ. If $\gamma = (\xi, \eta) \in \mathrm{Carc}(\Sigma)$ has domain $[a, b]$, write

$$\partial\gamma = (\xi(a), \xi(b)), \qquad \widehat{\partial}\gamma = (a, \xi(a), b, \xi(b)), \qquad J^{\Sigma,L}(\gamma) = \int_I L_\eta(\xi(t), t)\, dt.$$

Call a $\gamma_* = (\xi_*, \eta_*) \in \mathrm{Carc}(\Sigma)$ with domain $[a, b]$ a *minimizer* if

$$J^{\Sigma,L}(\gamma_*) \leq J^{\Sigma,L}(\gamma) \text{ for all } \gamma \in \mathrm{Carc}(\Sigma) \text{ such that } \widehat{\partial}\gamma = \widehat{\partial}\gamma^*. \tag{2.7}$$

The natural extension of Theorem 2.2 would then say that if γ_* is a minimizer then there is an absolutely continuous function $\zeta : [a, b] \to T^*M$ such that (i) $\zeta(t) \in T^*_{\xi(t)}M$ for each $t \in [a, b]$, (ii) the arc $t \to (\xi_*(t), \zeta(t))$ in T^*M is an integral curve of the time-varying vector field $t \to \vec{\mathcal{H}}^{\Sigma,L}_{\eta_*}$, and (iii) the control η_* minimizes the Hamiltonian, in the sense that

$$\mathcal{H}^{\Sigma,L}_{\eta_*}(\xi_*(t), \zeta(t), t) = \min\{\mathcal{H}^{\Sigma,L}_{\eta}(\xi_*(t), \zeta(t), t) : \eta \in \mathcal{U}\} \qquad \text{for a.e. } t \in [a, b]. \tag{2.8}$$

It turns out that the above statement is not quite correct in general, even under the most restrictive smoothness assumptions. However, a slightly modified version is correct. First, one has to to work not just with the Hamiltonian $\mathcal{H}^{\Sigma,L}$, but with all the Hamiltonians $\mathcal{H}^{\Sigma,\nu L}$, for all $\nu \geq 0$. (The new parameter ν is called the "abnormal multiplier.") And second, one has to substitute for the "strong" Hamiltonian minimization condition (2.8) the weaker property

$$(\forall \eta \in \mathcal{U})(\mathcal{H}^{\Sigma,\nu L}_{\eta_*}(\xi_*(t), \zeta(t), t) \leq \mathcal{H}^{\Sigma,\nu L}_{\eta}(\xi_*(t), \zeta(t), t) \qquad \text{for a.e. } t \in [a, b]). \tag{2.9}$$

If we write $\Xi = (\xi, \zeta)$ as before, use $T^{\#}M$ to denote the cotangent bundle of M with the zero section removed, and let π_{T^*M} denote the canonical projection from T^*M to M, then the correct statement of the Maximum Principle turns out to be—modulo technical assumptions—the following:

(MP) *Let $\Sigma = (M, \mathcal{U}, F)$ be a control system, and let L be a Lagrangian for Σ. Let $\gamma_* = (\xi_*, \eta_*) \in \mathrm{Carc}(\Sigma)$, $\widehat{\partial}\gamma_* = (a, x_-, b, x_+)$ be a minimizer. Then there exist a constant $\nu \geq 0$ and an integral curve $\Xi_* : [a, b] \to T^*M$ of the time-varying vector field $\vec{\mathcal{H}}^{\Sigma,\nu L}_{\eta_*}$, such that* (i) $\pi_{T^*M} \circ \Xi_* = \xi_*$, (ii) *either $\nu > 0$ or Ξ_* takes values in $T^{\#}M$, and* (iii) (2.9) *holds.*

This is not yet a rigorous theorem, because the technical assumptions have not been stated precisely. A precise statement of a more general result under minimal technical conditions will be given below in Section 8. However, some observations are already in order at this point:

1. Except for the presence of the "abnormal multiplier" ν and the use of "weak" rather than "strong" Hamiltonian minimization, **(MP)** is an almost verbatim transcription of Theorem 2.2 to our more general setting.
2. If we write $\Xi_*(t) = (\xi_*(t), \zeta_*(t))$, then Condition (i) of **(MP)** says, in coordinates, that ζ_* satisfies the "adjoint equation"
$$\dot{\zeta}_*(t) = -\zeta_*(t)\frac{\partial F_{\eta_*}}{\partial x}(\xi_*(t), t) - \nu\frac{\partial L_{\eta_*}}{\partial x}(\xi_*(t), t),$$
where we should think of ζ_* and $\frac{\partial L_\eta}{\partial x}$ as row vectors, and of $\frac{\partial F_\eta}{\partial x}$ as a square matrix. If we add to these the equation $\dot{\nu} = 0$, we obtain a system which is *linear* in (ζ_*, ν). Therefore the function $t \to (\zeta_*(t), \nu)$ either vanishes identically or vanishes nowhere, so the nontriviality condition (ii) of **(MP)** is equivalent to the requirement that $(\zeta_*(t), \nu) \neq 0$ for *some* t.
3. The pair $\zeta_*(t) = 0$, $\nu = 0$ always satisfies all the conditions of **(MP)** other than the nontriviality property (ii). Hence (ii) is essential, for without it **(MP)** would give no information.
4. If a pair (ζ_*, ν) satisfies the conditions of **(MP)**, and $\rho > 0$, then the

pair $(\rho\zeta_*, \rho\nu)$ also satisfies those conditions. Hence, if one can satisfy the conditions with $\nu \neq 0$ then it is in fact possible to satisfy them with $\nu = 1$. Hence the only real effect of introducing the abnormal multiplier is that we now have to consider *two* Hamiltonians, namely, $\mathcal{H}^{\Sigma,L}$ and $\mathcal{H}^{\Sigma,0}$.

5. Under fairly general conditions (cf. Sections 9 and 10 below) the weak form (2.9) of Hamiltonian minimization actually implies the strong version:

$$\mathcal{H}^{\Sigma,\nu L}_{\eta_*}(\Xi_*(t),t) = \min\{\mathcal{H}^{\Sigma,\nu L}_{\eta}(\Xi_*(t),t) : \eta \in \mathcal{U}\} \quad \text{for almost all } t \in [a,b]. \tag{2.10}$$

6. For the classical Calculus of Variations situation the possibility that $\nu = 0$ can be excluded. Indeed, in that case the Hamiltonian $\mathcal{H}$ is just $\langle z, \eta(t)\rangle + \nu L(x, \eta(t), t)$. If $\nu = 0$, then $\mathcal{H} = \langle z, \eta(t)\rangle$. For our reference arc ξ_* and adjoint vector ζ_*, this equals $\langle \zeta_*(t), \eta(t)\rangle$. An elementary argument shows that in this case (2.10) easily follows from (2.9). So the control value $\eta_*(t) = \dot{\xi}_*(t)$ minimizes the Hamiltonian among all possible choices of η, for almost all t. Since $\{\eta(t) : \eta \in \mathcal{U}\} = \mathbb{R}^n$, the vector $\eta_*(t)$ must minimize–for a.e. t–the linear functional $v \to \langle \zeta_*(t), v\rangle$ over all $v \in \mathbb{R}^n$. Therefore $\zeta_*(t) = 0$, because a nonzero linear functional on $\mathbb{R}^n$ does not have a minimum. So the assumption that $\nu = 0$ implies that $\zeta_* = 0$ as well, contradicting the nontriviality condition (iv).

It is well known that **(MP)** can be false if the abnormal multiplier is not included, as the following simple example shows:

Example 2.2. Consider the dynamics $\dot{x} = x^2 + u^2$, with a control $u \in \mathbb{R}$. Suppose we want to minimize $\int_0^1 u(t)\,dt$ among all trajectories that go from $x = 0$ to $x = 0$ in time 1 and are generated by integrable control functions $[0,1] \ni t \to u(t) \in \mathbb{R}$. Clearly, there is only one such trajectory, namely, $x_*(t) \equiv 0$, corresponding to $u_*(t) \equiv 0$. So x_* is obviously minimizing. The Hamiltonian is $H = \zeta(x^2 + u^2) + \nu u$, and $t \to \zeta(t)$ must satisfy $\dot{\zeta} = -\zeta x$, which becomes $\dot{\zeta} = 0$ along x_*, u_*, so $\zeta(t) = c =$ constant. The minimization of the Hamiltonian says that $u \to \zeta u^2 + \nu u$ must be minimized by $u = 0$. This is obviously impossible if $\nu \neq 0$. On the other hand, by taking $\nu = 0$ and $c = 1$ we see that the conclusion of **MP** holds. ∎

Remark 2.1. The Hamiltonian corresponding to $\nu = 0$ does not involve L. Notice that in Example 2.2 our reference trajectory is "rigid," in the sense that there is no other trajectory satisfying the desired endpoint conditions. This guarantees *a priori* that the reference trajectory must be optimal for *every* cost functional, so it is not surprising that this fact should result in a condition not involving the Lagrangian at all. ∎

3. Notations and Basic Definitions

A list of abbreviations used throughout the paper—of expressions whose meaning is either well known or will be defined later in the text—is provided in an appendix following the References.

An *interval* is a nonempty connected subset of $\mathbb{R}$, so intervals can be closed, open or half-closed, bounded or unbounded. We use $\mathcal{I}$ to denote the set of all nontrivial intervals. An interval is *nontrivial* if it is not reduced to a point. If $I \in \mathcal{I}$, then $L^1_{\mathrm{loc},+}(I)$ is the set of all nonnegative locally Lebesgue integrable functions on I.

We use $\mathcal{D}(f)$, $f\lceil S$ to denote, respectively, the domain of a map f and the restriction of f to a set S. The word "function" refers to an $\mathbb{R}$-valued map, unless we specify otherwise (e.g. by talking about a "vector-valued function"). A *family* of objects parametrized by a set S is the same as a map f with domain S, and the notation $\{f(s)\}_{s\in S}$ is just another name for the map f. However, the expression $\{f(s) : s \in S\}$ refers to the *set* of values of f for $s \in S$, i.e. the set $f(S)$. If S is a set, a *time-varying map* on S is a map whose domain is $S \times I$ for some $I \in \mathcal{I}$. If F is a time-varying map on S, and $\mathcal{D}(F) = S \times I$, then I is the *time domain* of F, and we write $I = \mathcal{TD}(F)$. If $F_1, \ldots, F_m$ are time-varying maps on the same set S, we write $\mathcal{TD}(F_1, \ldots, F_m) = \mathcal{TD}(F_1) \cap \ldots \cap \mathcal{TD}(F_m)$. For $t \in \mathcal{TD}(F)$, $x \in S$, we use F_t, F^x to denote, respectively, the partial maps $F(\cdot, t)$ and $F(x, \cdot)$, so $\mathcal{D}(F_t) = S$ and $\mathcal{D}(F^x) = \mathcal{TD}(F)$.

"Smooth" will always mean "of class C^∞." A *manifold* is always, by definition, smooth, finite-dimensional, Hausdorff, without boundary, and of pure dimension (i.e. all connected components have the same dimension). If M, N are manifolds, $f : N \to M$ is a smooth map, F is a smooth manifold, and $\Omega \subseteq M$ is open, a *smooth F-trivialization* of f over Ω is a smooth diffeomorphism $\Phi : \Omega \times F \to f^{-1}(\Omega)$ such that $f(\Phi(x, v)) = x$ for all $x \in \Omega$, $v \in F$. We call f *F-locally trivial* if M can be covered by open sets over which there is a smooth F-trivialization of f. A *smooth bundle over M with fiber F* is a manifold E endowed with an F-locally trivial smooth map $\pi_E : E \to M$. The submanifold $E(x) = \pi_E^{-1}(x)$ is the *fiber* of E over x. If $\mathcal{M}$ is a class of manifolds endowed with extra structure —e.g. linear spaces, affine spaces— and $F \in \mathcal{M}$, then an *$\mathcal{M}$-bundle with fiber F* is a smooth bundle E equipped with an $\mathcal{M}$-structure on each fiber $E(x)$, such that M can be covered by open sets Ω_α over which there are F-trivializations Φ_α such that each map $\Phi_\alpha^x : v \to \Phi_\alpha(x, v)$ is an $\mathcal{M}$-isomorphism between F and $E(x)$. If $\mathcal{M}$ is the class of finite-dimensional real linear spaces, then E is a *smooth vector bundle*. In that case, E^* denotes the dual of E. (When $E = TM$, the tangent bundle of M, we write T^*M rather than TM^*, so T^*M is the cotangent bundle of M, and we use T_xM, T_x^*M, rather than $TM(x)$, $T^*M(x)$ to denote the fibers, i.e. the tangent and cotangent spaces of M at $x \in M$. We write $T^\# M$ for the cotangent bundle of M with the zero section removed.) If $y \in E(x)$, $z \in E^*(x)$ (in particular if $y \in T_xM$ and $z \in T_x^*M$), we use indistinctly the notations $\langle z, y\rangle$, $z.y$ and $z(y)$ for the value at y of the linear functional z.

If M is a manifold, E is a smooth vector bundle over M, and $k \in \overline{\mathbb{N}}$, then $C^k(M)$, $\Gamma^k(E)$ denote, respectively, the space of real-valued functions of class C^k on M and the $C^k(M)$-module of sections of class C^k of E. In particular, $\Gamma^k(TM)$ and $\Gamma^k(T^*M)$ are, respectively, the spaces of vector fields and 1-forms of class C^k on M. If M is second countable, then the spaces $C^k(M)$ and $\Gamma^k(E)$ are separable Fréchet. Convergence in $C^k(M)$ (resp. $\Gamma^k(E)$) is uniform convergence on compact sets of the functions (resp. sections of E) together with all their partial derivatives of order $\leq k$. (Precisely: (i) a sequence $\{\varphi_j\}$ of functions in $C^k(M)$ converges in $C^k(M)$ to a function φ if for every i such that $0 \leq i < k+1$ and every i-tuple $(X_1, \ldots, X_i)$ of smooth vector fields on M, the functions $X_1 X_2 \ldots X_i \varphi_j$ converge to $X_1 X_2 \ldots X_i \varphi$ uniformly on compact sets, and (ii) a sequence $\{S_j\}$ in $\Gamma^k(E)$ converges in $\Gamma^k(E)$ to a section S if for every $S^* \in \Gamma^\infty(E^*)$ the scalar functions $\langle S^*, S_j\rangle$ converge to $\langle S^*, S\rangle$ in $C^k(M)$. The metrizability and separability of these spaces follow from the fact that M is second countable.) It is easy to see that $C^k(M)$ is dense in $C^i(M)$ whenever $k \leq i$.

If A is a topological space, we use $\mathrm{Comp}(A)$ to denote the set of all compact subsets of A. A *Carathéodory function* (CF) on a topological space A is a time-varying function φ on A such that (i) φ_t is continuous for every $t \in \mathcal{TD}(\varphi)$, and (ii) φ^x is Lebesgue measurable for every $x \in A$. (If A is separable metric, it then follows easily that φ must be jointly Borel$\times$Lebesgue measurable on $\mathcal{D}(\varphi)$.) We use $CF(A)$, $CF(I, A)$ to denote, respectively, the set of all CF's on A and the set of all $\varphi \in CF(A)$ such that $\mathcal{TD}(\varphi) = I$. If $\varphi \in CF(A)$, then φ will be called *locally integrably bounded* (LIB) if for every $K \in \mathrm{Comp}(A)$ there exists a $\psi \in L^1_{\mathrm{loc},+}(\mathcal{TD}(\varphi))$ such that $|\varphi(x,t)| \leq \psi(t)$ for all $(x,t) \in K \times \mathcal{TD}(\varphi)$. If A is metric, with distance function d, then a $\varphi \in CF(A)$ will be called *locally Lipschitz* (LL) if every φ_t is locally Lipschitz (i.e. φ satisfies a Lipschitz condition $|\varphi(x) - \varphi(y)| \leq C_{K,t} d(x,y)$ for $(x,y) \in K \times K$ for every $K \in \mathrm{Comp}(A)$, $t \in \mathcal{TD}(\varphi)$) and *locally integrably Lipschitz* (LIL) if it is LIB and LL, and for every $K \in \mathrm{Comp}(A)$ the Lipschitz constants $C_{K,t}$ can be chosen so that the function $t \to C_{K,t}$ is in $L^1_{\mathrm{loc},+}(\mathcal{TD}(\varphi))$. Notice that these definitions only involve the restriction of the metric to compact subsets of A, so the concepts of a LL and a LIL CF on A only depend on the class of d modulo local equivalence. (Two distance functions d_1, d_2 on a topological space A are *locally equivalent* if for every $K \in \mathrm{Comp}(A)$ there exist constants $c_1 > 0$, $c_2 > 0$ such that $c_1 d_1(x,y) \leq d_2(x,y) \leq c_2 d_1(x,y)$ for all $(x,y) \in K \times K$.) In particular, if A is a manifold, then all the distance functions arising from Riemannian metrics on M are locally equivalent, so the class of LIL CF's on M is well defined.

If M is a manifold, and $\varphi \in CF(M)$, we will call φ *of class* C^k if $\varphi_t \in C^k(M)$ for every $t \in \mathcal{TD}(\varphi)$, and *locally integrably of class* C^k— abbreviated LIC^k— if φ is of class C^k and the function $X_1X_2\ldots X_k\varphi$: $\mathcal{D}(\varphi) \to \mathbb{R}$ is LIB for every k-tuple $\mathbf{X} = (X_1, \ldots, X_k)$ of smooth vector fields on M. (For $k = 0$, this just amounts to saying that φ is LIB.)

If A is a topological space and $\mathcal{P}$ is any property of CF's that makes sense on A (e.g. LIB, LL and LIL if A is metric, C^k or LIC^k if A is a manifold), then $CF_{\mathcal{P}}(A)$, $CF_{\mathcal{P}}(I, A)$, will denote, respectively, the class of all $\varphi \in CF(A)$ that have Property $\mathcal{P}$, and the set of $\varphi \in CF_{\mathcal{P}}(A)$ with time domain I.

A *Carathéodory vector field*—or *control vector field*, abbreviated CVF — on M is a time-varying map F on M such that (i) $F(x,t) \in T_xM$ for every $(x,t) \in \mathcal{D}(F)$, and (ii) $F\varphi \in CF(M)$ for every $\varphi \in C^\infty(M)$. (By definition, $(F\varphi)(x,t) = F(x,t)\varphi$, so $\mathcal{D}(F\varphi) = \mathcal{D}(F)$.) We use $CVF(I, M)$ to denote the set of all CVF's on M with time domain I, and write $CVF(M) \stackrel{\text{def}}{=} \cup_I CVF(I, M)$. If $F \in CVF(M)$, we say that F is LL, of class C^k, LIB, LIL, LIC^k, if for every $\varphi \in C^\infty(M)$ the function $F\varphi$ is LL, of class C^k, LIB, LIL, LIC^k. If $\mathcal{P}$ is any property of $CVF's$, such as LL, C^k, LIB, LIL, LIC^k (or the LIBU property introduced below), we use $CVF_{\mathcal{P}}(I, M)$ to denote the set of those $F \in CVF(I, M)$ that have $\mathcal{P}$, and write $CVF_{\mathcal{P}}(M) = \cup_I CVF_{\mathcal{P}}(I, M)$. (Notice that $CVF_{C^0}(M)$ is just $CVF(M)$.)

A *curve* in M is a continuous map $\xi : I \to M$, where $I = \mathcal{D}(\xi)$ is an interval. The *graph* of a curve ξ is the subset $G(\xi) \stackrel{\text{def}}{=} \{(\xi(t), t) : t \in \mathcal{D}(\xi)\}$ of $M \times \mathbb{R}$. A curve $\xi : I \to M$ is *locally absolutely continuous* (LAC), *locally Lipschitz* (LL), *of class* C^k, if $\varphi \circ \xi$ is LAC, LL, or of class C^k for every $\varphi \in C^\infty(M)$. An *arc* is a LAC curve ξ such that $\mathcal{D}(\xi)$ is compact. In that case, ξ is in fact absolutely continuous (AC), i.e. $\varphi \circ \xi$ is AC for every $\varphi \in C^\infty(M)$. The self-explanatory notations $CRV(I, M)$, $CRV_{LAC}(I, M)$, $CRV(M)$, $CRV_{LAC}(M)$, $ARC(M)$ stand, respectively, for the sets of (i) all curves $\xi : I \to M$, (ii) all LAC members of $CRV(I, M)$, (iii) all curves in M, (iv) all LAC curves in M, (v) all arcs in M. When I is compact we write $ARC(I,M)$ rather than $CRV_{LAC}(I,M)$.

An *integral curve* (abbr. IC) of an $F \in CVF(M)$ is a LAC curve ξ such that $\mathcal{D}(\xi) \subseteq \mathcal{TD}(F)$ and $\dot{\xi}(t) = F(\xi(t),t)$ for almost all $t \in \mathcal{D}(\xi)$. We use $\mathrm{IC}(F)$ to denote the set of all IC's of F. A CVF F has the *existence property* if for every point $(\overline{x},\overline{t}) \in M \times \mathcal{TD}(F)$ there exists a $\xi \in \mathrm{IC}(F)$ such that (i) $\mathcal{D}(\xi)$ is a neighborhood of $\overline{t}$ relative to $\mathcal{TD}(F)$, and (ii) $\xi(\overline{t}) = \overline{x}$. We say that F has the *uniqueness property* if, whenever ξ_1, ξ_2 are IC's of F whose graphs intersect, it follows that $\xi_1(t) = \xi_2(t)$ for every $t \in \mathcal{D}(\xi_1) \cap \mathcal{D}(\xi_2)$.

The Carathéodory existence and uniqueness theorems assert that (i) every LIB CVF F has the existence property, and (ii) if F is LIL, then F has the uniqueness property. It is easy to see that, if F is a CVF that has the existence property, then through every point $(\overline{x},\overline{t})$ of $\mathcal{D}(F)$ there passes a *maximal IC* of F, i.e. an IC $\xi : I \to M$ that cannot be extended to an IC $\widetilde{\xi} : \widetilde{I} \to M$ of F defined on an interval $\widetilde{I}$ that strictly contains I. (If F is LIB, then it is easily shown that a $\xi \in \mathrm{IC}(F)$ is maximal if and only if (i) $G(\xi)$ is relatively closed in $\mathcal{D}(F)$ and (ii) $\mathcal{D}(\xi)$ is relatively open in $\mathcal{TD}(F)$.)

A *LIBU CVF* is a LIB CVF that has the uniqueness property.

A *controlled curve* in M is a pair $\gamma = (\xi, F)$ such that $F \in CVF(M)$ and $\xi \in \mathrm{IC}(F)$. The *domain* $\mathcal{D}(\gamma)$ of a controlled curve $\gamma = (\xi, F)$ is, by definition, the domain of ξ. We call a controlled curve *LIB-controlled* (resp. *LIBU-controlled*, *LIL-controlled*, *LIC^k-controlled*) if F is LIB (resp. LIBU, LIL, LIC^k).

If $\mathcal{F}$ is a set of CVF's in M, we write $\mathrm{IC}(\mathcal{F}) \stackrel{\mathrm{def}}{=} \cup_{F \in \mathcal{F}} \mathrm{IC}(F)$, and we use $\mathrm{Ccrv}(\mathcal{F})$ to denote the set of all controlled curves $\gamma = (\xi, F)$ such that $F \in \mathcal{F}$. If $\mathcal{P}$ is any property of $CVF's$, such as LL, C^k, LIB, LIL, LIC^k, LIBU, we use $\mathrm{Ccrv}_{\mathcal{P}}(\mathcal{F})$ to denote the set of those $\gamma \in \mathrm{Ccrv}_{\mathcal{P}}(\mathcal{F})$ that have $\mathcal{P}$. (In particular, $\mathrm{Ccrv}_{C^0}(\mathcal{F}) = \mathrm{Ccrv}_{LAC}(\mathcal{F}) = \mathrm{Ccrv}(\mathcal{F})$.)

If (N, Ω) is a symplectic manifold (i.e. a manifold N equipped with an everywhere nonsingular closed 2-form Ω), and $H \in C^k(N)$, $k \geq 1$, then we can associate with H a vector field $\vec{H} \in \Gamma^{k-1}(TN)$, known as the *Hamiltonian vector field* of H. More generally, if $H \in CF_{C^k}(N)$, $k \geq 1$, then $\vec{H}$ is a CVF of class C^{k-1}, with time domain $\mathcal{TD}(F)$, defined as follows. If $x \in N$, $t \in \mathcal{TD}(H)$, then $\vec{H}(x,t)$ is the vector $v \in T_xN$ such that $\langle dH_t(x), w\rangle = \Omega_x(v,w)$ for all $w \in T_xN$. The *Poisson bracket* $\{H,K\}$ of two functions $H, K \in C^1(N)$ is, by definition, the derivative of K along the IC's of $\vec{H}$, i.e. $\{H,K\} \stackrel{\text{def}}{=} \langle dK, \vec{H}\rangle$. More generally, if $H, K \in CF_{C^1}(M)$, then $\{H,K\}$ is defined on $M \times \mathcal{TD}(H,K)$ by $\{H,K\}(x,t) \stackrel{\text{def}}{=} \langle dK_t(x), \vec{H}(x,t)\rangle$, i.e. by $\{H,K\}(x,t) \stackrel{\text{def}}{=} \{H_t, K_t\}(x)$.

A *system of canonical coordinates* on a symplectic manifold (N,Ω) is a chart $\kappa = (x^1,\ldots,x^n,z_1,\ldots,z_n)$ of N such that $\Omega = \sum_{i=1}^n dz_i \wedge dx^i$ on $\mathcal{D}(\kappa)$. If κ is a system of canonical coordinates on (N,Ω), and $H \in C^1(N)$, then the expression $\vec{H}^\kappa$ of $\vec{H}$ with respect to κ is given, for $(x,z,t) \in \mathcal{D}(\kappa) \times \mathcal{TD}(H)$, by $\vec{H}^\kappa = \sum_{i=1}^n \left(\frac{\partial H}{\partial z_i}\frac{\partial}{\partial x^i} - \frac{\partial H}{\partial x^i}\frac{\partial}{\partial z_i}\right)$, so the IC's of $\vec{H}$ satisfy *Hamilton's equations* $\dot{x}^i = \frac{\partial H}{\partial z_i}(x,z,t)$, $\dot{z}_i = -\frac{\partial H}{\partial x^i}(x,z,t)$. The Poisson bracket of $H, K \in CF_{C^1}(N)$ is then given, on $\mathcal{D}(\kappa) \times \mathcal{TD}(H,K)$, by the formula $\{H,K\} = \sum_{i=1}^n \left(\frac{\partial H}{\partial z_i}\frac{\partial K}{\partial x^i} - \frac{\partial H}{\partial x^i}\frac{\partial K}{\partial z_i}\right)$.

An important special instance of the above situation arises when N is the cotangent bundle T^*M of a manifold M, in which case N is endowed with a canonical symplectic form $\Omega_M \stackrel{\text{def}}{=} d\omega_M$, where $\omega_M(w) \stackrel{\text{def}}{=} \langle z, d\pi_{T^*M}(w)\rangle$ for $z \in T_x^*M$, $w \in T_{(x,z)}T^*M$. In that case, whenever $F \in \Gamma^0(TM)$ and $L \in C^0(M)$—and, more generally, whenever $F \in CVF(M)$ and $L \in CF(M)$— we can associate with the pair (F,L) the *Hamiltonian* $H_{F,L} \in CF(T^*M)$, with time domain $\mathcal{TD}(H,L)$, defined by

$$H_{F,L}(x,z,t) \stackrel{\text{def}}{=} z.F(x,t) + L(x,t) \qquad \text{for } x \in M,\ z \in T_x^*M,\ t \in \mathcal{TD}(H,L).$$

It is clear that $H_{F,L}$ has any of the properties C^k, LIB, LIL, LIC^k, iff both F and L do. If $F \in CVF_{C^k}(M)$, $L \in CF_{C^k}(M)$, and $k \geq 1$, then $\vec{H}_{F,L} \in CVF_{C^{k-1}}(T^*M)$. The CVF $\vec{H}_{F,L}$ is the *inhomogeneous Hamiltonian lift of F with Lagrangian L*. When $L = 0$ we just write H_F, $\vec{H}_F$, and refer to H_F, $\vec{H}_F$ as the *Hamiltonian* and *Hamiltonian lift* of F. Naturally, if F and L are not time-varying, then $H_{F,L}$ is also independent of t, so in that case $H_{F,L}$ is just a function on T^*M.

If $\kappa = (x^1,\ldots,x^n)$ is a chart of M, and we let $z_i(x,z) = \left\langle z, \frac{\partial}{\partial x^i}\right\rangle$ for $z \in T_x^*M$, so that $z = \sum_{i=1}^n z_i(x,z)\,dx^i$, then $\kappa^* = (x^1,\ldots,x^n,z_1,\ldots,z_n)$ (where we are writing x^i rather than $x^i \circ \pi_{T^*M}$) is a system of canonical coordinates on T^*M, with domain $\pi_{T^*M}^{-1}(\mathcal{D}(\kappa))$, since $\omega_M = \sum_{i=1}^n z_i\,dx^i$. If $F \in CVF(M)$ and $L \in CF(M)$, then $H_{F,L}(x,z,t) = \sum_{i=1}^n z_i F^{\kappa,i}(x,t) + L(x,t)$ on $\mathcal{D}(\kappa^*)\times(\mathcal{T}\mathcal{D}(H,L))$, where $F^{\kappa,1},\ldots,F^{\kappa,n}$ are the κ-components of F, i.e. $F^{\kappa,i}(x,t) \stackrel{\mathrm{def}}{=} \langle dx^i, F(x,t)\rangle$. Therefore the IC's of $\vec{H}_{F,L}$ satisfy, with respect to the chart κ^*, the equations

$$(3.1)\quad \dot{x}^i(t) = F^{\kappa,i}(x(t),t), \qquad \dot{z}_i(t) = -z(t).\frac{\partial F^\kappa}{\partial x^i}(x(t),t) - \frac{\partial L}{\partial x^i}(x(t),t),$$

i.e. $\dot{x} = F^\kappa(x,t)$, $\dot{z} = -z\frac{\partial F^\kappa}{\partial x}(x,t) - dL_t(x)$, where F^κ is the vector-valued function $(x,t) \to (F^{\kappa,1}(x,t),\ldots,F^{\kappa,n}(x,t))$, and we are writing $F^\kappa(x,t)$ and the coordinates of x (also labeled x) as column vectors and the coordinates of z as a row vector, and letting $\frac{\partial F^\kappa}{\partial x}(x,t)$ be the square matrix whose columns are the vectors $\frac{\partial F^\kappa}{\partial x^i}(x,t)$, $i = 1,\ldots,n$. This shows that, if Ξ is an IC of $\vec{H}_{F,L}$, then its projection $\pi_{T^*M}\circ\Xi$ is an IC of F, justifying the name "Hamiltonian lift." It is clear that, if $F \in CVF_{LIC^k}(M)$, $L \in CF_{LIC^k}(M)$, and $k \geq 1$, then $\vec{H}_{F,L} \in CVF_{LIC^{k-1}}(T^*M)$, and $\vec{H}_{F,L}$ has existence and uniqueness of IC's even when $k = 1$. (For $k = 1$ $\vec{H}_{F,L}$ need not be locally Lipschitz, but (i) the equation $\dot{x} = F^\kappa(x,t)$ only involves the components of F, which are LIC^1, (ii) the equation $\dot{z} = -z\frac{\partial F^\kappa}{\partial x}(x(t),t) - dL_t(x(t))$ is linear affine in z for any given function $t \to x(t)$, and (iii) the matrix function $t \to \frac{\partial F^\kappa}{\partial x}(x(t),t)$ and the vector function $t \to dL_t(x(t))$ are integrable if $x(\cdot)$ is continuous, because F and L are LIC^1.)

4. Lie Differentiations; the Variational and Adjoint Equations

It is well known that, if M is a smooth manifold and $\xi : [a,b] \to M$ is a smooth curve, then there is no canonical way to define how to differentiate along ξ fields of vectors or covectors along ξ, unless some extra structure is associated with the pair (M,ξ). The most widely known kind of structure that makes such a differentiation possible is an affine connection on M, but this is by no means needed. It turns out that the "adjoint equation" of the MP involves a type of differentiation of a field of covectors along a curve that need *not* arise from a connection. The purpose of this section is to spell out in detail precisely what additional structure is needed for a curve ξ on a manifold M to give rise to a "Lie differentiation along ξ."

If M is a manifold, E is a smooth bundle over M, and $\xi \in CRV(I, M)$, then a *section of E along ξ* is a map σ with domain I that assigns to each $t \in I$ a member $\sigma(t)$ of $E(\xi(t))$. We use $\Gamma(E, \xi)$ to denote the set of sections of E along ξ. The members of $\Gamma(TM, \xi)$, $\Gamma(T^*M, \xi)$, are known, respectively, as *fields of vectors* and *fields of covectors* along ξ. We use $\Gamma^{meas}(E, \xi)$, $\Gamma^{LI}(E, \xi)$, $\Gamma^{LB}(E, \xi)$, $\Gamma^k(E, \xi)$ if $k \in \overline{\mathbb{N}}$, $\Gamma^{LAC}(E, \xi)$, $\Gamma^{LL}(E, \xi)$ to denote, respectively the set of all $\sigma \in \Gamma(E, \xi)$ that are measurable, LI, LB, of class C^k, LAC, LL. (Notice that $\pi_E \circ \sigma = \xi$, so the classes $\Gamma^k(E, \xi)$, $\Gamma^{LAC}(E, \xi)$, $\Gamma^{LL}(E, \xi)$ are void unless ξ itself is of class C^k, LAC, or LL.)

Let M be a manifold, and let $(x, v) \in TM$. A *covariant differentiation*, or *Lie differentiation*, at x in the direction of v is an $\mathbb{R}$-linear map $L : \Gamma^\infty(TM) \to T_xM$ that satisfies $L(\alpha X) = \alpha(x)LX + (v\alpha).X(x)$ for all $\alpha \in C^\infty(M)$, $X \in \Gamma^\infty(TM)$. We use $\mathbf{LD}(x, v)$ to denote the set of all Lie differentiations at x in the direction of v. It is easy to see that if $L \in \mathbf{LD}(x, v)$ then $L(X)$ only depends on the 1-jet of X at x, so L can also be thought of as a map from $J^1_x(TM)$ to T_xM, where $J^1_x(TM)$ is the space of 1-jets at x of vector fields on M. Clearly, if $L_i \in \mathbf{LD}(x, v_i)$ and $r_i \in \mathbb{R}$ for $i = 1, 2$, then $r_1L_1 + r_2L_2 \in \mathbf{LD}(x, r_1v_1 + r_2v_2)$. So each $\mathbf{LD}(x, v)$ is closed under affine combinations (i.e. under linear combinations with coefficients that add up to 1). We let $\mathbf{LD}(x) = \cup_{v \in T_xM}\mathbf{LD}(x, v)$, and $\mathbf{LD}[M] = \cup_{x \in M}\mathbf{LD}(x)$. Then $\mathbf{LD}(x)$ is a linear space, and there is a natural projection $\pi_{\mathbf{LD},x} : \mathbf{LD}(x) \to T_xM$ that assigns to each $L \in \mathbf{LD}(x)$ the unique $v \in T_xM$ such that $L \in \mathbf{LD}(x, v)$. If X is a smooth vector field on M, then X gives rise to an element $L_{x,X} \in \mathbf{LD}(x, X(x))$, defined by $L_{x,X}(Y) = [X, Y](x)$. The map $X \to L_{x,X}$ from $\Gamma^\infty(TM)$ to $\mathbf{LD}(x)$ is clearly linear. It is easily verified that the kernel of this map is precisely the set of vector fields that vanish at x together with their first derivatives. Moreover, the map $X \to L_{x,X}$ is clearly surjective. So $\mathbf{LD}(x)$ can be canonically identified with $J^1_x(TM)$. Therefore, if $n = \dim M$ then $\mathbf{LD}[M]$ is a smooth vector bundle over M, with fiber dimension $n^2 + n$, and we can associate with $\mathbf{LD}[M]$ a bundle $\mathbf{LD}[[M]]$ over TM whose fiber at each (x,v) is $\mathbf{LD}(x,v)$, so $\mathbf{LD}[[M]]$ is a bundle of affine (but not linear) n^2-dimensional spaces.

If $\kappa = (x^1,\ldots,x^n)$ is a chart near x, and $\partial_i^\kappa = \frac{\partial}{\partial x^i}$, then an $L \in \mathbf{LD}(x,v)$ is determined by an array $\{L_i^j\}_{1\leq i,j\leq n}$ of coefficients, defined by $L(\partial_i^\kappa) = \sum_{j=1}^n L_i^j \partial_j^\kappa(x)$. The Lie derivative LY of an arbitrary vector field $Y = \sum_j Y^j \partial_j^\kappa$ is then given by $LY = \sum_j (LY)^j \partial_j^\kappa(x)$, where $(LY)^j = vY^j + \sum_i Y^i(x)L_i^j$.

We now specialize to the case when X is a vector field of class C^1 defined near x, and $L = L_{x,X}$. (Since $L_{x,X}$ only depends on the 1-jet of X at x, it is well defined for any X which is of class C^1 near x.) If $Y \in \Gamma^\infty(TM)$, then the components of $[X,Y]$ are given by $[X,Y]^j(x) = X(x)Y^j - \sum_i \partial_i^\kappa X^j(x)Y^i(x)$. This says that, under the identification described earlier, the Lie differentiation $L_{x,X}$ that corresponds to X has coefficients $(L_{x,X})_i^j = -\partial_i^\kappa X^j(x)$.

A Lie differentiation $L \in \mathbf{LD}(x,v)$ gives rise to a similar operator—also denoted by L—acting on arbitrary tensor fields on M, which assigns to a section σ of the tensor bundle $T^{k,\ell}M \stackrel{\text{def}}{=} (\otimes^k TM) \otimes (\otimes^\ell T^*M)$ an element $L\sigma$ of $T_x^{k,\ell}M$, in such a way that the product differentiation rule $L(B(\sigma_1,\sigma_2)) = B(L\sigma_1,\sigma_2) + B(\sigma_1,L\sigma_2)$ holds when B is a contraction $\Gamma^\infty(T^{k_1,\ell_1}M) \to \Gamma^\infty(T^{k_1-1,\ell_1-1}M)$ of a covariant and a contravariant index or a canonical tensor product map $\Gamma^\infty(T^{k_1,\ell_1}M) \times \Gamma^\infty(T^{k_2,\ell_2}M) \to \Gamma^\infty(T^{k_1+k_2,\ell_1+\ell_2}M)$. Since this extension to general tensors will not be needed here, we will only describe the one case that will actually be used, namely, the extension to fields ω of covectors, i.e. 1-forms. If $\omega \in \Gamma^\infty(T^*M)$, and $L \in \mathbf{LD}(x,v)$, we define $L\omega \in T_x^*M$ by the formula $\langle L\omega, y\rangle = v\langle\omega,Y\rangle - \langle\omega(x),LY\rangle$, for $y \in T_xM$, $Y \in \Gamma^\infty(TM)$ such that $Y(x) = y$. The equation $L(\alpha Y) = \alpha(x)LY + (v\alpha)Y(x)$ implies that $\langle L\omega,y\rangle$ does not depend on Y, so $L\omega$ is well defined as a covector at x. Also, the identity $L(\alpha\omega) = (v\alpha)\omega(x) + \alpha(x)L\omega$ is easily verified, and the product rule $L\langle\omega,Y\rangle = \langle L\omega,Y(x)\rangle + \langle\omega(x),LY\rangle$ is a trivial consequence of the definition, if we interpret $L\varphi$ as $v\varphi$ when φ is a function. In coordinates, if $\omega = \sum_i \omega_i dx^i$, then $L\omega = \sum_i (L\omega)_i dx^i$, where $(L\omega)_i = v\omega_i - \sum_j L_i^j \omega_j(x)$. In particular, if $X \in \Gamma^\infty(TM)$, then $(L_{x,X}\omega)_i = v\omega_i + \sum_j \frac{\partial X^j}{\partial x^i}(x)\omega_j(x)$.

If $L \in \mathbf{LD}(\overline{x}, v)$, I is a nontrivial interval, $\xi : I \to M$ is a curve, $\overline{t} \in I$, ξ is differentiable at $\overline{t}$, $\xi(\overline{t}) = \overline{x}$, and $\dot{\xi}(\overline{t}) = v$, then we can define the Lie derivative Lw of any continuous field $I \ni t \to w(t) \in T_{\xi(t)}M$ of vectors along ξ which is differentiable at $\overline{t}$. We do this by first finding a time-varying vector field $M \times I \ni (x,t) \to W(x,t) \in T_x M$ which is smooth in x for each t, jointly differentiable at $(\overline{x}, \overline{t})$, and such that $W(\xi(t), t) = w(t)$ for $t \in I$, t near $\overline{t}$. (The existence of W is obvious: let $f_1, \ldots, f_n \in \Gamma^\infty(TM)$ be such that $f_1(\overline{x}), \ldots, f_n(\overline{x})$ is a basis of $T_{\overline{x}}M$. Write $w(t) = \sum_{i=1}^n w^i(t) f_i(\xi(t))$ for t near $\overline{t}$, where the functions w_i are continuous near $\overline{t}$, and differentiable at $\overline{t}$. Extend the w^i arbitrarily to continuous functions on I. Let $W(x,t) = \sum_{i=1}^n w^i(t) f_i(x)$ for $(x,t) \in M \times I$.) We then define

$$Lw \stackrel{\text{def}}{=} LW_{\overline{t}} + D_{\overline{t}} W(\overline{x}, t), \tag{4.1}$$

where $D_{\overline{t}} = \frac{d}{dt}\Big|_{t=\overline{t}}$. It is then easy to see that Lw does not depend on the choice of the extension W. (Indeed, choose a chart $\kappa = (x^1, \ldots, x^n)$ near $\overline{x}$, let $\xi(t)$ have components $\xi^1(t), \ldots, \xi^n(t)$, and write

$$w(t) = \sum_{j=1}^n w^j(t) \partial_j^\kappa(\xi(t)), \qquad W(x,t) = \sum_{j=1}^n W^j(x,t) \partial_j^\kappa(x),$$
$$v = \sum_{j=1}^n v^j \partial_j^\kappa(\overline{x}), \quad \text{and} \quad L(\partial_i^\kappa) = \sum_{j=1}^n L_i^j \partial_j^\kappa(\overline{x}).$$

Then $W^i(\xi(t), t) = w^i(t)$, and

$$\begin{aligned}
&LW_{\overline{t}} + D_{\overline{t}} W(\overline{x}, t) \\
&\quad = \sum_{j=1}^n \Big(D_{\overline{t}} W^j(\overline{x}, t) + \sum_{i=1}^n v^i \frac{\partial W^j}{\partial x^i}(\overline{x}, \overline{t}) + \sum_{i=1}^n L_i^j W^i(\overline{x}, \overline{t}) \Big) \partial_j^\kappa(\overline{x}) \\
&\quad = \sum_{j=1}^n \Big(D_{\overline{t}} W^j(\overline{x}, t) + \sum_{i=1}^n \dot{\xi}^i(\overline{t}) \frac{\partial W^j}{\partial x^i}(\overline{x}, \overline{t}) + \sum_{i=1}^n L_i^j W^i(\overline{x}, \overline{t}) \Big) \partial_j^\kappa(\overline{x}) \\
&\quad = \sum_{j=1}^n \Big(D_{\overline{t}} W^j(\xi(t), t) + \sum_{i=1}^n L_i^j W^i(\overline{x}, \overline{t}) \Big) \partial_j^\kappa(\overline{x}),
\end{aligned}$$

so that

$$
(4.2)\qquad
\begin{aligned}
Lw &= \sum_{j=1}^{n}\Big(D_{\bar t}w^j(t) + \sum_{i=1}^{n} L_i^j w^i(\bar t)\Big)\partial_j^\kappa(\bar x) \quad \text{if}\\
w(t) &= \sum_{j=1}^{n} w^j(t)\partial_j^\kappa(\xi(t)),
\end{aligned}
$$

proving that Lw does not depend on W.) In exactly the same way, we can define the Lie derivative $L\omega$ of a continuous field $I \ni t \to \omega(t) \in T^*_{\xi(t)}M$ of covectors along ξ which is differentiable at $\bar t$. The resulting formula, in coordinates, is

$$
(4.3)\qquad
\begin{aligned}
L\omega &= \sum_{i=1}^{n}\Big(D_{\bar t}\omega_i(t) - \sum_{j=1}^{n} L_i^j \omega_j(\bar t)\Big)\, dx^i(\bar x) \quad \text{if}\\
\omega(t) &= \sum_{i=1}^{n} \omega_i(t)dx^i(\xi(t)).
\end{aligned}
$$

With the above terminology, an *affine connection* on M is a smooth section ∇ of $\mathbf{LD}[[M]]$—i.e. a choice, for each $(x,v) \in TM$, of a Lie differentiation $\nabla_v^x \in \mathbf{LD}(x,v)$ depending smoothly on (x,v)—such that ∇_v^x, regarded as a member of the linear space $\mathbf{LD}(x)$, depends linearly on v for each x. On the other hand, if $k \geq 1$, then a vector field $X \in \Gamma^k(TM)$ gives rise to a section $\mathcal{L}_X$—called *Lie differentiation along* X—of class C^{k-1} of the bundle $\mathbf{LD}[M]$, defined by letting $\mathcal{L}_X(x) = L_{x,X}$. It then follows that $\pi_{\mathbf{LD}[[M]]} \circ \mathcal{L}_X = X$—i.e. $\pi_{\mathbf{LD},x}(\mathcal{L}_X(x)) = X(x)$ for all x—and $\mathcal{L}_X Y = [X,Y]$ for all $Y \in \Gamma^1(TM)$.

We can also define connections (or "Lie differentiations") along a LAC curve ξ in M defined on an interval I. Such a connection will be a map $\mathcal{L}$ that assigns to every $t \in I$ a member $\mathcal{L}_t$ of $\mathbf{LD}(\xi(t))$, in such a way that $\mathcal{L}_t \in \mathbf{LD}(\xi(t), \dot\xi(t))$. It clear that there is a well defined concept of measurablity of such maps, and we shall incorporate the measurability requirement into the definition. (Naturally, $\mathcal{L}$ is measurable iff for every $Y \in \Gamma^\infty(TM)$ the map $I \ni t \to \mathcal{L}_t Y \in TM$ is measurable.) Moreover, it is natural to identify two connections if they agree almost everywhere, and this leads us to the following

Definition 4.1. Let M be a smooth manifold, and assume that $\xi : I \to M$ is a LAC curve. A *connection*—or *Lie differentiation*, or *covariant differentiation*—along ξ is an equivalence class, modulo equality a.e., of Lebesgue measurable sections $\mathcal{L}$ of $\mathbf{LD}[M]$ along ξ such that $\pi_{LD[[M]]} \circ \mathcal{L} = \dot{\xi}$ a.e., i.e. $\mathcal{L}(t) \in \mathbf{LD}(\xi(t), \dot{\xi}(t))$ for a.e. $t \in I$. ∎

We use $\mathbf{LD}(\xi)$ to denote the set of all Lie differentiations along ξ and, for $\mathcal{L} \in \mathbf{LD}(\xi)$, we write $\mathcal{L}_t$ rather than $\mathcal{L}(t)$. It is clear that $\mathbf{LD}(\xi)$ is an affine (but not a linear) space.

Relative to any chart $\kappa = (x^1, \ldots, x^n)$ and any subinterval J of I such that $\xi(J) \subseteq \mathcal{D}(\kappa)$, an $\mathcal{L} \in \mathbf{LD}(\xi)$ has well defined *connection coefficients* $L_i^j : J \to \mathbb{R}$, which are measurable functions.

Let $\xi \in CRV_{LAC}(I, M)$, $\mathcal{L} \in \mathbf{LD}(\xi)$. If $v \in \Gamma^0(TM, \xi)$, $w \in \Gamma^0(T^*M, \xi)$, and v, w are a.e. differentiable, then the Lie derivatives $\mathcal{L}v$, $\mathcal{L}w$, are well defined a.e., and are given in coordinates by the obvious formulae

$$\begin{aligned} (\mathcal{L}v)^j(t) &= \dot{v}^j(t) + \sum_i L_i^j(t) v^i(t), \\ (\mathcal{L}w)_i(t) &= \dot{w}_i(t) - \sum_j L_i^j(t) w_j(t). \end{aligned} \tag{4.4}$$

Clearly, the product formula $\frac{d}{dt}(w.v) = \mathcal{L}w.v + w.\mathcal{L}v$ holds. In particular, the Lie derivatives $\mathcal{L}Y$, $\mathcal{L}\omega$ along ξ of C^1 vector fields Y and C^1 1-forms ω are well defined—via $\mathcal{L}Y \stackrel{\text{def}}{=} \mathcal{L}(Y \circ \xi)$ and $\mathcal{L}\omega \stackrel{\text{def}}{=} \mathcal{L}(\omega \circ \xi)$—as fields of vectors or covectors along ξ.

We call an $\mathcal{L} \in \mathbf{LD}(\xi)$ *locally integrable* (LI) if the scalar function $\omega.\mathcal{L}Y$ is LI for all $Y \in \Gamma^\infty(TM), \omega \in \Gamma^\infty(T^*M)$. (In view of the product rule, this is equivalent to saying that $\mathcal{L}\omega.Y$ is LI for all Y, ω, since $t \to (\omega.Y)(\xi(t))$ is LAC.) We use $\mathbf{LD}^{LI}(\xi)$ to denote the class of LI connections along ξ. It is clear that $\mathcal{L}$ is LI iff I can be covered by intervals J_α that are mapped to domains of coordinate charts relative to which the coefficients L_i^j are LI, and in that case the L_i^j are LI for *every* chart κ and every interval J such that $\xi(J) \subseteq \mathcal{D}(\kappa)$.

If $\mathcal{L} \in \mathbf{LD}^{LI}(\xi)$ and $v \in \Gamma^{LAC}(TM, \xi)$, then the components $(\mathcal{L}v)^j$ of $\mathcal{L}v$ relative to a chart are LI, so $\mathcal{L}v \in \Gamma^{LI}(TM, \xi)$. So $\mathcal{L}$ induces an $\mathbb{R}$-linear map—labeled $\mathcal{L}^{vec}$—from $\Gamma^{LAC}(TM, \xi)$ to $\Gamma^{LI}(TM, \xi)$. Moreover, given any $y \in \Gamma^{LI}(TM,\xi)$, $\bar{t} \in I$, $\bar{v} \in T_{\xi(\bar{t})}M$, there exists a unique $v \in \Gamma^{LAC}(TM,\xi)$ such that $\mathcal{L}^{vec}v = y$ and $v(\bar{t}) = \bar{v}$. (All this follows trivially if $\xi(I)$ is contained in the domain of a chart, because the system of equations $\dot{y}^j + \sum_i^j L_i^j(t)y^j(t) = v^j(t)$ is linear and the functions L_i^j are integrable, and then the conclusion can be easily globalized over all of I.) In particular, $\mathcal{L}^{vec}$ is onto, and $\dim \operatorname{Ker} \mathcal{L}^{vec} = \dim M$.

Similarly, $\mathcal{L}$ induces a surjective map

$$\mathcal{L}^{cov} : \Gamma^{LAC}(T^*M,\xi) \to \Gamma^{LI}(T^*M,\xi),$$

such that for every $z \in \Gamma^{LI}(T^*M,\xi)$, $\bar{t} \in I$, $\overline{w} \in T^*_{\xi(\bar{t})}M$, there exists a unique $\zeta \in \Gamma^{LAC}(T^*M,\xi)$ such that $\mathcal{L}^{cov}\zeta = z$ and $\zeta(\bar{t}) = \overline{w}$.

The equations $\mathcal{L}^{vec}v = y$, $\mathcal{L}^{cov}\zeta = z$ are known, respectively, as the *inhomogeneous variational equation* (IVE) and *inhomogeneous adjoint equation* (IAE) arising from ξ and $\mathcal{L}$, for given $y \in \Gamma^{LI}(TM,\xi)$ and $z \in \Gamma^{LI}(T^*M,\xi)$. (If y or z vanishes, then we just talk about the *variational equation* (VE) and *adjoint equation* (AE).) We will use $IVV(\xi,\mathcal{L},y)$, $IAV(\xi,\mathcal{L},z)$ to denote the sets of solutions of $\mathcal{L}^{vec}v = y$ and $\mathcal{L}^{cov}\zeta = z$, and write $VV(\xi,\mathcal{L}) \stackrel{\text{def}}{=} IVV(\xi,\mathcal{L},0)$, $AV(\xi,\mathcal{L}) \stackrel{\text{def}}{=} IAV(\xi,\mathcal{L},0)$. The members of $VV(\xi,\mathcal{L})$, $AV(\xi,\mathcal{L})$ are known, respectively, as *$\mathcal{L}$-variational vectors* and *$\mathcal{L}$-adjoint vectors* along ξ.

The product rule implies that, if v, ζ are solutions of $\mathcal{L}^{vec}v = y$, $\mathcal{L}^{cov}\zeta = z$, then $\frac{d}{dt}\zeta.v = \mathcal{L}\zeta.v + \zeta.\mathcal{L}v$. Since the function $\zeta.v$ is LAC, this implies the *integration by parts formula*

$$\zeta(b).v(b) - \zeta(a).v(a) = \int_a^b (\mathcal{L}\zeta.v + \zeta.\mathcal{L}v)(t)\,dt \quad \text{for } a,b \in I. \tag{4.5}$$

In particular,

(4.6) *the function $t \to \zeta(t).v(t)$ is constant if $v \in VV(\xi,\mathcal{L})$, $\zeta \in AV(\xi,\mathcal{L})$.*

One obvious way to define a connection along a LAC curve ξ is to construct a LIC[1] CVF $M \times \mathcal{D}(\xi) \ni (x,t) \to F(x,t)$, such that $\dot{\xi}(t) = F(\xi(t),t)$ for almost every t. Given F, we can define $\mathcal{L}^{\xi,F} \in \mathbf{LD}^{LI}(\xi)$ by $\mathcal{L}_t^{\xi,F} \stackrel{\text{def}}{=} L_{\xi(t),F_t}$, so that, if $Y \in \Gamma^1(TM)$, then

$$\mathcal{L}_t^{\xi,F}(Y) = [F_t, Y](\xi(t)) \quad \text{for a.e. } t \in \mathcal{D}(\xi). \tag{4.7}$$

Relative to a chart κ as above, we have

$$(4.8)\qquad (\mathcal{L}_t^{\xi,F}(Y))^j = \frac{d}{dt}(Y^j\circ\xi)(t) - \sum_{i=1}^n \left(\frac{\partial F_t^j}{\partial x^i}Y^i\right)(\xi(t)) \text{ for a.e. } t.$$

In the terminology of Section 3, constructing such a field F amounts to realizing ξ as the first component of a C^1-*controlled curve* (ξ, F). It is easy to see — using (4.4) and (4.8)— that the coefficients L_i^j fo $\mathcal{L}^{\xi,F}$ are given by $L_i^j = -\frac{\partial F_t^i}{\partial x^j}\circ\xi$. In particular, if F is LIC^1 then $\mathcal{L}^{\xi,F}\in\mathbf{LD}^{LI}(\xi)$. So *every LIC^1 CVF F such that $\xi\in\mathrm{IC}(F)$ gives rise to a LI Lie differentiation $\mathcal{L}^{\xi,F}$ along ξ*. Equivalently, *to every LIC^1-controlled curve $\gamma=(\xi,F)$ there corresponds in a canonical way a LI Lie differentiation $\mathcal{L}^{\xi,F}$ —or $\mathcal{L}^\gamma$— along ξ*. It is easy to show that, conversely, every LI connection along a LAC curve arises from some F, and F can in fact be taken to be LIC$^\infty$:

Proposition 4.1. *Let $\xi: I\to M$ be a LAC curve, and let $\mathcal{L}$ be a LI connection along ξ. Then $\mathcal{L}=\mathcal{L}^{\xi,F}$ for some LIC$^\infty$ control vector field F with domain $M\times I$.*

Proof. It clearly suffices to assume that ξ is contained in the domain Ω of a chart $\kappa=(x^1,\ldots,x^n)$, and I is compact. Let $L_i^j: I\to\mathbb{R}$ be the coefficients of $\mathcal{L}$ relative to κ. Let $\widetilde{F}_t(x)=\sum_{j=1}^n\Big(\dot{\xi}^j(t)-\sum_{i=1}^n L_i^j(t)(x^i-\xi^i(t))\Big)\partial_j^\kappa(x)$. Then $\widetilde{F}$ satisfies the desired condition, except for the fact that the $\widetilde{F}_t$ are only defined on Ω. To get a globally defined F, multiply $\widetilde{F}$ by a $\varphi\in C^\infty(M)$ that equals 1 near $\xi(I)$ and is such that support$(\varphi)\subseteq\Omega$. ∎

Let $\gamma=(\xi,F)$ be a LIC^1-controlled curve. Suppose $v\in\Gamma^{LAC}(TM,\xi)$, and φ is a function of class C^1 on some open set that contains $\xi(J)$. Let $t\in J$, $p=\xi(t)$. Then $(\mathcal{L}_t^\gamma v)\varphi = d\varphi(p).\mathcal{L}_t^\gamma v = \frac{d}{ds}\Big|_{s=t} d\varphi(\xi(s)).v(s) - (\mathcal{L}_t^\gamma d\varphi).v(t)$. On the other hand, if X is a smooth vector field, then

$$\begin{aligned}(\mathcal{L}_t^\gamma d\varphi).X(p) &= F_t(p)(d\varphi.X) - d\varphi(p).\mathcal{L}_t^\gamma X\\ &= F_t(p)(X\varphi) - d\varphi(p).[F_t,X](p)\\ &= (F_tX\varphi)(p) - ([F_t,X]\varphi)(p)\\ &= (XF_t\varphi)(p) = d(F_t\varphi)(p).X(p).\end{aligned}$$

Therefore $\mathcal{L}_t^\gamma d\varphi = d(F_t\varphi)(p)$, and then

$$(\mathcal{L}_t^\gamma v)\varphi = \frac{d}{ds}\Big|_{s=t} v(s)\varphi - v(t)(F_t\varphi). \tag{4.9}$$

So, if $\overline{y} \in T_pM$, and Φ is any set of C^1 functions φ near $\xi(t)$ whose differentials at $\xi(t)$ span $T^*_{\xi(t)}M$, we have

$$\mathcal{L}_t^\gamma v = \overline{y} \text{ iff } \frac{d}{ds}\Big|_{s=t} v(s)\varphi = v(t)(F_t\varphi) + \overline{y}\varphi \text{ for all } \varphi \in \Phi. \tag{4.10}$$

In particular, if $\kappa = (x^1, \ldots, x^n)$ is a chart near $\xi(t)$, and we apply the above with $\varphi = x^j$, we get the familiar formulation of the IVE in coordinates:

$$\dot{v}^j(t) = \sum_i \frac{\partial F_t^j}{\partial x^i}(\xi(t))v^i + \overline{y}^j. \tag{4.11}$$

Similarly, if $\zeta \in \Gamma^{LAC}(T^*M, \xi)$, and X is a vector field of class C^1 near $\xi(t)$, we have

$$\begin{aligned}(\mathcal{L}_t^\gamma\zeta).X(p) &= \frac{d}{ds}\Big|_{s=t} \zeta(s).X(\xi(s)) - \zeta(t).(\mathcal{L}_t^\gamma X) \\ &= \frac{d}{ds}\Big|_{s=t} \zeta(s).X(\xi(s)) - \zeta(t).[F_t, X](\xi(t)).\end{aligned} \tag{4.12}$$

So, if $\overline{z} \in T_p^*M$, and Ψ is any set of C^1 vector fields X near $\xi(t)$ whose values at $\xi(t)$ span $T_{\xi(t)}M$, we have

$$\begin{aligned}\mathcal{L}_t^\gamma\zeta = \overline{z} \text{ iff } &\frac{d}{ds}\Big|_{s=t} \zeta(s).X(\xi(s)) \\ &= \zeta(t).[F_t, X](\xi(t)) + \overline{z}.X(\xi(t)) \quad \text{ for all } X \in \Psi.\end{aligned} \tag{4.13}$$

In particular, for a chart κ near $\xi(t)$, by taking $X = \frac{\partial}{\partial x^i}$, we get the coordinate formulation of the IAE:

$$\dot{\zeta}_i(t) = -\sum_j \zeta_j(t)\frac{\partial F_t^j}{\partial x^i}(\xi(t)) + \overline{z}_i. \tag{4.14}$$

The preceding remarks contain the proof of most of the following:

Proposition 4.2. *Let $\gamma = (\xi, F)$ be a LIC^1-controlled curve in a manifold M, and let $L \in CF_{LIC^1}(M)$ be such that $\mathcal{D}(\gamma) \subseteq \mathcal{TD}(L)$.*

Let $\zeta \in \Gamma^{LAC}(T^*M,\xi)$. *Let* $\omega(t) = dL_t(\xi(t))$. *Write* $\Xi(t) = (\xi(t),\zeta(t))$, $I = \mathcal{D}(\xi)$. *Let* Φ *be a set of functions of class* C^1 *near* $\Xi(I)$, *such that* $\{dK(\Xi(t)) : K \in \Phi\}$ *spans* $T^*_{\Xi(t)}T^*M$ *for every* $t \in I$. *Let* Ψ *be a set of vector fields of class* C^1 *near* $\xi(I)$, *such that* $\{X(\xi(t)) : X \in \Psi\}$ *spans* $T_{\xi(t)}M$ *for every* $t \in I$. *Then the following conditions are equivalent*:

(i) ζ *satisfies the IAE* $\mathcal{L}^\gamma\zeta + \omega = 0$,

(ii) $\frac{d}{dt}\zeta(t).X(\xi(t)) = \zeta(t).[F_t,X](\xi(t)) - XL_t(\xi(t))$ *a.e. for all vector fields* X *that are of class* C^1 *on some open set containing* $\xi(I)$,

(iii) $\frac{d}{dt}\zeta(t).X(\xi(t)) = \zeta(t).[F_t,X](\xi(t)) - XL_t(\xi(t))$ *a.e. for all* $X \in \Psi$,

(iv) Ξ *is an integral curve of* $\vec{H}_{F,L}$,

(v) $\frac{d}{dt}K(\Xi(t)) = \{H_{F_t,L_t},K\}(\Xi(t))$ *a.e. for all functions* K *that are of class* C^1 *on a neighborhood of* Ξ *in* T^*M,

(vi) $\frac{d}{dt}K(\Xi(t)) = \{H_{F_t,L_t},K\}(\Xi(t))$ *a.e. for all* $K \in \Phi$.

Proof. The equivalence of (i), (ii) and (iii) has already been proved. The equivalence of (i) and (iv) follows by comparing the coordinate expression (4.14) of the IAE (with $\bar{z}_i = -\frac{\partial L}{\partial x^i}(\xi(t)))$ with the formula (3.1) for the IC's of $\vec{H}_{F,L}$. The implication (iv)$\Rightarrow$(v) is a trivial consequence of the definition of the Poisson bracket, and (v)$\Rightarrow$(vi) is trivial. Finally, (vi)$\Rightarrow$(iv) because, if (vi) holds and $t \in I$ is such that $\frac{d}{dt}K(\Xi(t)) = \{H_{F_t,L_t},K\}(\Xi(t))$ for all $K \in \Phi$, and Ξ is differentiable at t, then $dK(\Xi(t)).\dot{\Xi}(t) = dK(\Xi(t)).\vec{H}_{F,L}(\Xi(t))$ for all $K \in \Psi$, so $\dot{\Xi}(t) = \vec{H}_{F,L}(\Xi(t))$. ∎

Definition 4.2. If M is a manifold, $\gamma = (\xi,F)$ is a LIC^1-controlled curve in M, and $L \in CF(M)$ is LIC^1 and such that $\mathcal{D}(\xi) \subseteq \mathcal{TD}(L)$, then a solution ζ of the IAE $\mathcal{L}^\gamma\zeta + \omega = 0$—where $\omega(t) = dL_t(\xi(t))$—is called an L-*adjoint vector along* γ. If $L = 0$ we just call ζ an *adjoint vector along* γ. ∎

We use $\mathrm{Adj}(\xi,F,L)$, or $\mathrm{Adj}(\gamma,L)$, to denote the set of all L-adjoint vectors along γ. If $L = 0$, we just write $\mathrm{Adj}(\xi,F)$, or $\mathrm{Adj}(\gamma)$. If $n = \dim M$, then it is clear that $\mathrm{Adj}(\gamma,L)$ is an n-dimensional affine space, and $\mathrm{Adj}(\gamma)$ is linear.

We stress that we have *not* defined "adjoint vectors along a curve" ξ in

M. The definition of an adjoint vector makes sense for a LIC^1-controlled curve. If $\gamma = (\xi, F)$, then the space $\mathrm{Adj}(\gamma)$ will in general depend on F and not on ξ alone.

Of the various equivalent conditions of Proposition 4.2, the most important one is (ii), so we will reformulate in terms of "generalized switching functions." The *generalized switching function* $\varphi_{X,\gamma,\zeta}$ associated with a vector field X along (γ, ζ) is the real-valued function on $\mathcal{D}(\gamma)$ given by $\varphi_{X,\gamma,\zeta}(t) \stackrel{\text{def}}{=} \langle \zeta(t), X(\xi(t)) \rangle$. The formula of Condition (ii) then says that

$$\dot{\varphi}_{X,\gamma,\zeta}(t) = \varphi_{[F_t,X],\gamma,\zeta}(t) - XL_t(\xi(t)) \quad \text{for a.e. } t \in \mathcal{D}(\xi). \tag{4.15}$$

Therefore:

AV If $\gamma = (\xi, F)$ is a LIC^1-controlled curve in M, $L \in CF_{C^1}(M)$, $I = \mathcal{D}(\xi) \subseteq \mathcal{TD}(L)$, and ζ is a LAC field of covectors along γ, then $\zeta \in \mathrm{Adj}(\gamma, L)$ *iff* (4.15) *holds for every vector field X of class C^1 on a neighborhood of ξ.*

We conclude this section by describing the generalization of the AVE to the case of a LIL-controlled curve.

If M and N are smooth manifolds, and $\psi : M \to N$ is a locally Lipschitz map, then the well-known Rademacher Theorem says that ψ is differentiable almost everywhere. If we let $\mathrm{Diff}(\psi)$ be the set of points where ψ is differentiable, then $D\psi(x)$ is a well defined linear map from T_xM to $T_{\psi(x)}N$ for each $x \in \mathrm{Diff}(\psi)$. The graph $G(D(\psi))$ of $D\psi$ is the set $\{(x, \psi(x), D\psi(x)) : x \in \mathrm{Diff}(M)\}$, which is a subset of the bundle $J^1C^\infty(M,N)$ of 1-jets of smooth maps from M to N. (If we use $\mathrm{Hom}(X,Y)$ to denote the set of linear maps from X to Y, whenever X, Y are linear spaces, then an element q of $J^1C^\infty(M,N)$ can be thought of in a natural way as a triple (x,y,L), where $x \in M$, $y \in N$, and $L \in \mathrm{Hom}(T_xM, T_yN)$. So $J^1C^\infty(M,N)$ can be identified with the bundle $\mathrm{Hom}(TM,TN)$ over $M \times N$ whose fiber at (x,y) is $\mathrm{Hom}(T_xM, T_yN)$.)

The *Clarke generalized Jacobian* of a locally Lipschitz map $\psi : M \to N$ at a point $x \in M$ is the subset $\partial\psi(x)$ of $\mathrm{Hom}(T_xM, T_{\psi(x)}N)$ which is the convex hull of the set of all limits $\lim_{j\to\infty} D\psi(x_j)$, ranging over all sequences $\{x_j\}$ in $\mathrm{Diff}(\psi)$ such that $x_j \to x$ and $\lim_{j\to\infty} D\psi(x_j)$ exists. It is clear that $\partial\psi(x)$ is compact and convex.

In the special case when $N = TM$ and ψ is a locally Lipschitz vector field on M, then each $D\psi(x)$, $x \in \mathrm{Diff}(\psi)$, can be regarded as a 1-jet of vector fields on M, i.e. as a member of $\mathbf{LD}(x)$, so $\partial\psi(x)$ is a compact convex subset of $\mathbf{LD}(x)$.

Definition 4.3. Let $\gamma = (\xi, F)$ be a LIL-controlled curve in M. A *connection* associated with γ is a measurable map $\mathcal{D}(\xi) \ni t \to \mathcal{L}_t \in \mathbf{LD}(\xi(t))$ such that $\mathcal{L}_t \in \partial F_t(\xi(t))$ for a.e. $t \in \mathcal{D}(\xi)$. We write $\boldsymbol{\mathcal{L}}^\gamma$ to denote the set of all connections along γ. ■

The LIL analogue of the IAE is the "inhomogeneous adjoint inclusion" (abbr. IAI), defined as follows. Suppose $\gamma = (\xi, F)$ is a LIL-controlled curve in M, and $L \in CF(M)$ is LIL and such that $I = \mathcal{D}(\gamma) \subseteq \mathcal{TD}(L)$. Then we can define the "augmented" map $F^L : M \times \mathcal{TD}(F, L) \to TM \times \mathbb{R}$ given by $F^L(x,t) = (F(x,t), L(x,t))$. For fixed t, the map F_t^L is Lipschitz, so it has a well defined Clarke generalized Jacobian ∂F_t^L. For each $x \in M$, $\partial F_t^L(x)$ is a compact convex subset of the product $\mathbf{LD}(x) \times T_x^*M$. We will use $\boldsymbol{\mathcal{L}}^{\gamma,L}$ to denote the set of all measurable selections of $\partial F^L \circ \gamma$, so the members of $\boldsymbol{\mathcal{L}}^{\gamma,L}$ are the measurable maps $I \ni t \to (\mathcal{L}_t, \omega_t) \in \mathbf{LD}(\xi(t)) \times T^*_{\xi(t)}M$ such that $(\mathcal{L}_t, \omega_t) \in \partial F_t^L(\xi(t))$ for almost all $t \in I$.

Definition 4.4. If M is a manifold, $\gamma = (\xi, F)$ is a LIL-controlled curve in M, and $L \in CF(M)$ is LIL and such that $\mathcal{D}(\xi) \subseteq \mathcal{TD}(L)$, then a *solution of the IAI* for γ and L is a LAC field of covectors ζ along ξ such that

[IAIS] *there exists a selection* $(\mathcal{L}, \omega) \in \boldsymbol{\mathcal{L}}^{\gamma,L}$ *for which*

$$-(\mathcal{L}\zeta)(t) = \omega_t(\xi(t)) \qquad \text{for a.e. } t \in \mathcal{D}(\xi).$$

A solution ζ of the IAI for γ and L will be called an *L-adjoint vector along* γ. If $L = 0$ we just call ζ a *solution of the AI* (adjoint inclusion) or an *adjoint vector* along γ. We use $\mathrm{Adj}(\xi, F, L)$, or $\mathrm{Adj}(\gamma, L)$, to denote the set of all L-adjoint vectors along γ. If $L = 0$, we just write $\mathrm{Adj}(\xi, F)$, or $\mathrm{Adj}(\gamma)$. ■

Remark 4.1. If ζ is a solution of the IAI for γ and L, corresponding to a selection $(\mathcal{L},\omega) \in \boldsymbol{\mathcal{L}}^{\gamma,L}$, then in particular $\mathcal{L} \in \boldsymbol{\mathcal{L}}^{\gamma}$ and ω is a selection of $\partial L \circ \gamma$. So

[IAIS*] $\quad$ *ζ satisfies the equation $-\mathcal{L}_t\zeta \in \partial L_t$ a.e. for some $(\mathcal{L},\omega) \in \boldsymbol{\mathcal{L}}^{\gamma,L}$.*

Notice, however, that [IAIS*] is in general strictly weaker than [IAIS], since $\partial F_t^L(x) \subseteq \partial F_t(x) \times \partial L_t(x)$, but the inclusion need not be an equality. ∎

5. Control Systems, Reachable Sets, Lagrangians

Definition 5.1. A *control system* is a triple $\Sigma = (M,\mathcal{U},F)$ such that (i) M is a manifold, (ii) $\mathcal{U}$ is a set, and (iii) $F = \{F_\eta\}_{\eta\in\mathcal{U}}$ is a family of control vector fields on M. The manifold M is the *state space* of Σ, $\mathcal{U}$ is the *open-loop control space*, the elements of $\mathcal{U}$ are the *open-loop controls* or, simply, *controls*, and F is the *dynamics*. ∎

If $\Sigma = (M,\mathcal{U},F)$ is a control system, we use $\mathcal{F}(\Sigma)$, $\widehat{\mathcal{F}}(\Sigma)$ to denote, respectively, the sets $\{F_\eta : \eta \in \mathcal{U}\}$ and $\{\widehat{F} : (\exists F \in \mathcal{F}(\Sigma))(\exists J \in \mathcal{I})(J \subseteq \mathcal{TD}(F) \wedge \widehat{F} = F\lceil M \times J)\}$. We say that a CVF $G \in CVF(M)$ *occurs* in Σ if $G \in \widehat{\mathcal{F}}(\Sigma)$. If $\widehat{\mathcal{U}}$ is any subset of $\mathcal{U}$, we will use $\Sigma\lceil\widehat{\mathcal{U}}$ to denote the *restriction* of Σ to $\widehat{\mathcal{U}}$, i.e. the system $(M,\widehat{\mathcal{U}},F\lceil\widehat{\mathcal{U}})$, where $F\lceil\widehat{\mathcal{U}} = \{F_\eta : \eta \in \widehat{\mathcal{U}}\}$. If $\mathcal{P}$ is any property of CVF's (for example, LIB, LIBU, LIL, or LIC^k), then a control $\eta \in \mathcal{U}$ will be called *$\mathcal{P}$-admissible* (for example, *LIB-admissible*) for Σ if the corresponding CVF F_η satisfies $\mathcal{P}$. We will use $\mathcal{U}_{\mathcal{P},F}$, or $\mathcal{U}_{\mathcal{P},\Sigma}$, to denote the class of all $\mathcal{P}$-admissible $\eta \in \mathcal{U}$, and write $\Sigma\lceil\mathcal{P}$ for $\Sigma\lceil\mathcal{U}_{\mathcal{P},\Sigma}$.

Definition 5.2. A *controlled trajectory* of a system $\Sigma = (M,\mathcal{U},F)$ is a pair $\gamma = (\xi,\eta)$ such that $\eta \in \mathcal{U}$ and (ξ,F_η) is a controlled curve. If $\mathcal{D}(\xi)$ is compact, then γ is a *controlled arc* of Σ. The *domain* $\mathcal{D}(\gamma)$ of γ is the domain of ξ. We use $\mathrm{Ctraj}(\Sigma)$ (resp. $\mathrm{Carc}(\Sigma)$) to denote the set of all controlled trajectories (resp. controlled arcs) of Σ. A *trajectory* (resp. *arc*) of Σ is a ξ such that $(\xi,\eta) \in \mathrm{Ctraj}(\Sigma)$ (resp. $(\xi,\eta) \in \mathrm{Carc}(\Sigma)$) for some $\eta \in \mathcal{U}$. If $\mathcal{P}$ is any property of CVF's, such as LIB or LIL or LIC^1, a $\gamma = (\xi,\eta) \in \mathrm{Ctraj}(\Sigma)$ is *$\mathcal{P}$-controlled* if η is $\mathcal{P}$-admissible. We use $\mathrm{Ctraj}_{\mathcal{P}}(\Sigma)$

(resp. $\mathrm{Carc}_{\mathcal{P}}(\Sigma)$) to denote the set of all $\mathcal{P}$-controlled trajectories (resp. $\mathcal{P}$-controlled arcs) of Σ. ■

Often, we will use the self-explanatory expression "the control system $\dot{x} = F_\eta(x,t)$, $x \in M$, $\eta \in \mathcal{U}$" as a name for the control system $(M, \mathcal{U}, F)$.

Definition 5.3. Let $\Sigma = (M, \mathcal{U}, F)$ be a control system and let $a, b \in \mathbb{R}$, $a \leq b$. If $p, x \in M$, we say that x is *Σ-reachable from p over $[a,b]$* if there exists a controlled arc (ξ, η) of Σ such that $\xi(a) = p$ and $\xi(b) = x$. We say that x is *Σ-reachable from p* if it is Σ-reachable from p over a, b for some a, b. We define $\mathcal{R}^\Sigma(p)$ (resp. $\mathcal{R}^\Sigma_{[a,b]}(p)$), the *$\Sigma$-reachable set from p* (resp. the *Σ-reachable set from p over $[a,b]$*) to be the set of all points $x \in M$ that are Σ-reachable from p (resp. Σ-reachable from p over $[a,b]$). ■

Definition 5.4. If $\Sigma = (M, \mathcal{U}, F)$ is a control system, a *Lagrangian* for Σ is a family $\{L_\eta\}_{\eta \in \mathcal{U}}$ of Carathéodory functions on M such that $\mathcal{TD}(L_\eta) = \mathcal{TD}(F_\eta)$ for all $\eta \in \mathcal{U}$. ■

Given a Lagrangian L for $\Sigma = (M, \mathcal{U}, F)$, we would like to define the *cost functional* $J^{L,\Sigma} : \mathrm{Carc}(\Sigma) \to \mathbb{R}$ by letting $J^{L,\Sigma}(\gamma) = \int_{\mathcal{D}(\gamma)} L_\eta(\xi(t), t)\,dt$ if $\gamma = (\xi, \eta)$. This is possible provided that the integral makes sense. If we want to use Lebesgue integration, then the minimal requirement—since the integrand is guaranteed to be measurable—is that the function $\psi_{\gamma,L} : t \to L_\eta(\xi(t), t)$ be bounded below or above by an integrable function. (In that case, the integral makes sense, although its value could be $+\infty$ or $-\infty$. [5]) Let us call a $\gamma \in \mathrm{Ctraj}(\Sigma)$ *weakly acceptable for L* if at least one of the functions $\psi^+_{\gamma,L} : t \to \max(L_\eta(\xi(t), t), 0)$, $\psi^-_{\gamma,L} : t \to -\min(L_\eta(\xi(t), t), 0)$ is locally integrable, and *acceptable for L* if $\psi_{\gamma,L}$ is locally integrable. If γ is weakly acceptable for L and has domain $[a,b]$, then the function $\rho_{\gamma,L} : [a,b] \to \overline{\mathbb{R}}$ given by $\rho_{\gamma,L}(t) = \int_a^t L_\eta(\xi(s), s)\,ds$ is the *running cost* along γ. (More generally, any function $\rho : \mathcal{D}(\gamma) \to \overline{\mathbb{R}}$ such that $\rho(t) \equiv \int_a^t L_\eta(\gamma(s), s)\,ds + r$ for some finite constant r and some $a \in \mathcal{D}(\gamma)$ will be called a *running cost* for γ.) We define $J^{L,\Sigma}(\gamma) = \rho_{\gamma,L}(b)$. So $J^{L,\Sigma}$

[5] Naturally, one could also allow integrability in the sense of some more general conditionally convergent integration theory, but this possibility appears to be without interest and will not be pursued here.

is well defined as a map from $\mathrm{Carc}_{wa}(\Sigma, L)$ to $\overline{\mathbb{R}}$, where $\mathrm{Carc}_{wa}(\Sigma, L)$ is the set of all $\gamma \in \mathrm{Carc}(\Sigma)$ that are weakly acceptable for L.

The *augmented system* Σ^L associated with a system $\Sigma = (M, \mathcal{U}, F)$ and a Lagrangian $L = \{L_\eta : \eta \in \mathcal{U}\}$ is the system

$$\dot{x} = F_\eta(x,t), \qquad \dot{x}^0 = L_\eta(x,t), \qquad \eta \in \mathcal{U}, \tag{5.1}$$

in which "the running cost x^0 has been added as a new state variable." That is, Σ^L is the triple $(M \times \mathbb{R}, \mathcal{U}, F^L)$ where, for $\eta \in \mathcal{U}$, $F^L_\eta \in CVF(M \times \mathbb{R})$ is given by $F^L_\eta(x, x^0, t) = (F_\eta(x,t), L_\eta(x,t))$, with domain $\{(x, x^0, t) : x \in M,\ x^0 \in \mathbb{R},\ t \in \mathcal{TD}(F_\eta)\}$. According to our general definitions, a trajectory $\widehat{\xi}$ of the augmented system for a control $\eta \in \mathcal{U}$ is a pair (ξ, ξ^0) such that ξ is a trajectory of Σ for η, (ξ, η) is acceptable for L, and ξ^0 is a running cost for (ξ, η). In particular, a controlled arc $\gamma = (\xi, \eta)$ of Σ may fail to give rise to controlled arcs (ξ, ξ^0, η) of Σ^L, since γ may fail to be acceptable for L.

We now define the key substitution properties needed to construct "control variations," as required by the proof of the MP sketched in [22] and [23].

If $F \in CVF(M)$ and J is a subinterval of $\mathcal{TD}(F)$, then $F\lceil J$ will denote the restriction of F to $M \times J$, so $\mathcal{TD}(F\lceil J) = J$.

Definition 5.5. Let F, G be CVF's on M, and let $E \subseteq \mathbb{R}$. Then

1. $SUB(E, F, G)$ (i.e. "the CVF obtained from F by substituting G over E") is defined if $E \subseteq \mathcal{TD}(F, G)$. In that case, $SUB(E, F, G)$ has time domain $\mathcal{TD}(F)$, and

$$SUB(E,F,G)(x,t) = \chi_E(t)G(x,t) + (1 - \chi_E(t))F(x,t) \quad \text{for } (x,t) \in \mathcal{D}(F). \tag{5.2}$$

2. Assume that $E = [a, b]$ is a subinterval of $\mathcal{TD}(F)$. Then $DEL(E, F)$ (i.e. "the CVF obtained from F by removing $F\lceil E$") is the CVF on M with time domain

$$T = ((-\infty, a] \cap \mathcal{TD}(F)) \cup \{t : t \geq a,\ t + b - a \in \mathcal{TD}(F)\},$$

such that, for $(x,t) \in M \times T$,

$$DEL(E,F)(x,t) \stackrel{\text{def}}{=} \chi_{(-\infty,a]}(t)F(x,t) + (1-\chi_{(a,\infty)}(t))F(x,t+b-a) \tag{5.3}$$

3. Assume that $E = [a,b]$ is a compact interval, and let $r = b - a$. Then $INS(E,F,G)$ (i.e. "the CVF obtained from F by inserting G over E") is defined if $a \in \mathcal{TD}(F)$ and $E \subseteq \mathcal{TD}(G)$. In that case, the time domain of $INS(E,F,G)$ is

$$T = ((-\infty,a] \cap \mathcal{TD}(F)) \cup [a,b] \cup \{t : t \geq b,\ t-(b-a) \in \mathcal{TD}(F)\} \tag{5.4}$$

and the value of $INS(E,F,G)$ at $(x,t) \in M \times T$ is

$$INS(E,F,G)(x,t) \stackrel{\text{def}}{=} \chi_{(-\infty,a]}(t)F(x,t) + \chi_{(a,b]}(t)G(x,t) + \chi_{(b,\infty]}(t)F(x,t-r). \ \blacksquare$$

Definition 5.6. Let $\Sigma = (M, \mathcal{U}, F)$ be a control system. Then

1. Σ is *closed under substitutions* over members of a class $\mathcal{S}$ of subsets of the real line if, whenever η_1, η_2 belong to $\mathcal{U}$, $E \in \mathcal{S}$, and $SUB(E, F_{\eta_1}, F_{\eta_2})$ is defined, it follows that $SUB(E, F_{\eta_1}, F_{\eta_2})$ occurs in Σ.
2. If Σ is closed under substitutions over members of the class of all compact subsets (resp. measurable subsets, subintervals, compact subintervals) of $\mathbb{R}$, then Σ will be said to have the *compact* (resp. *measurable, interval, compact interval*) *substitution property*.
3. Σ is *closed under interval deletions* if, whenever $\eta \in \mathcal{U}$ and E is a compact subinterval of $\mathcal{TD}(F_\eta)$, it follows that $DEL(E, F_\eta)$ occurs in Σ.
4. $\Sigma = (M, \mathcal{U}, F)$ is *closed under interval insertions* if, whenever $\eta_1, \eta_2 \in \mathcal{U}$ are two controls, and E is a compact interval such that $INS(E, F_{\eta_1}, F_{\eta_2})$ is defined, it follows that $INS(E, F_{\eta_1}, F_{\eta_2})$ occurs in Σ. $\blacksquare$

6. Hamiltonian Minimization and Separability

Throughout this section, we fix a control system $\Sigma = (M, \mathcal{U}, F)$, and a Lagrangian L for Σ. We associate with the pair (Σ, L) the *Hamiltonian* $\mathcal{H}^{\Sigma,L}$, which is the family $\{H_\eta^{\Sigma,L}\}_{\eta\in\mathcal{U}}$ of Hamiltonians of all the pairs (F_η, L_η). In other words, for each $\eta \in \mathcal{U}$ we define $\mathcal{H}_\eta^{\Sigma,L} \in CF(T^*N)$ by $\mathcal{H}_\eta^{\Sigma,L}(x,z,t) \stackrel{\text{def}}{=} z.F_\eta(x,t) + L_\eta(x,t)$, so the time domain of $\mathcal{H}_\eta^{\Sigma,L}$ is $\mathcal{TD}(F_\eta)$. Alternatively, we may view $\mathcal{H}^{\Sigma,L}$ as a partially defined function on $T^*M \times \mathbb{R} \times \mathcal{U}$, with domain $\{(x,z,t,\eta) \in T^*M \times \mathbb{R} \times \mathcal{U} : t \in \mathcal{TD}(F_\eta)\}$, given by

$$\mathcal{H}^{\Sigma,L}(x,z,t,\eta) \stackrel{\text{def}}{=} z.F_\eta(x,t) + L_\eta(x,t). \tag{6.1}$$

When $L = 0$, we omit the L and write $\mathcal{H}^\Sigma$, $\mathcal{H}_\eta^\Sigma$.

If Ξ is a curve in T^*M, and $\eta \in \mathcal{U}$, we call η *weakly minimizing along* Ξ *for* (Σ, L) if (i) $\mathcal{D}(\Xi) \subseteq \mathcal{TD}(F_\eta)$, and (ii) for every $\widetilde{\eta} \in \mathcal{U}$ there is a subset $E(\widetilde{\eta})$ of $\mathcal{D}(\Xi) \cap \mathcal{TD}(F_{\tilde{\eta}})$ of Lebesgue measure zero such that $\mathcal{H}^{\Sigma,L}(\Xi(t),t,\eta) \le \mathcal{H}^{\Sigma,L}(\Xi(t),t,\widetilde{\eta})$ for all t belonging to $(\mathcal{D}(\Xi) \cap \mathcal{TD}(F_{\tilde{\eta}})) \setminus E(\widetilde{\eta})$. We call η *strongly minimizing along* ξ *for* (Σ, L) if

$$\mathcal{H}^{\Sigma,L}(x,z,t,\eta) = \min\{\mathcal{H}^{\Sigma,L}(x,z,t,\widetilde{\eta}) : \widetilde{\eta} \in \mathcal{U}\} \quad \text{for a.e. } t \in \mathcal{TD}(F_{\tilde{\eta}}) \cap \mathcal{D}(\Xi). \tag{6.2}$$

Equivalently, η is strongly minimizing along ξ for (Σ, L) iff it is weakly minimizing and, in addition, the set $E(\widetilde{\eta})$ of Condition (ii) can be chosen to be independent of $\widetilde{\eta}$. If $\gamma = (\xi, \eta)$ is a LIL-controlled trajectory of a system $\Sigma = (M, \mathcal{U}, F)$, and L is a Lagrangian for Σ, then a field of covectors ζ along γ will be called *weakly* (resp. *strongly*) *minimizing* for γ for (Σ, L), if η is weakly (resp. strongly) minimizing for (Σ, L) along the curve $\Xi = (\xi, \zeta)$.

We write $\mathrm{Adj}^\Sigma(\gamma, L) \stackrel{\text{def}}{=} \mathrm{Adj}(\xi, F_\eta, L_\eta)$ if $\Sigma = (M, \mathcal{U}, F)$ is a control system, L is a Lagrangian for Σ, and $\gamma = (\xi, \eta) \in \mathrm{Ctraj}_{LIL}(\Sigma)$. (The elements of $\mathrm{Adj}^\Sigma(\gamma, L)$ are called Σ-L-*adjoint vectors along* γ, or simply L-*adjoint vectors along* γ if the meaning of Σ is clear from the context.)

Also, we write $\mathrm{Adj}^{\Sigma}_{wmin}(\gamma, L)$, $\mathrm{Adj}^{\Sigma}_{smin}(\gamma, L)$, to denote the classes of $\zeta \in \mathrm{Adj}^{\Sigma}(\gamma, L)$ that are weakly and strongly minimizing.

As usual, when $L = 0$ we drop the reference to L, and talk about controls η that are minimizing for Σ along a curve in T^*M, or fields of covectors that are minimizing for Σ for a controlled trajectory γ. We then write $\mathrm{Adj}^{\Sigma}_{wmin}(\gamma) \stackrel{\mathrm{def}}{=} \mathrm{Adj}^{\Sigma}_{wmin}(\gamma, 0)$ and $\mathrm{Adj}^{\Sigma}_{smin}(\gamma) \stackrel{\mathrm{def}}{=} \mathrm{Adj}^{\Sigma}_{smin}(\gamma, 0)$

The conclusion of the Maximum Principle asserts the existence of a nontrivial $\zeta \in \mathrm{Adj}^{\Sigma}_{wmin}(\gamma, L)$ along a controlled trajectory γ, having certain additional properties. It is important to know when this implies the stronger property that ζ is strongly minimizing. It turns out that the answer is affirmative whenever a very general "separability" condition holds. To make this precise, we introduce some definitions.

Definition 6.1. A control system $\Sigma = (M, \mathcal{U}, F)$ is *separable along an* arc $\xi : [a, b] \to M$ if there exists a countable subset $\mathcal{U}_0$ of $\mathcal{U}$, such that (i) $[a, b] \subseteq \mathcal{TD}(F_\eta)$ for every $\eta \in \mathcal{U}_0$, and (ii) for every $t \in [a, b]$, the set of values $\{F_\eta(\xi(t), t) : \eta \in \mathcal{U}_0\}$ is dense in $\{F_\eta(\xi(t), t) : \eta \in \mathcal{U},\ [a, b] \subseteq \mathcal{TD}(F_\eta)\}$. The system Σ is *separable* if it is separable along every trajectory of Σ. ∎

The following observation is completely trivial, but crucial for the passage from weak to strong minimization:

Proposition 6.1. *For a separable system, weak and strong Hamiltonian minimization are equivalent.* ∎

7. Separation, Approximating Cones, Transversality

Definition 7.1. Let S_1, S_2 be subsets of a topological space A, and let $q \in A$. We say that S_1 and S_2 are *separated at* q if $S_1 \cap S_2 = \{q\}$, and that S_1 and S_2 are *locally separated at* q if there is a neighborhood U of q such that $S_1 \cap S_2 \cap U = \{q\}$. ∎

Recall that a *cone* in a real linear space V is a nonempty subset C of V such that $v \in C$, $r \geq 0$ implies $rv \in C$.

Definition 7.2. Let S be a subset of a manifold M, and let $x \in S$. A *strongly approximating cone* of S at x is a closed convex cone C in the tangent space T_xM with the property that there exist a neighborhood N of 0 in T_xM and a continuous map $F : N \cap C \to S$ such that

$$F(v) = x + v + o(\|v\|) \text{ as } v \to 0 \text{ through values in } N \cap C. \tag{7.1}$$

(Formula (7.1) has a clear meaning in local coordinates, and it is easy to see that it is in fact independent of the choice of coordinates.)

An *approximating cone* of S at x is a closed convex cone which is the closure of an increasing union of strongly approximating cones. ∎

It is easy to verify that, if S is a submanifold of M, then the ordinary tangent space T_xS is an approximating cone of S at x. If S is a convex set near x relative to some chart κ (i.e. if S is such that $\kappa(S)$ is convex for some chart κ near x), then the set of limiting directions of S at x is an approximating cone. (A *limiting direction of* S at $x \in S$ is a vector $v \in T_xM$ such that either $v = 0$ or there exists a sequence $\{x_k\}_{k=1}^{\infty}$ of points of S and a sequence of real numbers $h_k > 0$ that satisfy $x_k \neq x$, $h_k \to 0$, and $x_k = x + h_k v + o(h_k)$ as $k \to \infty$. This in principle depends on the choice of a chart κ, but is easily shown to be in fact independent of κ.) It is easy to see that every vector v belonging to an approximating cone of S at x must be a limiting direction of S at x. Let us call a set S *AC-regular* at x if the set of all limiting directions of S at x is an approximating cone of S at x. Then the preceding examples simply say that every set which is convex relative to some chart is AC-regular at x.

Definition 7.3. If V is a finite-dimensional real linear space, a vector $z \in V^*$ is *transversal*—or *polar*— to a subset $C \subseteq V$ if $\langle z, v\rangle \leq 0$ for all $v \in C$. We use $C^\dagger$ to denote the set of all vectors polar to C. ∎

It is clear that $C^\dagger$ is always a closed convex cone. If C itself is a closed convex cone, then it is well known that $C^{\dagger\dagger} = C$.

8. The Maximum Principle

We are now ready to state the MP. We will first present the *separation* *theorem* that constitutes the most basic form of the principle, and then show how other results—such as necessary conditions for controllability and optimality—follow from it.

In its simplest form, the separation theorem gives a necessary condition for a "reference trajectory" $\xi_* : [a,b] \to M$ of a control system $\Sigma = (M, \mathcal{U}, F)$ to be such that, if $p_* = \xi_*(a)$ and $q_* = \xi_*(b)$, then the reachable set $\mathcal{R}$ from p_* is locally separated from a set S at q. However, it will be convenient to formulate the basic separation result in a slightly more general form, by considering in addition a map g from a neighborhood U of q_* to some other manifold N, and taking S to be a subset of N, in which case the result becomes a necessary condition for $g(\mathcal{R})$ to be locally separated from S at $g(q_*)$. Moreover, we will consider together the case when $\mathcal{R}$ is the "fixed time interval" reachable set $\mathcal{R}^{\Sigma}_{[a,b]}(p)$ and that when $\mathcal{R}$ is the "variable time interval" reachable set $\mathcal{R}^{\Sigma}(p)$. The necessary condition is given in terms of ξ_* *together with a control* η_* *that generates* ξ_*, and says that there must exist a field of covectors ζ along ξ_* that satisfies four requirements, namely: (i) the "adjoint equation," which says that ζ is an adjoint vector along the controlled trajectory $\gamma_* = (\xi_*, \eta_*)$, (ii) the *minimization condition*, that says that ζ is weakly minimizing (so that, if Σ is separable, then ζ is actually strongly minimizing), (iii) *nontriviality*, that is, $\zeta \neq 0$ and, finally, (iv) *transversality*, which is the condition that relates ζ to the set S. In addition, there are some minor variations. For example, when the system is closed under deletions and insertions, we can add the requirement that the Hamiltonian $\langle \zeta(t), F_{\eta_*}(\xi(t), t)\rangle$ be constant as a function of t and, in the variable time-interval case, this constant has to be 0.

From the separation theorem one can easily deduce necessary conditions for *noncontrollability along* ξ_* (i.e. for the point q to lie on the boundary of $\mathcal{R}$), and for *optimality*, i.e. for γ_* to minimize some Lagrangian cost functional $J^{L,\Sigma}$—or, more generally, $J^{L,\Sigma,\varphi}$, cf. below—among all trajectories ξ that satisfy an endpoint condition $\partial\xi \in S$, where S is a given subset of $M \times M$. The noncontrollability condition is exactly like the condition for separation, except that in that case there is no set S, and the transversality requirement disappears. (Again, when the system is closed under deletions and insertions one gets a constant Hamiltonian, and if the time interval is free then the constant must vanish.) The optimality condition is exactly like the one for separation, except that now one has to use the Hamiltonian $\mathcal{H}^{\Sigma,\nu L}$for some $\nu \geq 0$.

In order to state the Maximum Principle, we make the following assumptions:

A1. $\Sigma = (M,\mathcal{U},F)$ is a control system;

A2. either (i) Σ has the compact substitution property, or (ii) Σ has the interval substitution property and every F_η, $\eta \in \mathcal{U}$, is LIB;

A3. $a, b \in \mathbb{R}$, $a \leq b$, $p_*, q_* \in M$, $\xi_* : [a,b] \to M$ is a trajectory of Σ, $p_* = \xi_*(a)$, and $q_* = \xi_*(b)$;

A4. $\eta_* \in \mathcal{U}$ is such that $\gamma_* = (\xi_*, \eta_*)$ is a LIL-controlled arc of Σ;

A5.a. N is a smooth manifold, W is a neighborhood of q_* in M, and g is a map from W to N;

A5.b. g is Lipschitz, and $\mathbf{G} = \partial g(q_*)$ is the Clarke generalized Jacobian of g at q_*;

A6. S is a subset of N and C is an approximating cone of S at $g(q_*)$;

A7. either (i) C is not a linear subspace of $T_{g(q_*)}N$, or (ii) $C = \{0\}$ but $g(q)$ belongs to $\mathrm{Clos}(S - \{g(q_*)\})$.

We will call ξ_*, η_*, γ_* the *reference arc*, *reference control*, and *reference controlled arc*, respectively.

In addition to the basic hypotheses A1,. . .,A7, we will want to consider the special case when Σ satisfies the stronger condition

A8. Σ is closed under interval insertions and deletions.

The separation properties of interest to us are

(SEP-F) the sets $g(\mathcal{R}^{\Sigma}_{[a,b]}(p_*))$ and S are locally separated at $g(q_*)$,

(SEP-V) the sets $g(\mathcal{R}^{\Sigma}(p_*))$ and S are locally separated at $g(q_*)$.

Actually, the necessary condition for SEP-V is in fact necessary for the following weaker separation property:

(SEP-V$^{\#}$) for some $\alpha > 0$, the sets $g(\mathcal{R}^{\Sigma}_{[a,b\pm\alpha]}(p_*))$ and S are locally separated at $g(q_*)$,

where

$$\mathcal{R}^{\Sigma}_{[a,b\pm\alpha]}(p) = \bigcup\{\mathcal{R}^{\Sigma}_{[a,c]}(p) : b-\alpha \le c \le b+\alpha\}. \tag{8.1}$$

We recall that the Clarke generalized Jacobian (cf. Section 4) $\mathbf{G}$ of $g : W \to N$ at q_* is a set of linear maps from $T_{q_*}M$ to $T_{g(q_*)}N$, so every $G \in \mathbf{G}$ has a well defined transpose $G^{\dagger}$, which is a linear map from $T^*_{g(q_*)}N$ to $T^*_{q_*}M$.

Theorem 8.1 (The basic form of the Maximum Principle). *Suppose that the data Σ, M, $\mathcal{U}$, F, a, b, ξ_*, η_*, γ_*, p_*, q_*, N, W, g, $\mathbf{G}$, S and C are such that A1,…,A7 hold. Assume that (SEP-F) holds. Then there exist a $\theta \in T^*_{g(q_*)}N$, a $G \in \mathbf{G}$, and a weakly minimizing adjoint vector $\zeta \in \mathrm{Adj}^{\Sigma}(\gamma_*)$ along γ_* such that $\theta \neq 0$, θ belongs to the polar cone $C^{\dagger}$ of C, and $\zeta(b) = G^{\dagger}\theta$. If in addition A8 holds, then ζ can be chosen such that the Hamiltonian $t \to H^{\Sigma}_{\eta_*}(\xi_*(t), \zeta(t), t)$ is a.e. constant. If both A8 and (SEP-V$^{\#}$) hold, then ζ can be chosen such that $H^{\Sigma}_{\eta_*}(\xi_*(t), \zeta(t), t) = 0$ a.e.* ■

A fairly self-contained outline of the proof of Theorem 8.1 is given in [22] and [23], for the case when the state space is an open subset of a Euclidean space $\mathbb{R}^n$. However, since the result of [22] and [23] allows systems with with jumps, the manifold case is an immediate corollary, obtained by applying the Euclidean space result to a system where the jumps are coordinate changes.

A special case of Theorem 8.1 is a result on *controllability*. Let us say that a reference trajectory $\xi_* : [a,b] \to M$ is a *fixed-interval boundary trajectory* if $\xi_*(b)$ is not an interior point of the reachable set $\mathcal{R}^{\Sigma}_{[a,b]}(\xi_*(a))$. (The negation of this is *fixed-interval local controllability along* ξ_*, i.e. the property that $\xi_*(b) \in \mathrm{Int}(\mathcal{R}^{\Sigma}_{[a,b]}(\xi_*(a)))$.) Similarly, we say that ξ_* is a *variable-interval boundary trajectory* if for some $\alpha > 0$ $\xi_*(b)$ is not an interior point of the reachable set $\mathcal{R}^{\Sigma}_{[a,b\pm\alpha]}(\xi_*(a))$. (The negation of this is *variable-interval local controllability along* ξ_*.) It turns out that

Theorem 8.1 implies that a necessary condition for $\xi_* : [a,b] \to M$ to be a fixed-interval boundary trajectory is that there exist a nontrivial minimizing adjoint vector ζ, with the additional constancy of the Hamiltonian property if A8 holds, and the vanishing of the Hamiltonian if ξ_* is a variable-interval boundary trajectory. We state a slightly more general result, that contains the controllability condition as a particular case, obtained by taking $N = W = M$ and $g =$ identity.

Theorem 8.2. *Let Σ, M, $\mathcal{U}$, F, a, b, ξ_*, η_*, γ_*, p_*, q_*, N, W, g and $\mathbf{G}$ be such that A1,...,A5 hold. Assume that $g(q_*)$ is not an interior point of $g(\mathcal{R}^{\Sigma}_{[a,b]}(p_*))$. Then there exist a $\theta \in T^*_{g(q_*)}N$, a $G \in \mathbf{G}$, and a weakly minimizing adjoint vector $\zeta \in \mathrm{Adj}^{\Sigma}(\gamma_*)$ along γ_* such that $\theta \neq 0$ and $\zeta(b) = G^{\dagger}\theta$. If in addition A8 holds, then ζ can be chosen such that $t \to H^{\Sigma}_{\eta_*}(\xi_*(t), \zeta(t), t)$ is a.e. constant. If A8 holds and $g(q_*)$ is not an interior point of $g(\mathcal{R}^{\Sigma}_{[a,b\pm\alpha]}(p_*))$, then ζ can be chosen such that $H^{\Sigma}_{\eta_*}(\xi_*(t), \zeta(t), t) = 0$ a.e.*

Proof. If suffices to apply Theorem 8.1 to the case where the set S is just a sequence of points converging to $g(q_*)$ but not equal to $g(q_*)$, and $C = \{0\}$. ■

Remark 8.1. In Theorems 8.1 and 8.2, Assumption A5.b can be weakened, using the concept of a "semidifferentiable map," introduced in [22] and [23]. It suffices to assume that

A5.b$^{\#}$. g is continuous on W and semidifferentiable at q_*, and $\mathbf{G}$ is a semidifferential of g at q_*.

As explained in [22] and [23], this includes in particular two cases that are not directly covered by Theorem 8.1, namely, (a) when g is continuous near q_* and differentiable at q_*, and $\mathbf{G} = \{dg(q_*)\}$; (b) when g is Lipschitz and $\mathbf{G}$ is a Warga derivate container of g at q_*. ■

Finally, we turn to the optimality result. We now assume that we are given a collection of data Σ, M, $\mathcal{U}$, F, a, b, ξ_*, η_*, γ_*, p_*, q_* satisfying A1, A2, A3, A4 as in the previous theorems, and in addition:

OPT1.a. W is a neighborhood of (p_*, q_*) in $M \times M$, and $\varphi : W \to \mathbb{R}$ is a function,

OPT1.b. φ is Lipschitz, and $\mathbf{G} = \partial\varphi(p_*, q_*)$ is the Clarke generalized gradient of φ at (p_*, q_*),

OPT2. S is a subset of $M \times M$ such that $(p_*, q_*) \in S$, and $C \subseteq T_{p_*}M \times T_{q_*}M$ is an approximating cone of S at (p_*, q_*),

OPT3. L is a Lagrangian for Σ, and L_{η_*} is LIL.

Let us write $\mathrm{Carc}^L(\Sigma)$ to denote the set of $\gamma \in \mathrm{Carc}(\Sigma)$ that are acceptable for L.

We can then consider the *variable-interval optimal control problem*:

PROBLEM $\mathcal{P}(\Sigma, L, S, \varphi)$: *minimize the cost functional*

$$J^{L,\Sigma,\varphi}(\gamma) \stackrel{\mathrm{def}}{=} \int_a^b L_\eta(\xi(t), t)\, dt + \varphi(\xi(a), \xi(b)) \tag{8.2}$$

in the set of all triples (a, b, γ) *such that* $\gamma = (\xi, \eta) \in \mathrm{Carc}^L(\Sigma)$, $\mathcal{D}(\xi) = [a, b]$, *and* $(\xi(a), \xi(b)) \in S$,

and, for a given a, b, the fixed-interval optimal control problem:

PROBLEM $\mathcal{P}_{[a,b]}(\Sigma, L, S, \varphi)$: *minimize* $J^{L,\Sigma,\varphi}(\gamma)$ *in the set of all controlled arcs* $\gamma = (\xi, \eta) \in \mathrm{Carc}^L(\Sigma)$ *such that* $\mathcal{D}(\xi) = [a, b]$ *and* $(\xi(a), \xi(b)) \in S$.

An important particular case of variable-interval optimal control is *minimum time control*, in which the Lagrangian is equal to 1 and the function φ is zero. (Usually, such problems are considered for autonomous systems, in which case one can always take a to be 0.) We will use $\mathcal{P}_{time}(\Sigma, S)$ to denote the problem $\mathcal{P}(\Sigma, 1, S, 0)$.

To get necessary conditions for these problems, we will apply Theorem 8.1 to an auxiliary system $\widetilde{\Sigma}$. This will require that Assumption A2 be satisfied for the augmented system Σ^L rather than just for Σ, i.e. that

OPT4. Either: (i) Σ^L has the compact substitution property, or (ii) Σ^L has the interval substitution property and F_η and L_η are LIB for all $\eta \in \mathcal{U}$.

Similarly, we will need the analogue of A8 for Σ^L:

OPT5. Σ^L is closed under interval insertions and deletions.

With this terminology, we have

Theorem 8.3 (The optimal control version of the Maximum Principle). *Assume that* Σ, M, $\mathcal{U}$, F, L, γ_*, ξ_*, η_*, a, b, p_*, q_*, S, C, φ *and* $\mathbf{G}$ *are such that Hypotheses* A1, A3, A4, OPT1, OPT2, OPT3 *and* OPT4 *hold, and* γ_* *is a solution of the minimization problem* $\mathcal{P}_{[a,b]}(\Sigma, L, S, \varphi)$. *Then there exist a constant* $\nu \geq 0$ *and a* νL_{η_*}*-adjoint vector* ζ *along* γ_* *such that* (i) ζ *is weakly minimizing for* $(\Sigma, \nu L)$ *along* γ_*, (ii) ζ *satisfies the transversality condition*

$$(-\zeta(a), \zeta(b)) \in \nu\mathbf{G} + C^\dagger, \tag{8.3}$$

and (iii) *either* $\zeta(b) \neq 0$ *or* $\nu > 0$. *If* OPT5 *holds, then* ζ *and* ν *can be chosen so that the Hamiltonian* $\langle \zeta(t), F_{\eta_*}(\xi_*(t), t)\rangle + \nu L_{\eta_*}(\xi_*(t), t)$ *is a.e. constant. If OPT6 holds and* (a, b, γ_*) *is a solution of the free-interval problem* $\mathcal{P}(\Sigma, L, S, \varphi)$, *then* ζ *and* ν *can be chosen so that the constant value of the Hamiltonian is zero.*

Proof. We apply Theorem 8.1 to an auxiliary control system $\widetilde{\Sigma}$, defined on the interval $[a-1, b]$, whose state consists of a pair x_1, x_2 of variables in M, as well as a real variable y, and whose dynamical law is

$$\dot{x}_1 = \dot{x}_2 = v, \quad \dot{y} = 0, \ v \text{ arbitrary}, \qquad \text{for } a-1 \leq t < a,$$

$$\dot{x}_1 = 0, \quad \dot{x}_2 = F_\eta(x_2, t), \quad \dot{y} = L_\eta(x_2, t), \qquad \text{for } a \leq t \leq b.$$

Precisely, we define $\widetilde{M} = M \times M \times \mathbb{R}$, and $\widetilde{\mathcal{U}} = \check{\mathcal{U}} \times \mathcal{V}$, where $\check{\mathcal{U}} = \{\eta \in \mathcal{U} : [a, b] \subseteq \mathcal{TD}(F_\eta)\}$, and $\mathcal{V}$ is the set of all piecewise constant maps from the interval $[a-1, a]$ to the space $\Gamma^\infty(TM)$ of smooth vector fields on M. For $\tilde{\eta} = (\eta, V) \in \widetilde{\mathcal{U}}$, and $(x_1, x_2, y) \in \widetilde{M}$, we let $\widetilde{F}_{\tilde{\eta}}(x_1, x_2, y, , t) = (V(t)(x_1), V(t)(x_2), 0)$ for $a - 1 \leq t < a$, and define $\widetilde{F}_{\tilde{\eta}}(x_1, x_2, y, , t) = (0, F_\eta(x_2, t), L_\eta(x_2, t))$ for $a \leq t \leq b$. We then take $\widetilde{\Sigma} = (\widetilde{M}, \widetilde{\mathcal{U}}, \widetilde{F})$.

We let $\sigma_*(t) = \int_a^t L_{\eta_*}(\xi_*(s), s)ds$ for $a \le t \le b$ (so σ_* is the running cost along γ_*). We then let $\widetilde{\eta}_* = (\eta_*, 0)$, and define $\widetilde{\xi}_* : [a-1, b] \to \widetilde{M}$ by letting $\widetilde{\xi}_*(t) = (p_*, p_*, 0)$ for $a - 1 \le t < a$, and $\widetilde{\xi}_*(t) = (p_*, \xi_*(t), \sigma_*(t))$ for $a \le t \le b$. We then let $\widetilde{\gamma}_* = (\widetilde{\xi}_*, \widetilde{\eta}_*)$. Then $\widetilde{\gamma}_*$ is a controlled arc of $\widetilde{\Sigma}$, such that $\widetilde{\xi}_*(a-1) = \widetilde{\xi}_*(a) = (p_*, p_*, 0)$ and $\widetilde{\xi}_*(b) = (p_*, q_*, c)$, where $c = \sigma_*(b)$.

We let $N = \widetilde{M}$, and define g by $g(x_1, x_2, y) = (x_1, x_2, y + \varphi(x_1, x_2))$. Then g is defined on a neighborhood $\widetilde{W}$ of (p_*, q_*, c) and is Lipschitz. Pick a smooth function $\psi : M \times M \to \mathbb{R}$ such that $\psi(p_*, q_*) = 0$ and $\psi(x_1, x_2) > 0$ for $(x_1, x_2) \neq (p_*, q_*)$. Then define $c_* = c + \varphi(p_*, q_*)$, and

$$\widetilde{S} = \{(x_1, x_2, r) \in N : (x_1, x_2) \in S \wedge r \le c_* - \psi(x_1, x_2)\}. \tag{8.4}$$

Then $g(p_*, q_*, c) = (p_*, q_*, c_*) \in \widetilde{S}$.

If $(\widetilde{\xi}, \widetilde{\eta}) \in \mathrm{Carc}^L(\widetilde{\Sigma})$ is such that $\mathcal{D}(\widetilde{\xi}) = [a-1, b]$, $\widetilde{\xi}(a-1) = (p_*, p_*, 0)$, and in addition $\widetilde{\xi}(b)$ is of the form (q_1, q_2, y) with $g(q_1, q_2, y) \in \widetilde{S}$, then $(q_1, q_2) \in S$ and $y + \varphi(q_1, q_2) \le c + \varphi(p_*, q_*) - \psi(q_1, q_2)$. On the other hand, the fact that $\widetilde{\gamma}_*$ is a solution of $\mathcal{P}_{[a,b]}(\Sigma, L, S, \varphi)$ easily implies that $y + \varphi(q_1, q_2) \ge c + \varphi(p_*, q_*)$. These two inequalities imply $\psi(q_1, q_2) = 0$, so $q_1 = p_*$, $q_2 = q_*$. Then $y = c$ as well, so $g(q_1, q_2, y) = (p_*, q_*, c_*)$. Therefore the sets $g(\mathcal{R}^{\tilde{\Sigma}}_{[a-1,b]}(p_*, p_*, 0))$ and $\widetilde{S}$ are separated at (p_*, q_*, c_*). Moreover, if $\widetilde{\gamma}_*$ actually solves the free interval problem, then for any $\alpha > 0$ the sets $g(\mathcal{R}^{\tilde{\Sigma}}_{[a-1,b\pm\alpha]}(p_*, p_*, 0))$ and $\widetilde{S}$ are separated at (p_*, q_*, c_*).

It is easy to see that $\widetilde{C} = C \times (-\infty, 0]$ is an approximating cone for $\widetilde{S}$ at (p_*, q_*, c_*). Moreover, $\widetilde{C}$ is not a linear space, so Theorem 8.1 can be applied to obtain a weakly minimizing adjoint vector $\widetilde{\zeta}$ along $\widetilde{\gamma}_*$ for the system $\widetilde{\Sigma}$, a $\widetilde{\theta} \in T^*_{p_*,q_*,c_*}N \setminus \{0\}$ such that $\widetilde{\theta} \in \widetilde{C}^\dagger$, and a $G \in \partial g(p_*, q_*, c)$ such that $\widetilde{\zeta}(b) = G^\dagger \widetilde{\theta}$.

Now $\widetilde{\zeta}(t) \in T^*_{\tilde{\xi}_*(t)}(M \times M \times \mathbb{R})$, and $\widetilde{\xi}_*(t) = (p_*, \xi_*(t), \sigma_*(t))$ for $a \le t \le b$, $\widetilde{\xi}_*(t) = (p_*, p_*, 0)$ for $a - 1 \le t < a$. Then $\widetilde{\zeta}(t) = (\zeta_1(t), \zeta_2(t), \nu(t))$, where $\zeta_1(t) \in T^*_{p_*}M$ for $t \in [a-1, b]$, $\zeta_2(t) \in T^*_{p_*}M$ for $t \in [a-1, a]$, $\zeta_2(t) \in T^*_{\xi_*(t)}M$ for $t \in [a, b]$, and $\nu(t) \in \mathbb{R}$. On the interval $[a, b]$ the adjoint equation says that $\dot{\nu}(t) = 0$, $\dot{\zeta}_1(t) = 0$, and $-\dot{\zeta}_2(t) \in (\zeta_2(t), \nu(t)) \cdot \partial(F^L_{\eta_*,t}(\xi_*(t))$, where $F^L_{\eta,t}(x) = (F_{\eta_*}(x, t), L_{\eta_*}(x, t))$. So ζ_1 and ν are constant on $[a,b]$. On $[a-1,a]$ the adjoint equation says that ζ_1, ζ_2 and ν are constant. Moreover, the Hamiltonian minimization condition on $[a-1,a]$ implies that $\zeta_1(t)+\zeta_2(t)=0$ for $t\in[a-1,a]$.

If we define $\zeta(t)=\zeta_2(t)$ for $t\in[a,b]$, it follows easily that $\zeta\in \mathrm{Adj}^{\Sigma}(\gamma_*,\nu L_{\eta_*})$. On the other hand, the Clarke generalized Jacobian $\partial g(p_*,q_*,c)$ is easily seen to be equal to $\{G_\lambda : \lambda\in\partial\varphi(p_*,q_*)\}$, where $G_\lambda : T_{p_*}M\times T_{q_*}M\times\mathbb{R}\to T_{p_*}M\times T_{q_*}M\times\mathbb{R}$ is the map $(v_1,v_2,r)\to(v_1,v_2,r+\lambda(v_1,v_2))$. (Recall that each λ is a map from $T_{p_*}M\times T_{q_*}M$ to $\mathbb{R}$.) For $\lambda\in\partial\varphi(p_*,q_*)$, the transpose $G_\lambda^\dagger$ is therefore the map from $T^*_{p_*}M\times T^*_{q_*}M\times\mathbb{R}$ to $T^*_{p_*}M\times T^*_{q_*}M\times\mathbb{R}$ given by

$$G_\lambda^\dagger(\theta_1,\theta_2,\rho)=((\theta_1,\theta_2)+\rho\lambda,\rho)\in T^*_{(p_*,q_*)}(M\times M)\times\mathbb{R}\sim T^*_{p_*}M\times T^*_{q_*}M\times\mathbb{R}. \tag{8.5}$$

In particular, if we let $\widetilde{\theta}=(\theta_1,\theta_2,\rho)$, and pick λ such that $G_\lambda=G$, then $G^\dagger(\widetilde{\theta})=((\theta_1,\theta_2)+\rho\lambda,\rho)$. Therefore $(-\zeta(a),\zeta(b),\nu)=(-\zeta_2(a),\zeta_2(b),\nu)=(\zeta_1(a),\zeta_2(b),\nu)=(\zeta_1(b),\zeta_2(b),\nu)=\widetilde{\zeta}(b)=G^\dagger(\widetilde{\theta})=((\theta_1,\theta_2)+\rho\lambda,\rho)$. Therefore $\nu=\rho$ and $(-\zeta(a),\zeta(b))=(\theta_1,\theta_2)+\nu\lambda$.

The transversality condition $\widetilde{\theta}\in\widetilde{C}^\dagger$ says that $(\theta_1,\theta_2)w+\rho\beta\le 0$ for all $w\in C$, $\beta\le 0$. Taking $\alpha=0$, $\beta=-1$, we get $\rho\ge 0$, i.e. $\nu\ge 0$. Taking $w\in C$, $\beta=0$, we get $(\theta_1,\theta_2)\in C^\dagger$. So $(-\zeta(a),\zeta(b))\in\nu\partial\varphi(p_*,q_*)+C^\dagger$, as stated.

Finally, the fact that $\widetilde{\theta}\neq 0$ implies that either $\nu\neq 0$ or $(-\zeta(a),\zeta(b))\neq(0,0)$. If $\nu=0$, then $\zeta\in\mathrm{Adj}^{\Sigma}(\gamma_*)$, so either $\zeta\equiv 0$ or $\zeta(t)\neq 0$ for all t. Since $(-\zeta(a),\zeta(b))\neq(0,0)$, we find that $\zeta(b)\neq 0$.

The last two assertions then follow from the similar conclusions of Theorem 8.1 , and our proof is complete. ∎

Remark 8.2. If ζ and ν satisfy the conclusions of Theorem 8.3, and $r>0$, then it is easy to see that $\rho\zeta$ and ν satisfy the conclusions as well. Therefore *the number ν can always be assumed to be equal to* 1 *or to* 0. ∎

Remark 8.3. Notice that in Theorem 8.3 we no longer need the special hypothesis A7, because the cone $\widetilde{C}$ is not a linear subspace, even if C is. ∎

Remark 8.4. The transversality condition says that there is a $\lambda\in$

$\partial\varphi(p_*, q_*)$ such that $\langle(-\zeta(a), \zeta(b)) - \nu\lambda, v\rangle \le 0$ for all $v \in C$. In the special case when C is a linear subspace, this actually implies that $\langle(-\zeta(a), \zeta(b)) - \nu\lambda, v\rangle = 0$ for all $v \in C$. Therefore, in this case the transversality condition simply says that $(-\zeta(a), \zeta(b))\lceil C = \nu\lambda\lceil C$. This happens, in particular, if S is a submanifold of M, in which case C can be taken to be the tangent space to S at q_*. In the special case when $\varphi = 0$ and S is a submanifold of $M \times M$, the transversality condition says that $(-\zeta(a), \zeta(b))$ is orthogonal to $T_{(p_*, q_*)}S$. If S is of the form $\{p_*\} \times Z$, where Z is a submanifold of M (i.e. if $\xi(a)$ is fixed and we only have a terminal point constraint $\xi(b) \in Z$), then the transversality condition says that $\zeta(b)$ is orthogonal to $T_{q_*}Z$. ■

Remark 8.5. If $S = \{p_*\} \times M$ and φ is just a function of q, then we have a *free terminal point problem*. In that case, the transversality condition simply says that $\zeta(b) \in \nu\partial\varphi(q_*)$. If $\nu = 0$ then $\zeta(b) = 0$, contradicting the fact that $\zeta(b)$ and ν cannot both vanish. So *for a free terminal point problem we can always assume that* $\nu = 1$. ■

Remark 8.6. The possibility that $\nu = 0$ can also be excluded in other cases. For example, for the classical Calculus of Variations, in which the dynamics is just $\dot{x} = u$, the Hamiltonian in the sense of Theorem 8.3 is given by $\langle z, u\rangle + \nu L(x, u, t)$. If ν was equal to zero, the minimization condition would require the control $u(t)$ being used to minimize the linear functional $v \to \langle\zeta(t), v\rangle$ for v ranging over the whole space $\mathbb{R}^n$. Clearly, such a functional does not have a minimum unless $\zeta(t) = 0$. But this possibility is excluded by the nontriviality condition. So *for a classical Calculus of Variations problem we can always assume that* $\nu = 1$. In particular, *Theorem* 8.3 *implies Theorem* 2.2. ■

An important particular case of the previous theorem is:

Theorem 8.4 (The Maximum Principle for time-optimal control)**.** *Assume that the data* Σ, M, $\mathcal{U}$, F, ξ_*, η_*, γ_*, a, b, p_*, q_*, S *and* C *are such that* A1, A2, A3, A4, A8 *and* OPT2 *hold. Assume that* ξ_* *is a solution of the minimization problem* $\mathcal{P}_{time}(\Sigma, S)$. *Then there exists a nonzero weakly minimizing adjoint vector* ζ *along* γ_* *such that* $(-\zeta(a), \zeta(b)) \in C^\dagger$ *and the Hamiltonian* $\langle\zeta(t), F_{\eta_*}(\xi_*(t), t)\rangle$ *is constant and* ≤ 0 *almost everywhere.*

Proof. This is a trivial consequence of Theorem 8.3, taking g = identity and $\varphi \equiv 0$. Indeed, the Hamiltonian $\mathcal{H}_\eta^{\Sigma,\nu}$ is given by $\mathcal{H}_\eta^{\Sigma,\nu}(x,z,t) = \langle z, F_\eta(x,t)\rangle + \nu$. So Theorem 8.3 gives a $\nu \geq 0$ and a $\zeta \in \mathrm{Adj}^\Sigma(\gamma_*,\nu)$ such that $\zeta(b)$ and ν cannot both vanish, $\mathcal{H}_{\eta_*}^{\Sigma,\nu}$ vanishes a.e. along Ξ (where $\Xi(t) = (\xi_*(t), \zeta(t))$), and ζ satisfies the appropriate transversality condition. The condition that $\zeta \in \mathrm{Adj}^\Sigma(\gamma_*,\nu)$ just says that $\zeta \in \mathrm{Adj}^\Sigma(\gamma_*)$. In terms of the Hamiltonian $\mathcal{H}_\eta^\Sigma(x,z,t) = \langle z, F_\eta(x,t)\rangle$ of Σ, the condition that $\mathcal{H}_{\eta_*}^{\Sigma,\nu} = 0$ along Ξ just says that $\mathcal{H}_{\eta_*}^\Sigma$ is constant and ≤ 0 along Ξ. If $\zeta \equiv 0$, then the vanishing of $\mathcal{H}_{\eta_*}^{\Sigma,\nu}$ along Ξ implies that $\nu = 0$ as well, contradicting the nontriviality condition. So $\zeta \neq 0$. Finally, since $\varphi = 0$, the transversality condition just says that $(-\zeta(a), \zeta(b)) \in C^\dagger$. ∎

Definition 8.1. A controlled arc that satisfies the necessary condition of Theorem 8.1, 8.2, 8.3, or 8.4, for the corresponding problem (separation, controllability, optimal control or minimum time control) is a *controlled extremal* for the problem. ∎

9. Classical Systems: the Autonomous Case

In this and the next section we show how the Maximum Principle applies to systems of the "classical" form $\dot{x} = f(x,t,u)$, and in particular we explain how in this case one can obtain strong rather than weak minimization in the conclusion, under rather general technical conditions.

We will do this in two cases, namely, (a) that of systems of the form $\dot{x} = f(x,u)$, $u \in U$, and (b) systems $\dot{x} = f(x,t,u)$, $u \in U(t)$. In the former case, we will allow the control set U to be a completely general set, with no structure whatsoever, and will only require $f(x,u)$ to be continuous in x for each fixed u. In the latter case, we select a setting similar to that of F. Clarke's version of the MP, presented in [3], [4].

We now discuss Case (a). Here the relevant objects of interest will be systems generated by a family $f = \{f_u\}_{u\in U}$ of ordinary vector fields on M, and we will have to define in a precise way how such objects give rise to control systems $\Sigma = (M, \mathcal{U}, F)$ in the sense of Definition 5.1. Roughly speaking, if we write $f(x,u) = f_u(x)$, then Σ ought to be "the control system $\dot{x} = f(x, u)$, $u \in U$." This, however, is still vague, because it fails to to specify the "class $\mathcal{U}$ of admissible open-loop controls," i.e. the set of functions $t \to u(t)$ that can actually be plugged into the formal expression $\dot{x} = f(x, u)$ to produce a true differential equation $\dot{x} = f(x, u(t))$. For this reason, the system Σ associated with a given f will not be unique, because there is no canonical choice of $\mathcal{U}$.

Definition 9.1. A *vector field family* on a manifold M is a parametrized family $f = \{f_u\}_{u \in U}$ of vector fields on M (i.e. continuous sections of TM). A *vector field system* is a triple $\mathcal{V} = (M, U, f)$ such that $f = \{f_u\}_{u \in U}$ is a vector field family on M. ■

If $\mathcal{V} = (M, U, f)$ is a vector field system, we let $\widetilde{\mathcal{V}}$ denote the set of all vector fields f_u, $u \in U$. We say that an $F \in CVF(M)$ is *generated by* $\mathcal{V}$ if $F_t \in \widetilde{\mathcal{V}}$ for every $t \in \mathcal{TD}(F)$. We call a CVF F *piecewise constant* if there is a locally finite partition $\mathcal{P}$ of $\mathcal{TD}(F)$ into intervals such that the map $t \to F_t$ is constant on each $J \in \mathcal{P}$.

Definition 9.2. A control system $\Sigma = (M, \mathcal{U}, F)$ is *generated* by a vector field system $\mathcal{V} = (\widetilde{M}, U, f)$ if (i) $M = \widetilde{M}$, (ii) F_η is generated by $\mathcal{V}$ for every $\eta \in \mathcal{U}$, and (iii) whenever I is any interval such that $I = \mathcal{TD}(F_{\eta_0})$, for some $\eta_0 \in \mathcal{U}$, and G is a piecewise constant CVF generated by $\mathcal{V}$ with time domain I, then there exists an $\eta \in \mathcal{U}$ such that $G = F_\eta$ on $M \times I$. ■

Remark 9.1. Condition (iii) is a slightly weaker form of the natural requirement that "all control vector fields arising from piecewise constant U-valued controls" should occur in Σ. We choose only to impose this weaker condition because we want to allow, for example, systems with a fixed time interval. ■

The following general result, together with Proposition 6.1, says that for systems generated by vector field systems the Maximum Principle can always be applied in the "strong" form, i.e. with strong rather than weak minimization.

Theorem 9.1. *Every control system $\Sigma = (M, \mathcal{U}, F)$ generated by a vector field system $\mathcal{V} = (M, U, f)$ is separable.*

Proof. Let $\xi : [a,b] \to M$ be a trajectory of Σ corresponding to a control η_0. Let $I = \mathcal{TD}(F_{\eta_0})$, so $[a,b] \subseteq I$. Let $K \in \mathrm{Comp}(M)$ be such that $\{\xi(t) : a \leq t \leq b\} \subseteq \mathrm{Int}(K)$. Let $\Gamma^0(K)$ denote the space of continuous vector fields on K, i.e. the set of restrictions to K of continuous sections of TM (or, equivalently, the set of continuous maps $X : K \to TM$ such that $\pi_{TM} \circ X = \mathrm{id}_K$). Then $\Gamma^0(K)$ is a separable Banach space, if we pick a Riemannian metric on M and let $\|X\| = \sup\{\|X(x)\| : x \in K\}$. So the subspace $W = \{f_u \lceil K : u \in U\}$ is separable as well. Let W_0 be a countable dense subset of W, and pick for each $g \in W_0$ a $u_0 = \mu(g)$ such that $g = f_{u_0} \lceil K$. Let $U_0 = \{\mu(g) : g \in W_0\}$. Then U_0 is countable. Let $\mathcal{W}_0$ be the set of all piecewise constant U_0-valued maps $\theta : I \to U_0$ that are constant on the intervals of a finite partition $\mathcal{P}_\theta$ of I with rational partition points. Then $\mathcal{W}_0$ is countable. If $\theta \in \mathcal{W}_0$, then Definition 9.2 implies that there exists an $\eta = \nu(\theta)$ in $\mathcal{U}$ such that $I \subseteq \mathcal{TD}(F_\eta)$ and $f_{\theta(t)}(x) = F_\eta(x,t)$ for $(x,t) \in M \times I$. Let $\mathcal{U}_0 = \{\nu(\theta) : \theta \in \mathcal{W}_0\}$. Then $\mathcal{U}_0$ is countable as well. If $\eta \in \mathcal{U}_0$, then of course $I \subseteq \mathcal{TD}(F_\eta)$. Moreover, if $t \in [a,b]$, then $\{F_\eta(\xi(t),t) : \eta \in \mathcal{U}_0\}$ is clearly dense in $\{f_u(\xi(t)) : u \in U\}$, which contains the set $\{F_\eta(\xi(t),t) : \eta \in \mathcal{U}, [a,b] \subseteq \mathcal{TD}(F_\eta)\}$. So Σ is separable along ξ. ∎

Remark 9.2. At first sight, it may appear surprising that a general separability result can be proved in the setting where the control set U is completely arbitrary, and in particular is not required to be a separable metric space. However, our definition of separability only involves the values of the vector fields $F_\eta(\cdot,t)$, and the control values are just parameters. So all that really matters is the separability of the Fréchet space of vector fields. ∎

For a vector field system $\mathcal{V} = (M,U,f)$, we define a *Lagrangian* to be a family $L = \{L_u\}_{u\in U}$ of continuous functions on M. The *augmented vector field system* $\mathcal{V}^L$ is the system $(M \times \mathbb{R}, U, f^L)$, where $f^L = \{f^L_u : u \in U\}$, $f^L_u(x,x^0) = (f_u(x), L_u(x))$. We can then consider various control systems generated by $\mathcal{V}^L$.

10. Classical Systems: the Time-Varying Case

We now turn to the second case, namely, systems $\dot{x} = f(x,t,u)$. We select a setting similar to that of F. Clarke's version of the MP, presented in [3], [4].

Recall that (a) the *Borel σ-algebra* $\mathcal{B}(X)$ of a metric space X is the σ-algebra generated by the open sets, and (b) a *Polish space* is a separable topological space X whose topology arises from a metric d such that (X,d) is complete. If I is an interval, we use $\mathcal{L}(I)$ to denote the σ-algebra of Lebesgue measurable subsets of I.

If P is a Polish space, $I \in \mathcal{I}$, and M is a smooth manifold, we define a *classical control dynamics* (abbr. CCD) *on M with time domain I and ambient control space P* to be a pair (U,f) such that

(I) U is an $\mathcal{L}(I) \times \mathcal{B}(P)$-measurable subset of $I \times P$, such that $U(t) = \{u : (t,u) \in U\}$ is nonempty for every $t \in I$;

(II) $(x,t,u) \to f(x,t,u)$ is a map from $M \times U$ to TM, such that $f(x,t,u) \in T_xM$ for all $(x,t,u) \in M \times U$,

(III) every map $f(\cdot,t,u)$, $(t,u) \in U$, is continuous,

(IV) every map $f(x,\cdot,\cdot)$, $x \in M$, is $\mathcal{L}(I) \times \mathcal{B}(P)$-measurable.

(In particular, (I), (III) and (IV) imply that f is $\mathcal{B}(M) \times \mathcal{L}(I) \times \mathcal{B}(P)$-measurable.)

A *control* is a Lebesgue measurable map $\eta : J \to P$ such that J is a subinterval of I and $\eta(t) \in U(t)$ for each $t \in J$.

If we use $\mathcal{U}$ to denote the class of all controls, and then write $F_\eta(x,t) = f(x,t,\eta(t))$, $F = \{F_\eta\}_{\eta \in \mathcal{U}}$, then $\Sigma = (M,\mathcal{U},F)$ is a control system. A Σ arising in this way from a CCD (U,f) will be called a *classical control system*.

Theorem 10.1. *A classical control system is separable and has the measurable substitution property.*

Proof. The substitution property is trivial. Separability is proved in [23]. ∎

It follows from Theorem 10.1 that for a classical control system the MP applies, with the extra conclusion that the adjoint vector is *strongly minimizing*. (For the optimality results, we have to require that the augmented system Σ^L, obtained by adding the running cost as a new state variable, satisfy the same technical assumptions. That is, the Lagrangian $(x,t,u) \to L(x,t,u)$ has to be continuous in x for each (t,u), and $\mathcal{L}(I) \times \mathcal{B}(P)$-measurable in (t,u) for each x, exactly as in [4].)

For Theorem 8.1 to apply, the reference CVF $f_* : (x,t) \to f(x,t,u_*(t))$ has to be LIL. This is a weaker assumption than those of [4], Theorems 5.1.2 and 5.2.1, where (a) $f(x,t,u)$ has to be Lipschitz with respect to x for each t,u, with a Lipschitz constant $k(t,u)$, and (b) the function $t \to k(t,u_*(t))$ has to be integrable. Condition (b) is our requirement that f_* be LIL—and this hypothesis is essential, as can be shown by simple examples—cf. Example 15.5 below—but *Condition* (a) *is no longer needed*.

11. An Example: the Derivation of the Geodesic Equation

Our first example of an application of the Maximum Principle will be the derivation of the well-known equations of the Riemannian geodesics. This will illustrate both the power and the simplicity of the control theory method, especially when it is combined with an intrinsically geometric formulation and with the approach of regarding a control system as a—possibly very large—collection of vector fields. In particular, we work with the geometric formulation of the adjoint equation in terms of a connection along a curve, which turns out to be tailor-made for use in a differential-geometric setting.

Let M be a Riemannian manifold, with metric tensor G. Our goal is to characterize those arcs $\xi : [a,b] \to M$ that minimize length among all arcs $\widetilde{\xi}$ such that $\partial\widetilde{\xi} = \partial\xi$. (Recall that if $\mathcal{D}(\xi) = [a,b]$ then $\partial\xi \stackrel{\text{def}}{=} (\xi(a), \xi(b))$). It is well known that every such arc can be reparametrized by a constant times arc-length in such a way that $a = 0, b = 1$, and that the minimum-length arcs $\xi : [0,1] \to M$ that are so parametrized are precisely those that minimize the integral $J(\xi) = \frac{1}{2}\int_0^1 G(\dot{\xi}(t), \dot{\xi}(t))\, dt$. So it suffices to characterize the arcs $\xi : [0,1] \to M$ that minimize $J(\xi)$ among all $\widetilde{\xi}$ such that $\partial\widetilde{\xi} = \partial\xi$.

Let $\mathcal{A}$ be the set of all arcs $\xi : [0,1] \to M$ for which $J(\xi) < \infty$, so $\mathcal{A}$ is the Sobolev space $W^{1,2}([0,1], M)$. For $p, q \in M$, let $\mathcal{M}(p,q) = \{\xi \in \mathcal{A} : \partial\xi = (p,q)\}$, and let $\mathcal{M} = \cup_{p,q}\mathcal{M}(p,q)$. Our goal is then to characterize the class $\mathcal{M}$ by deriving a necessary condition for a $\xi \in \mathcal{A}$ to belong to $\mathcal{M}$. Naturally, the condition is going to be the familiar equation for the geodesics.

To apply the Maximum Principle, we first realize the class $\mathcal{A}$ as the set of all trajectories of a control system generated by a vector field system (U, f). The simplest way to do this is by letting $U = \Gamma^\infty(TM)$, the space of all smooth vector fields u on M. Let us introduce the somewhat redundant notation f_u for a vector field $u \in U$, so $f_u(x)$ is just equal to $u(x)$, and let us also write $f(x,u) \stackrel{\text{def}}{=} f_u(x)$. With these notations, we can think of the family of vector fields $\{f_u\}_{u \in U}$ as a control system $\dot{x} = f(x,u)$ in the traditional control theory sense, except that our choice of control space U is somewhat unorthodox, since U is a very large space of vector fields. A *control* will be, for us, a LIC^∞ control vector field $(x,t) \to F(x,t)$ on M with time domain $[0,1]$. Equivalently, we may think of a control as a function $[0,1] \ni t \to \eta(t) \in U$ such that the map $(t,x) \to f_{\eta(t)}(x)$ —i.e. the map $(t,x) \to \eta(t)(x)$—is measurable with respect to t for each fixed x, and is bounded uniformly with respect to t together with all x-derivatives of all orders on each compact subset of M. (That is, a control is a bounded measurable map from $[0,1]$ to the Fréchet space $\Gamma^\infty(TM)$.) We let $\mathcal{U}$ be the class of all controls, and write $F_\eta(x,t) = \eta(t)(x)$ for $\eta \in \mathcal{U}$. Then $\Sigma = (M, \mathcal{U}, F)$ is a control system in the sense of our definition. It is clear that every trajectory of a control is in $\mathcal{A}$ and the converse—that is, that every arc $\xi \in \mathcal{A}$ arises as a trajectory for some control η—is trivially true. So $\mathcal{A}$ is exactly the class of all trajectories of Σ.

For any control $\eta \in \mathcal{U}$, define $L_\eta : M \times [0,1] \to \mathbb{R}$ by $L_\eta(x,t) = \frac{1}{2}G(\eta(t)(x), \eta(t)(x))$ for $x \in M$. Then $L = \{L_\eta : \eta \in \mathcal{U}\}$ is a Lagrangian for Σ in the sense of our definitions. Given $p, q \in M$, the members of

$\mathcal{M}(p,q)$ are exactly the solutions of the optimal control problem with fixed endpoints in which one seeks to minimize the integral $\int_0^1 L_\eta(\xi(t),t)dt$ among all controlled trajectories $\gamma=(\xi,\eta)$ of Σ for which $\xi(0)=p$, $\xi(1)=q$. (In the terminology of Section 8, this is Problem $\mathcal{P}_{[0,1]}(\Sigma,L,\{(p,q)\},0)$.)

The Maximum Principle then gives the following necessary condition for a $\xi_*\in\mathcal{A}$, arising from a control η_*, to belong to $\mathcal{M}(p,q)$: there exist a field of covectors $t\to\zeta(t)$ along ξ_* and a $\nu\geq 0$ that satisfy the following five conditions: (i) nontriviality, i.e. $(\zeta(t),\nu)\neq(0,0)$ for $t\in[0,1]$, (ii) absolute continuity, (iii) the adjoint equation, (iv) the Hamiltonian minimization condition, which says that for almost every $t\in[0,1]$ the function $\eta\to\langle\zeta(t),\eta(t)(\xi_*(t))\rangle+\nu L_\eta(\xi_*(t),t)$ is minimized by η_*, and (v) the constancy of the Hamiltonian condition, which says that there is a constant ρ such that $t\to\langle\zeta(t),\eta_*(t)(\xi_*(t))\rangle+\nu L_{\eta_*}(\xi_*(t),t)=\rho$ a.e. (Our system Σ is clearly invariant under deletions and insertions —so ζ can be chosen so as to satisfy (v)— and separable, so (iv) holds.)

We now show how these conditions imply the geodesic equation. First of all, it is clear that the Hamiltonian minimization condition holds for a particular t iff $\eta_*(t)(\xi_*(t))$ —i.e. $\dot{\xi}_*(t)$— minimizes the functional $T_{\xi_*(t)}M\ni y\to\langle\zeta(t),y\rangle+\frac{\nu}{2}G(y,y)$. This functional does not have a minimum if $\nu=0$, $\zeta(t)\neq 0$, so ν cannot vanish. Therefore we may assume that $\nu=1$. Next, recall that the metric G gives rise to an isomorphism $G(x):T_xM\to T_x^*M$ such that $\langle G(x)v,w\rangle=\langle v,w\rangle_G$ whenever $v,w\in T_xM$. If we let $w(t)=G(\xi_*(t))^{-1}(\zeta(t))$, then the vector $\dot{\xi}_*(t)$ minimizes the functional $T_{\xi_*(t)}M\ni y\to G(w(t),y)+\frac{1}{2}G(y,y)$. Now, if we have an inner product $\langle\cdot,\cdot\rangle$ on a finite-dimensional linear space S, and a nonzero vector $w\in S$, it is clear that the functional $v\to\langle w,v\rangle+\frac{1}{2}\langle v,v\rangle$ is minimized by $v=-w$. So $\dot{\xi}_*(t)=\eta_*(t)(\xi_*(t))=-w(t)$, i.e. $\dot{\xi}_*(t)=-G(\xi_*(t))^{-1}(\zeta(t))$.

Condition (ii) then implies that $\dot{\xi}_*(t)$ is absolutely continuous. The constant ρ is equal to $\langle\zeta(t),\eta_*(t)(\xi_*(t))\rangle+L_{\eta_*}(\xi_*(t),t)$, i.e. to $\langle w(t),\dot{\xi}_*(t)\rangle+\frac{1}{2}G(\dot{\xi}_*(t),\dot{\xi}_*(t))$, which equals $-\frac{1}{2}\|\dot{\xi}_*(t)\|_G^2$. So $\|\dot{\xi}_*(t)\|_G$ is equal to a constant σ.

We now use the remaining piece of information, namely, the adjoint equation. This says that $\mathcal{L}^\gamma\zeta=-\omega$, where $\mathcal{L}^\gamma$ is the connection along

ξ_* canonically associated with the controlled arc $\gamma = (\xi_*, \eta_*)$, $\omega(t) = dL_{\eta_*,t}(\xi_*(t))$, and $L_{\eta_*,t}(x) \stackrel{\text{def}}{=} L_{\eta_*}(x,t)$.

So far, η_* is an arbitrary control that generates ξ_*, and we still have the freedom to choose it. Let us recall that, according to Proposition 4.1, every LI connection $\mathcal{L}$ along ξ_* is of the form $\mathcal{L}^{\xi_*,F}$ for some LIC$^\infty$ CVF F with domain $M \times [0,1]$. In particular, we can take $\mathcal{L}$ to be the pullback to ξ_* of the *Riemannian connection*, or *Levi-Civita connection* of G, i.e. of the unique torsion-free affine connection ∇ such that $\nabla G = 0$. (The torsion-free condition says that $\nabla_X Y - \nabla_Y X = [X,Y]$ for every $X, Y \in \Gamma^\infty(TM)$.) The adjoint equation then becomes $\nabla_{\dot{\xi}_*(t)}(\zeta) = -dL_{\eta_*,t}(\xi_*(t))$, where $\eta_*(t) = F_t$, $F_t(x) = F(x,t)$, and F is a LIC$^\infty$ CVF such that ξ_* is an integral curve of F and $\mathcal{L}_t^{\xi_*,F} = \nabla_{\dot{\xi}_*(t)}$ for almost all t. In that case, $2L_{\eta_*}(x,t) = G(F_t(x), F_t(x))$.

Now, if $X, Y \in \Gamma^\infty(TM)$, then $XG(Y,Y) = 2G(\nabla_X Y, Y)$, because $\nabla_X G = 0$. Then $XG(Y,Y) = 2G(\nabla_Y X + [X,Y], Y)$, because ∇ is torsion-free. On the other hand, if $0 \leq t \leq 1$, then $\mathcal{L}_t = \mathcal{L}_t^{\xi_*,F}$, so $\mathcal{L}_t X = [F_t, X](\xi_*(t))$ for every $X \in \Gamma^\infty(TM)$. Therefore $XG(F_t,F_t)(\xi_*(t)) = 2G(\nabla_{F_t} X + [X, F_t], F_t)(\xi_*(t))$. Since $\mathcal{L}_t X = \nabla_{F_t} X(\xi_*(t))$, we have

$$XG(F_t,F_t)(\xi_*(t)) = 2G(\mathcal{L}_t X + [X,F_t](\xi_*(t)), F_t(\xi_*(t))) = 0. \tag{11.1}$$

Since X is arbitrary, we have in fact shown that $dG(F_t,F_t)(\xi_*(t)) = 0$, i.e. that $dL_{\eta_*,t}(\xi_*(t)) = 0$. So the adjoint equation just says that $\nabla_{\dot{\xi}_*(t)}(\zeta) = 0$. But then

$$\nabla_{\dot{\xi}_*(t)}(\dot{\xi}_*) = \nabla_{\dot{\xi}_*(t)}(G^{-1}\zeta) = G^{-1}(\nabla_{\dot{\xi}_*(t)}(\zeta)) = 0,$$

using again the fact that G is invariant under ∇. So we have shown that

$$\nabla_{\dot{\xi}_*}\dot{\xi}_* = 0,$$

which is the geodesic equation. ∎

Remark 11.1. Notice that in the above derivation we never took coordinates, and everything was done directly in an invariant way. In particular, we only used the two properties characterizing the Riemannian connection, but did not write the explicit formulas for the Christoffel symbols, and we applied all the conditions of the Maximum Principle, but did not express any of them —not even the adjoint equation— in local coordinates. ∎

12. Single-Input Control-Affine Minimum Time Problems

Our second example of an application of the MP is in many ways diametrically opposed to the first one. Whereas the problem of Section 11 is one with a kinematics consisting of all curves, and the interesting structure arises from the Lagrangian —i.e. from the metric— we will now look at problems where the Lagrangian L is $\equiv 1$, which is as trivial as a function can be, and all the geometric structure comes from the kinematics. We will show how the MP can be applied to get strong conclusions about the structure of optimal controls.

We consider the simplest type of kinematics that involves nontrivial control. Clearly, this requires a minimum of two vector fields, so we will focus on "two-vector field" control systems, i.e. systems in which at each point x there are two directions of motion, $X(x)$ and $Y(x)$. (For general reasons —e.g. because we want to have existence of optimal controls, and/or we prefer that the reachable sets be closed— it is always a good idea to have convexity of the set of directions of motion at a point, so we will consider systems where the dynamics is such that $\dot{x}$ belongs to the convex hull of $X(x)$ and $Y(x)$.) Writing $f = \frac{1}{2}(X+Y)$, $g = \frac{1}{2}(X-Y)$, we end up with a system of the form $\dot{x} = f(x) + ug(x)$, $|u| \leq 1$. The fact that $L \equiv 1$ means, of course, that we will be looking at *minimum time control*.

So we will consider control systems $\dot{x} = f(x) + ug(x)$, $|u| \leq 1$, where x takes values in a manifold M, and f, g are smooth vector fields on M. The controls are measurable functions $u(\cdot) : [a,b] \to [-1,1]$, defined on some compact interval $[a,b]$. We would like to understand the structure of the time-optimal trajectories and their corresponding controls. Naturally, the MP can be applied, but this is only the first step in the analysis, since our real problem is to how translate the information provided by the MP into interesting facts about the optimal controls.

The MP, applied to this case, says that if $\xi_* : [a,b] \to M$ is a time-optimal trajectory corresponding to a control $u_* : [a,b] \to [-1,1]$, then there exists a nowhere vanishing AC field of covectors ζ along ξ_* that satisfies the adjoint equation (AE) and, for some constant $\rho \geq 0$, the Hamiltonian minimization condition:

$$
\begin{aligned}
-\rho &= \zeta(t).(f + u_*(t)g)(\xi_*(t)) \\
A &= \min\{\zeta(t).(f + vg)(\xi_*(t)) : v \in [-1,1]\} \qquad \text{for a.e. } t \in [a,b].
\end{aligned}
\tag{12.1}
$$

The problem is to interpret the above conditions and translate them into facts about u_*. Define the *switching function* $\varphi : [a,b] \to \mathbb{R}$ by letting $\varphi(t) = \zeta(t)g(\xi(t))$. Then (12.1) clearly implies that $u_* = 1$ as long as $\varphi < 0$, and $u_* = -1$ as long as $\varphi > 0$, so the switchings of $u_*(t)$ occur at the zeros of φ, and therefore the properties of these zeros translate into information about the optimal controls. For example, if the switching function φ could be proved to have a discrete set of zeros, then it would follow that the optimal controls are "bang-bang," i.e. piecewise constant with values 1 or -1.

One possible course of action now would be to write the adjoint equation in local coordinates. This might seem reasonable, at least in the Euclidean case, i.e. when M is an open subset of $\mathbb{R}^n$, so that there is a natural choice of coordinates. However, we shall argue that, *even in the Euclidean case, the resulting system of equations, namely,* $\dot{z}_i = -\frac{\partial H}{\partial x_i}(x,z,u)$, *with* $x = \xi_*(t)$, $u = u_*(t)$, *is unnecessarily complicated, and fails to display the most relevant information about the problem.* This is so because (a) *the components* $\zeta_i(t)$ *of the adjoint vector* $\zeta(t)$ *are the inner products* $\zeta(t).e_i$ *of* $\zeta(t)$ *with the members* e_i *of the canonical basis of* $\mathbb{R}^n$, and (b) *the inner products* $\zeta(t).e_i$ *are much less significant than the switching function* φ, *which is the inner product of* ζ *with the vector field* g. (For example, we have explained above how the structure of the set of zeros of φ directly implies properties of the optimal controls.) For this reason, it is more convenient —even in the Euclidean case— to use a formulation of the MP where the canonical basis of $\mathbb{R}^n$ is not singled out for preferential treatment, and the products $\zeta.X$ of ζ with all possible vector fields X are treated alike. This is precisely what the intrinsic formulation of the AE given in Section 4 does, since it gives a condition on the derivative of the function $t \to \zeta(t).X$ for all possible vector fields X. The AE then turns out to say that the derivative of $\zeta(t).X$ is another function of the same kind, where X is replaced by the Lie bracket $[f + u_*(t)g, X]$. That is,

$$\text{(12.2)} \quad \frac{d}{dt}(\zeta(t).X(\xi_*(t))) = \zeta(t).[f + u_*(t)g, X](\xi_*(t)) \quad \text{for a.e. } t \in [a, b]$$

for *all* smooth vector fields X on M.

Example 12.1. According to (12.2), the time derivative $\dot\varphi$ of φ is the function ψ given by $\psi(t) = \langle \zeta(t), [f + u(t)g, g](\xi(t)) \rangle$. Since $[g, g] = 0$, we have in fact $\psi(t) = \langle \zeta(t), [f, g](\xi(t)) \rangle$. This already gives interesting information about the optimal trajectories. For example, suppose that M is two-dimensional. Then, *if $g(x)$ and $[f, g](x)$ are linearly independent at every point $x \in M$, it follows that every time-optimal trajectory is bang-bang.* (Indeed, since $\zeta(t) \neq 0$, the functions φ and $\dot\varphi$ cannot vanish simultaneously, so the set of zeros of φ is discrete.) More generally, this shows that the occurrence of non bang-bang optimal arcs is intimately related to the structure of the "singular" set, i.e. the set S of points $x \in M$ such that $g(x)$ and $[f, g](x)$ are linearly dependent. Let us call a trajectory $\xi_* : [a, b] \to M$ arising from a control $u_*(t)$ *singular* on a subinterval J of $[a, b]$ if $-1 < u_*(t) < 1$ almost everywhere on J. Then (assuming still that $\dim M = 2$) *if ξ_* is singular on J, then $\xi_*(J) \subseteq S$.* In other words, "singular optimal arcs" if they exist, have to lie in S. This observation is the first step of a long analysis that can be pursued all the way, until one obtains a complete characterization of the structure of optimal trajectories for general real-analytic problems of the form $\dot x = f(x) + ug(x)$, $|u| \leq 1$, in two dimensions (cf. [17], [18], [19]). It is easy to see why real analyticity should make a difference at this point: when f and g are real-analytic, the set S is the set of zeros of a real-analytic function —the determinant of f and $[f, g]$— so the structure of S is particularly simple, since S is a locally finite union of points and analytic arcs. ∎

The following example shows that "singular optimal arcs," as defined in Example 12.1, can indeed occur in very simple situations.

Example 12.2. (Example 12.1, continued.) Consider the control system Σ in the plane whose dynamical equations are given by $\dot{x}_1 = 1 - x_2^2$, $\dot{x}_2 = u$, where x_1, x_2 are real variables and the control u is required to satisfy $|u| \leq 1$. Let $\xi_* : [0,1] \to \mathbb{R}^2$ be the trajectory given by $\xi_*(t) = (t,0)$, corresponding to the control $u_*(t) \equiv 0$. Then it is easy to see that ξ_* is time-optimal, because ξ_* goes from $(0,0)$ to $(1,0)$ in time 1, and the dynamics of Σ does not allow motion from left to right with horizontal speed > 1. On the other hand, it is clear that ξ_* is singular in the sense of Example 12.1. Our system is of the form $\dot{x} = f(x) + ug(x)$, with $f = (1-x_2^2)\frac{\partial}{\partial x_1}$, $g = \frac{\partial}{\partial x_2}$. A simple calculation shows that the Lie bracket $[f,g]$ is equal to $2x_2\frac{\partial}{\partial x_2}$. So g and $[f,g]$ are linearly dependent on the set $\{(x_1,x_2) : x_2 = 0\}$, which is precisely where ξ lies, in agreement with the analysis of Example 12.1. ∎

Example 12.3. A *single-input linear system* is a system of the form $\dot{x} = f(x) + ug(x)$ on $\mathbb{R}^n$, for which the vector field f is linear —i.e. $f = V_A$ for some square matrix A, where $V_A(x) \stackrel{\text{def}}{=} Ax$— and g is constant, i.e. $g = W_b$ for some $b \in \mathbb{R}^n$, where $W_b(x) \stackrel{\text{def}}{=} b$. We will consider the case when the set of control values is the interval $[-1,1]$. The adjoint equation reduces in this case to $\dot{\zeta} = -\zeta.A$, and the switching function φ is given by $\varphi(t) = \zeta(t).b$, so the higher order derivatives $\varphi^{(k)}$ of φ are $\varphi^{(k)}(t) = (-1)^k\zeta(t).A^kb$. If ν is the smallest positive integer such that $A^\nu b \in$ linear span $(b, Ab, \ldots, A^{\nu-1}b)$, then φ satisfies a linear constant-coefficient O.D.E. of the the form $\varphi^{(\nu)} + \sum_{k=0}^{\nu-1} a_j\varphi^{(j)} = 0$. This implies that the set of zeros of φ is discrete —and therefore u_* is bang-bang— provided that φ does not vanish identically. A system $\dot{x} = Ax + ub$ in $\mathbb{R}^n$ is called *controllable* if $\text{linearspan}(b, Ab, \ldots, A^{n-1}b) = \mathbb{R}^n$. For a controllable system, $\nu = n$, and φ can only vanish identically if $\zeta(t).A^kb = 0$ for $k = 0, \ldots, n-1$, in which case $\zeta(t) \equiv 0$, contradicting the nontriviality condition. So *for a controllable single-input linear system, every time-optimal trajectory is bang-bang.* ∎

Example 12.4 (Example 12.3, continued)**.** For a noncontrollable single-input linear system, all pathology would appear to be possible. For example, consider the system $\dot{x}_1 = x_1$, $\dot{x}_2 = u$, $|u| \le 1$, which is clearly not controllable, since $b = (0,1)^\dagger$ and $A = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 0 \end{smallmatrix}\right)$, so that $Ab = 0$. (Here "$\dagger$" denotes transpose.) Let $u_*(\cdot) : [0,T] \to [-1,1]$ be an arbitrary measurable function. Let ξ_* be the corresponding trajectory, with initial condition $\xi_*(0) = (1,0)^\dagger$. Then $\xi_*(T) = (e^T, y)^\dagger$, where y depends on u_*, but it is clear that for any other trajectory $\xi : [0,\tau] \to \mathbb{R}^2$ going from $(1,0)^\dagger$ to $(e^T, y)^\dagger$ the time τ has to be equal to T. Therefore all trajectories that start at $(1,0)^\dagger$ are optimal. This means that optimal controls for this problem can be as pathological as arbitrary controls, i.e. arbitrary measurable functions.

On the other hand, the appearance of unlimited pathology is somewhat deceptive, at least for this problem, since one can slightly modify the question of strong piecewise regularity for optimal controls, in a very reasonable way, so as to get a positive answer. Let us call a set $\mathcal{U}$ of controls *sufficient for time-optimality* if, whenever $\overline{x}$, $\widehat{x}$ are such that there is a time-optimal trajectory ξ from $\overline{x}$ to $\widehat{x}$, it follows that there is a time-optimal trajectory $\widetilde{\xi}$ from $\overline{x}$ to $\widehat{x}$ that arises from a control in $\mathcal{U}$. (In other words, $\mathcal{U}$ is sufficient if for the purpose of optimally controlling our system one can forget about the controls that do not belong to $\mathcal{U}$, because all the optimal controlling that can be done with general controls can also be done with controls in $\mathcal{U}$.) Then it is easy to see that the class $\mathcal{U}$ consisting of all bang-bang controls with at most one discontinuity is sufficient for our example, since the effect of a control $u : [a,b] \to [-1,1]$ is completely determined by its integral $\int_a^b u(t)dt$, and all values of the integral that can be realized with general controls can also be realized with controls in $\mathcal{U}$. ■

The results of the previous examples can be generalized to arbitrary single-input linear systems of the form $\dot{x} = Ax + ub$ in $\mathbb{R}^n$, and to multi-input linear systems $\dot{x} = Ax + Bu$, where the control u is now a vector in $\mathbb{R}^m$, A and B are, respectively, an $n \times n$ and an $n \times m$ matrix, and the set U of control values is a convex polyhedron K. Under these conditions, it can be proved that *the controls that are bang-bang* (*i.e. piecewise constant with values in the set of vertices of* K) *form a sufficient class for time-optimality*.

The preceding discussion shows that the statement that "time-optimal controls are bang-bang" is true —after a slight reinterpretation— for linear control problems, but can fail in a rather strong way as soon as the slightest deviation from linearity is permitted, as in Example 12.2. (It is easy to see that in Example 12.2 the *only* control that goes from $(0,0)$ to $(1,0)$ in time 1 is u_*, so the class of bang-bang controls is not sufficient.) Stronger nonlinearities can produce even stronger pathologies, as we now show.

Example 12.5. Let $u_* : [0,T] \to [-1,1]$ be an arbitrary measurable function. Let $v(t) = \int_0^t u(s)\,ds$. Let K be the graph of v, i.e. $K = \{(t, v(t)) : 0 \le t \le T\}$. Let $\theta : \mathbb{R}^2 \to \mathbb{R}$ be a function of class C^∞ such that $\theta \equiv 0$ on K and $\theta > 0$ on $\mathbb{R}^2 \setminus K$. Consider the control system in $\mathbb{R}^3$ given by the equations $\dot{x}_1 = 1$, $\dot{x}_2 = u$, $\dot{x}_3 = \theta(x_1, x_2)$. As before, the values of the control u are required to belong to $[-1,1]$. Let $r = v(T)$. Let $\xi_*(t) = (t, v(t), 0)$. Then it is easy to see that ξ_* goes from $(0,0,0)$ to $(T, r, 0)$ in time T, and *every trajectory ξ that goes from $(0,0,0)$ to $(T,r,0)$ does so in time T and coincides with ξ_* or one of its time translates*. In particular, *u_* and its time translates are the only time-optimal controls that give rise to trajectories going from $(0,0,0)$ to $(T,r,0)$*. So *no class $\mathcal{U}$ can be sufficient for time-optimality unless it contains u_* or one of its translates.* ■

The point of Example 12.5 is that u_* is arbitrary, so one can get arbitrary unavoidable pathology with just simple examples in $\mathbb{R}^3$. This means, in particular, that *the only class of $[-1,1]$-valued controls defined on intervals of the form $[0,T]$ that is universally sufficient (i.e. sufficient for all systems) is the class of all controls.*

Notice, however, that the above construction does not work if one insists on staying within the class of *real-analytic* systems, since in that case the function θ will not exist unless K is very special. It turns out that

> *For real-analytic systems one can prove the existence of nontrivial universally sufficient classes. For example, the class of all controls that are real-analytic on an open dense subset of their domain is universally sufficient* (cf. [20], [21]).

So far, everything that has been said would suggest that linearity of a system has much to do with the question whether time-optimal controls are bang-bang. The next example shows that the situation is not so simple.

Example 12.6. Consider the system

$$\begin{aligned}\dot{x}_1 &= x_1^2 + u,\\ \dot{x}_2 &= 1 + x_1 + x_1^3 x_2 + 2x_1^3 x_3 - x_1^5 x_3 + 2x_1 x_3 u,\\ \dot{x}_3 &= x_1 x_2 - x_1^3 x_3,\end{aligned} \tag{12.3}$$

with control constraint $|u| \leq 1$. The system equations obviously look highly nonlinear. However, if one writes them in the $\dot{x} = f(x) + ug(x)$ form, as before, and computes iterated Lie brackets of f and g, one finds that $[g,[f,g]]$ is a linear combination of g and $[f,g]$ with smooth coefficients: $[g,[f,g]] = \alpha g + \beta[f,g]$ (actually, $[g,[f,g]] = 2g$), and that the vector fields g, $[f,g]$ and $[f,[f,g]]$ are linearly independent at every point. These facts can be exploited, using the adjoint equation in its intrinsic form (12.2). Let ξ_* be time-optimal, let u_* be a control giving rise to ξ_*, and let ζ be as in the conclusion of the MP. One then finds that

$$\dot{\varphi}(t) = \zeta(t).[f,g](\xi_*(t)) \tag{12.4}$$

and

$$\ddot{\varphi}(t) = \zeta(t).[f,[f,g]](\xi_*(t)) + u_*(t)\zeta(t).[g,[f,g]](\xi_*(t))$$

so

$$\ddot{\varphi}(t) = \sigma(t) + u_*(t)\alpha(\xi_*(t))\varphi(t) + u_*(t)\beta(\xi_*(t))\dot{\varphi}(t), \tag{12.5}$$

where $\sigma(t) = \zeta(t).[f,[f,g]](\xi_*(t))$. This equation implies that the set of zeros of φ is discrete. (Indeed: the absolute continuity of ζ together with (12.4) imply that φ is C^1 and $\dot{\varphi}$ is absolutely continuous. If $\varphi(t) = 0$ and $\dot{\varphi}(t) \neq 0$ then it is clear that t is an isolated zero of φ. If $\varphi(t) = 0$ and $\dot{\varphi}(t) = 0$ then $\sigma(t) \neq 0$, because $\zeta(t) \neq 0$ and g, $[f,g]$ and $[f,[f,g]]$ are linearly independent at $\xi_*(t)$. So (12.5) implies that $\ddot{\varphi}$ has constant sign and is bounded away from zero in a neighborhood of t. Therefore t is, once again, an isolated zero of φ.) Since all zeros of φ are isolated, u_* is bang-bang. So *all optimal controls are bang-bang*, exactly as in the linear case. ■

Examples 12.2, 12.3, 12.6 show that the amount of "nonlinearity" of a system is not a very good indicator of its control-theoretic properties. This should not be surprising, since properties such as the bang-bang nature of time-optimal arcs are clearly invariant under arbitrary nonlinear coordinate changes, whereas linearity itself is not. This means that one can easily produce plenty of examples of nonlinear systems whose time-optimal arcs are bang-bang by just taking linear systems and changing coordinates to render them "nonlinear."

It is clear that facts about the behavior of a system that are invariant under general nonlinear coordinate changes should be related to properties that have the same invariance. As Example 12.6 suggest, what matters is the *structural relations among the Lie brackets of f and g*, i.e. whether or not certain iterated brackets are linear combinations of some other brackets. The coordinate-free formulation in terms of Equation (12.6) is the one that makes it easiest to relate these structural properties to the necessary conditions of the MP.

We illustrate this by proving the following "simplified bang-bang theorem."

Theorem 12.1. *Consider a single-input control system $\Sigma : \dot{x} = f(x) + ug(x)$, $|u| \leq 1$, where f and g are smooth vector fields on a manifold M. Let g_k be defined inductively by $g_0 = g$, $g_{k+1} = [f, g_k]$. Assume that* (i) *for every k the vector field $[g, g_k]$ is a linear combination of $g_0, \ldots, g_k$ with smooth coefficients and* (ii) *for some $\overline{k}$ the vectors $g_0(x), \ldots, g_{\bar{k}}(x)$ span T_xM for all x. Then all time-optimal arcs of Σ are bang-bang.*

Proof. Write $[g, g_k] = \sum_{j=0}^{k} \alpha_k^j g_j$. For an optimal arc ξ_* arising from a control u_*, let ζ be a nontrivial minimizing adjoint vector. Define $\varphi_k(t) = \langle \zeta(t).g_k(\xi_*(t)) \rangle$, and write $\Phi = (\varphi_0, \ldots, \varphi_{\bar{k}})^\dagger$. Then $\Phi(t) = 0$ iff $\zeta(t) = 0$, so $\Phi(t) \neq 0$ for all t. On the other hand,

$$\dot{\varphi}_k = \varphi_{k+1} + u_* \sum_{j=0}^{k} \alpha_k^j \varphi_j \quad \text{for } k \in \{0, \ldots, \overline{k}\} \tag{12.6}$$

along $\Xi_* = (\xi_*, \zeta)$. Given $t \in \mathcal{D}(\xi_*)$ such that $\varphi_0(t) = 0$, let $k \in \{1, \ldots, \overline{k}\}$ be such that $\varphi_k(t) = c \neq 0$ but $\varphi_j(t) = 0$ for $j < k$. Since the φ_j are bounded, (12.6) implies that $\varphi_j(t+h) = O(|h|)$ as $|h| \to 0$ for $j < k$, and $\varphi_k(t+h) = c + O(|h|)$. But then (12.6) implies that $\varphi_{k-1}(t+h) = ch + O(h^2)$, and $\varphi_j(t+h) = O(h^2)$ for $j < k-1$. Repeated applications of (12.6) then yield $\varphi_{k-i}(t+h) = \frac{ch^i}{i!} + O(|h|^{i+1})$ for $i = 0, \ldots, k$. In particular, setting $i = k$, we find that $\varphi_0(t+h) = ah^k + o(|h|^k)$ for some $a \neq 0$. So t is an isolated zero of φ_0. Therefore all the zeros of the switching function φ_0 are isolated. As explained before, this implies that u_* is bang-bang. ∎

Remark 12.1. Theorem 12.1 can be extended and generalized in various ways. For example, for real-analytic systems one can almost dispense with Condition (ii). It suffices to make the much weaker hypothesis that (ii') *for every $x \in M$ the vector $f(x)$ belongs to the linear span of* $\{g_k(x) : k \in \mathbb{N}\}$. Indeed, every trajectory of Σ is contained in an integral manifold S of the Lie algebra of vector fields generated by f and g, so one can always work with the restriction of Σ to S, i.e. assume that M itself is an integral manifold. Then (i) and (iii) imply (ii).

Condition (iii) is essential, as shown by the system $\dot{x}_1 = 1$, $\dot{x}_2 = u$, for which $g_k = 0$ for $k > 0$, so $[g, g_k] = 0$ for all k, and therefore (i) holds. Clearly, for this system *every* trajectory is time-optimal, so there are in particular optimal controls that are not bang-bang.

On the other hand, if one changes the question as discussed earlier, and asks instead whether bang-bang optimal arcs form a *sufficient family* for time-optimality, then the answer is "yes" for analytic systems, assuming only Hypothesis (i). (In our trivial example $\dot{x}_1 = 1$, $\dot{x}_2 = u$, it is clear that one can always pick u_* to be bang-bang with at most one switching.) A stronger version of this result, under a hypothesis weaker than (i), and yielding a stronger conclusion (local uniform bounds on the number of switchings) is proved in [16]. The proof of the main result of [16] is based on a technical lemma with a very long proof. A much simpler proof of the lemma has been found by Dmitruk (personal communication). ∎

Remark 12.2. For single-input linear systems, $[g, g_k] = 0$ for all k. Theorem 12.1 therefore contains and generalizes the result of Example 12.3. Moreover, the property that $[g, g_k] = 0$ for all k is invariant under general nonlinear coordinate changes, whereas linearity itself is not. Since the bang-bang property is itself invariant, it can be argued that an "invariant reason" for a system to have this property provides a more satisfactory explanation than a noninvariant one. If so, then Theorem 12.1 can be seen as explaining, from a nonlinear perspective, why linear systems have the bang-bang property. ∎

13. Invariant Systems on Lie Groups

Let G be a Lie group, and let $L(G)$ denote the Lie algebra of G, thought of as the set of left-invariant vector fields on G. A *left-invariant vector field system* on G is a parametrized family $f = \{f_u\}_{u \in U}$ of members of $L(G)$. We will think of f as a map $U \to L(G)$, and write $f(u)$ instead of f_u.

A *control* for f is a function η from some interval to U such that $f \circ \eta$ is measurable. A control $\eta : I \to U$ gives rise to a time-varying vector field f_η with domain $G \times I$, defined by $f_\eta(g, t) = f(\eta(t))(g)$ for $g \in G$. The vector field f_η is LIB iff η is locally Lebesgue integrable, and in that case $f_\eta \in CVF_{\mathrm{LIC}^\infty}(G)$.

A control system generated by such an f will be, according to our general definitions, a triple $(G, \mathcal{U}, F)$ such that (i) whenever an interval I is the time domain of some control vector field that occurs in F, then every f_η arising from a piecewise constant control $\eta : I \to U$ occurs in the family F, and (ii) every CVF that occurs in F is the restriction to its domain of some f_η. We will limit ourselves to systems that are generated by f in the following stronger sense: (I) the parameter set $\mathcal{U}$ actually is a set of controls, and (II) for every $\eta \in \mathcal{U}$ the control vector field F_η is actually f_η. Specifying such a system amounts to specifying a set $\mathbf{I}$ of intervals and a set $\mathcal{U}$ of U-valued controls in such a way that $\mathcal{D}(\eta) \in \mathbf{I}$ for every $\eta \in \mathcal{U}$, and $\mathcal{U}$ contains all piecewise constant U-valued controls whose domain is in $\mathbf{I}$.

An "invariant Lagrangian" should then be a family of functions $\{L_u\}_{u\in U}$ such that each L_u is invariant under left translations. This, however, just means that each L_u is a constant function on G. So an *invariant Lagrangian* for an invariant family f is, simply, a real-valued function $u \to L(u)$ on U. In accordance with our general definitions, a controlled arc $\gamma = (\xi, \eta)$ will be acceptable for L if the function $L \circ \eta$ is locally integrable. So in this case acceptability is just a property of η, and does not involve ξ.

Given an invariant system $\Sigma = (G, \mathcal{U}, F)$ generated by f, and an invariant Lagrangian L for f, we can define $L_\eta(t) = L(\eta(t))$, and $J_L(\gamma) = \int_{\mathcal{D}(\xi)} L(\eta(t))\, dt$ for a controlled arc $\gamma = (\xi, \eta)$ of Σ such that η is acceptable.

For invariant systems with an invariant Lagrangian, the concepts of an *adjoint vector* and an L-adjoint vector for some L coincide, since the "Lagrangian part" of the Hamiltonian is independent of $g \in G$. Let $\eta : I \to U$ be such that $f \circ \eta$ is locally integrable, and let $\xi : I \to G$ be a trajectory for η. Let $\gamma = (\xi, \eta)$. An adjoint vector along γ is an absolutely continuous map $\zeta : I \to T^*G$ such that each $\zeta(t)$ is in $T^*_{\xi(t)}G$ and ζ satisfies the adjoint equation. The cotangent spaces T^*_gG can all be identified with the dual $L(G)^*$ of $L(G)$, via the maps $\theta_g : T^*_gG \to L(G)^*$ that assign to each $z \in T^*_gG$ the linear functional $\theta_g(z) : L(G) \to \mathbb{R}$ defined by $\theta_g(z)(X) = \langle z, X(g)\rangle$. With this identification, we can regard fields of covectors along a trajectory as $L(G)^*$-valued functions. Let ζ be a field of covectors along ξ, and let ζ_θ be the $L(G)^*$-valued function that corresponds to ζ under this identification, so that $\zeta_\theta(t) = \theta_{\xi(t)}(\zeta(t))$, i.e. $\langle \zeta_\theta(t), X\rangle = \langle \zeta(t), X(\xi(t))\rangle$ whenever $X \in L(G)$. The adjoint equation says that

$$\frac{d}{dt}\langle \zeta(t), X(\xi(t))\rangle = \langle \zeta(t), [f(\eta(t)), X](\xi(t))\rangle \tag{13.1}$$

for every smooth vector field X on G. If X is left-invariant, then $[f(\eta(t)), X]$ is left-invariant as well, so we can rewrite (13.1) as

$$\frac{d}{dt}\langle \zeta_\theta(t), X\rangle = \langle \zeta_\theta(t), [f(\eta(t)), X]\rangle. \tag{13.2}$$

So if ζ is an adjoint vector along γ then ζ_θ must satisfy (13.2) for all $X \in L(G)$. It is easy to show that, conversely, if ζ_θ satisfies (13.2) for all

$X \in L(G)$ then ζ is an adjoint vector along γ. From now on, we identify $\mathrm{Adj}^{\Sigma}(\gamma)$ with a space of $L(G)^*$-valued functions, so we think of ζ_θ, rather than ζ, as an "adjoint vector." We can then drop the subscript θ, and simply state that

An absolutely continuous map $\zeta : \mathcal{D}(\xi) \to L(G)^*$ is an adjoint vector along a controlled trajectory (ξ, η) of a system Σ generated by a left-invariant vector field system $f = \{f_u : u \in U\}$ on the Lie group G, if and only if

$$(13.3)\quad \frac{d}{dt}\langle \zeta(t), X\rangle = \langle \zeta(t), [f(\eta(t)), X]\rangle \qquad \text{for a.e. } t \text{ for all } X \in L(G).$$

This condition can also be expressed in terms of the dual of the adjoint representation of $L(G)$. For $Y \in L(G)$, let ad_Y denote the map $X \to [Y, X]$ from $L(G)$ to $L(G)$, and let $\mathrm{ad}_Y^* : L(G)^* \to L(G)^*$ denote the dual map. Then $\langle \zeta, [f(\eta(t)), X]\rangle = \langle \zeta, \mathrm{ad}_{f(\eta(t))}(X)\rangle = \langle \mathrm{ad}^*_{f(\eta(t))}(\zeta(t)), X\rangle$. So we have shown that

(13.3) *holds iff* $\dot\zeta(t) = \mathrm{ad}^*_{f(\eta(t))}(\zeta(t))$ *for a.e.* $t \in \mathcal{D}(\xi)$.

Example 13.1. Let A_1, A_2 be two linearly independent skew-symmetric 3×3 matrices. Consider the system $\dot R = R(u_1A_1 + u_2A_2)$, evolving in $G = SO(3)$, the group of 3×3 orthogonal matrices R such that $\det R = 1$, with control constraint $u_1^2 + u_2^2 \le 1$. A *control* is a measurable function $I \ni t \to (\eta_1(t), \eta_2(t))$ with values in $\{(u_1, u_2) : u_1^2 + u_2^2 \le 1\}$, whose domain I is a compact interval. Then $L(G)$ is naturally identified with $so(3)$, the Lie algebra of 3×3 skew-symmetric matrices, by assigning to each $A \in so(3)$ the vector field X_A given by $X_A(R) = RA$. Then $[X_A, X_B] = X_{[A,B]}$, where $[A, B] = AB - BA$.

Let us study the minimum time problem. If $[a, b] \ni t \to R(t) \in G$ is a minimum time trajectory, corresponding to a control $\eta = (\eta_1, \eta_2)$, then by the MP there exists a minimizing adjoint vector $\zeta : [a, b] \to L(G)^*$ such that $\zeta(t) \neq 0$ for all t, and the Hamiltonian is ≤ 0 along $(R(\cdot), \zeta)$. The Hamiltonian is $H(R, z, u) = u_1\langle z, X_{A_1}(R)\rangle + u_2\langle z, X_{A_2}(R)\rangle$. Then $H(R(t), \zeta(t), u) = u_1\langle \zeta(t), X_{A_1}(R(t))\rangle + u_2\langle \zeta(t), X_{A_2}(R(t))\rangle$. We rewrite this as $H(R(\cdot), \zeta(\cdot), u) = u_1\varphi_1 + u_2\varphi_2$, where the "switching functions" φ_i

are given by $\varphi_i(t) = \langle \zeta(t), X_{A_i}(R(t)) \rangle$. If we let $\rho(t) = \sqrt{\varphi_1(t)^2 + \varphi_2(t)^2}$, then the Hamiltonian minimization condition easily implies that $\eta_i(t) = -\frac{\varphi_i(t)}{\rho(t)}$ if $\rho(t) \neq 0$. The value of the Hamiltonian for $u = \eta$ then turns out to be $-\rho(t)$. The constancy of the Hamiltonian then says that ρ is constant. The adjoint equations say that $\dot{\varphi}_1(t) = -\eta_2(t)\varphi_3(t)$ and $\dot{\varphi}_2(t) = \eta_1(t)\varphi_3(t)$, where $\varphi_3(t) = \langle \zeta(t), X_{A_3} \rangle$, and $A_3 = [A_1, A_2]$.

If the constant ρ was equal to 0, then it would follow that φ_1 and φ_2 vanish identically. Since the three vector fields X_{A_i} form a basis of $L(G)$, and $\zeta(t) \neq 0$, the function φ_3 never vanishes. But then the equations $\dot{\varphi}_1 = -\eta_2\varphi_3$, $\dot{\varphi}_2 = \eta_1\varphi_3$ imply that $\eta_1(t) \equiv \eta_2(t) \equiv 0$. So our trajectory $R(\cdot)$ is just a constant G-valued function, and is therefore not time-optimal unless $b = a$.

If $b > a$, we have shown that ρ must be > 0. But then, if we write $\varphi_i = \rho\psi_i$ for $i = 1, 2, 3$, we see that $\eta_i = -\psi_i$ for $i = 1, 2$, and $\dot{\psi}_1 = -\eta_2\psi_3$, $\dot{\psi}_2 = \eta_1\psi_3$. On the other hand, the derivative of φ_3 can be computed using the adjoint equation once again. Since $\varphi_3 = \langle \zeta, X_{A_3} \rangle$, we have $\dot{\varphi}_3 = \eta_1\langle \zeta, X_{[A_1,A_3]} \rangle + \eta_2\langle \zeta, X_{[A_2,A_3]} \rangle$. Now, $[A_1, A_3] = c_{11}A_1 + c_{12}A_2$ and $[A_2, A_3] = c_{21}A_1 + c_{22}A_2$, for some constants c_{ij}. (Recall that $so(3)$ is isomorphic to $\mathbb{R}^3$ via an isomorphism that maps the Lie bracket to the cross product. Since the cross product $v_3 = v_1 \times v_2$ of two vectors is orthogonal to v_1 and v_2, it follows that the cross products $v_1 \times v_3$ and $v_2 \times v_3$ are orthogonal to v_3, so they are linear combinations of v_1 and v_2.) So

$$\dot{\varphi}_3 = c_{11}\eta_1\varphi_1 + c_{12}\eta_1\varphi_2 + c_{21}\eta_2\varphi_1 + c_{22}\eta_2\varphi_2. \tag{13.4}$$

If we divide through by ρ, use $\eta_i = -\psi_i$ for $i = 1, 2$, and define $\eta_3 = -\psi_3$, we see that the three functions η_i satisfy:

$$\dot{\eta}_1 = -\eta_2\eta_3, \qquad \dot{\eta}_2 = \eta_1\eta_3, \qquad \dot{\eta}_3 = \sum_{i,j=1,2} c_{ij}\eta_i\eta_j. \tag{13.5}$$

We have shown that *every optimal control* (η_1, η_2) *that is defined on an interval* $[a, b]$ *with* $b > a$ *satisfies the system* (13.5) *for some function* η_3. In particular, this implies that all optimal controls are smooth. ∎

Remark 13.1. Example 13.1 is a special case of the more general sub-Riemannian minimization problems to be considered in the Section 14. For these problems, minimizers are "normal" or "abnormal" extremals, and normal extremals are always smooth. Example 13.1 is "strongly bracket-generating" (cf. Strichartz [15]), and this implies that all minimizers are normal extremals and hence smooth. ■

14. Sub-Riemannian Minimizers

If M is a manifold, and E is a subbundle of TM, we use $L(E)$ to denote the Lie algebra of vector fields generated by $\Gamma^\infty(E)$. The subbundle E is called *bracket-generating* if $\{X(x) : X \in L(E)\} = T_xM$ for all $x \in M$. A *Riemannian metric* on E is a smooth map G that assigns to each $x \in M$ a strictly positive definite symmetric bilinear form $G(x)$ on $E(x)$. We will regard $G(x)$ as linear map from $E(x)$ to $E(x)^*$, so $G(x)v \in E(x)^*$ if $v \in E(x)$, and the $G(x)$-inner product of $v, w \in E(x)$ is then $\langle G(x)v, w\rangle$. We will write $||v||_G$ to denote the length $\sqrt{\langle G(x)v, v\rangle}$ of a vector $v \in E(x)$.

A *sub-Riemannian manifold* is a triple (M, E, G) such that M is a manifold, E is a bracket-generating subbundle of the tangent bundle TM of M, and G is a Riemannian metric on E. If $S = (M, E, G)$ is a sub-Riemannian manifold, then we define $\mathcal{A}(S)$, the class of S-*admissible* (or E-*admissible*) arcs, whose elements are the arcs ξ in M that satisfy $\dot{\xi}(t) \in E(\xi(t))$ for almost every $t \in \mathcal{D}(\xi)$. For an arc $\xi \in \mathcal{A}(S)$, we define the *length*

$$\Lambda(\xi) \stackrel{\text{def}}{=} \int_{\mathcal{D}(\xi)} ||\dot{\xi}(t)||_G \, dt. \tag{14.1}$$

An arc $\xi : [a, b] \to M$ that belongs to $\mathcal{A}(S)$ is a *minimizer* if $\Lambda(\xi) \leq \Lambda(\xi')$ for all $\xi' \in \mathcal{A}(S)$ that start at $\xi(a)$ and end at $\xi(b)$. We let $\mathcal{M}(S)$ be the set of all minimizers. We would like to find necessary conditions for an arc in $\mathcal{A}(S)$ to be a minimizer.

An arc $\xi \in \mathcal{A}(S)$ is *parametrized by constant times arc-length* (abbr. PCAL) if there exists a constant c such that $||\dot{\xi}(t)||_G = c$ for almost all $t \in \mathcal{D}(\xi)$. It is easy to see that every arc coincides, up to a time reparametrization, with a PCAL arc $\xi : [0, 1] \to M$. We let $\mathcal{A}([0, 1], S) = \{\xi \in$

$\mathcal{A}(S) : \mathcal{D}(\xi) = [0,1]\}$, and use $\mathcal{M}_{PCAL}(S)$ denote the class of all PCAL minimizers $\xi \in \mathcal{A}([0,1], S)$. Then it is easily shown that $\xi_* \in \mathcal{M}_{PCAL}(S)$ iff ξ_* minimizes the functional $\xi \to J_G(\xi) = \int_0^1 \|\dot{\xi}(t)\|_G^2 \, dt$ subject to $\partial\xi = \partial\xi_*$. So we will study the minimizers of J_G instead.

To be able to apply the MP, we proceed as in Section 11 and realize $\mathcal{A}([0,1], S)$ as the class of trajectories of a control system. As in Section 11, we do this by using a vector field system with a very large space of vector fields. Precisely, we take $U = \Gamma^\infty(E)$, and write $f_u(x) = f(x,u) = u(x)$. Then $f = \{f_u : u \in U\}$ is a vector field system. We then let $\mathcal{U}$ denote the set of all LIC$^\infty$ maps $\eta : [0,1] \to U$. (This means that our space $\mathcal{U}$ of controls is exactly $L^1([0,1], U)$, the space of integrable functions with values in the Fréchet space U.) If $\eta \in \mathcal{U}$, we write $F_\eta = \eta$, so $F_\eta(x,t) = \eta(t)(x)$. Now let $\Sigma = (M, \mathcal{U}, F)$. Then Σ is a control system generated by (M, U, f) in the sense of our definitions. It is easy to see that the set of trajectories of Σ is exactly $\mathcal{A}([0,1], S)$.

Next define $L_\eta(x,t) = \frac{1}{2} G(\eta(t)(x), \eta(t)(x))$. Then L is a Lagrangian for Σ, and an arc $\xi_* \in \mathcal{A}([0,1], S)$, arising from a control η_*, belongs to $\mathcal{M}_{PCAL}(S)$ iff it minimizes $\int_0^1 L_\eta(\xi(t), t) dt$ subject to $(\xi, \eta) \in \mathrm{Carc}(\Sigma)$, $\xi(0) = \xi_*(0)$, $\xi(1) = \xi_*(1)$.

Now let $\xi_* : [0,1] \to M$ belong to $\mathcal{M}_{PCAL}(S)$, and pick an $\eta_* : [0,1] \to U$ such that $\gamma_* = (\xi_*, \eta_*) \in \mathrm{Carc}(\Sigma)$. Then the MP tells us that there exist an absolutely continuous field ζ of covectors along ξ_* and a $\nu \in \{0,1\}$ such that $(\zeta, \nu) \neq (0,0)$ and ζ is a νL-adjoint vector along γ_* that satisfies the minimization condition. Therefore the following hold for a.e. $t \in [a,b]$:

$$\text{(MC)} \qquad \begin{aligned} \langle \zeta(t), \eta_*(t)(\xi_*(t)) \rangle + \nu L_{\eta_*}(\xi_*(t), t) &\le \langle \zeta(t), \eta(t)(\xi_*(t)) \rangle \\ &\quad + \nu L_\eta(\xi_*(t), t) \quad \text{for all } \eta \in \mathcal{U}, \end{aligned}$$

$$\text{(AE)} \qquad \begin{aligned} \frac{d}{dt}\langle \zeta(t), X(\xi_*(t)) \rangle &= \langle \zeta(t), [\eta_*(t), X](\xi_*(t)) \rangle \\ &\quad - \nu X L_{\eta_*,t}(\xi_*(t)) \quad \text{for all } X \in \Gamma^\infty. \end{aligned}$$

So far, the conditions on ζ appear to depend on the choice of a control η_* that generates ξ_*. We now seek to get rid of this dependence, as we did in Section 11 for the Riemannian case. This is most easily done for the minimization condition (MC). Indeed, since $\eta_*(t)(\xi_*(t)) = \dot{\xi}_*(t)$, (MC) is equivalent to

$$\text{(MC.b)} \qquad \langle \zeta(t), \dot{\xi}_*(t)\rangle + \frac{\nu}{2}G(\dot{\xi}_*(t), \dot{\xi}_*(t)) \leq \langle \zeta(t), v\rangle + \frac{\nu}{2}G(v,v) \quad \text{for all } v \in E(\xi_*(t)),$$

which does not involve η_* at all.

Next, we use the fact that G gives rise, for each x, to a linear map $G^{\#}(x) : T_x^*M \to E(x)$, characterized by the identity $G(G^{\#}(x)(z), v) = z(v)$ for $z \in T_x^*M$, $v \in E(x)$. (That is, $G^{\#}(x)(z)$ is the vector $w \in E(x)$ such that $G(w,\cdot) = z\lceil E(x)$. In the Riemannian case, $G(x)$ maps T_xM to T_x^*M and is invertible, and $G^{\#}(x)$ is precisely $G(x)^{-1}$, i.e. "the metric tensor with raised indices," usually denoted by g^{ij}. In the general sub-Riemannian case, $G(x)$ is invertible as a map from $E(x)$ to $E^*(x)$, so $G(x)^{-1}$ is a well defined map from $E^*(x)$ to $E(x)$, and $G^{\#}(x)$ is the composite of the restriction map $z \to z\lceil E(x)$ with $G(x)^{-1}$. In more traditional language, all this amounts to saying that the g^{ij} are well defined in the sub-Riemannian case even though the g_{ij} are not.)

Define $w(t) = G^{\#}(\xi_*(t))(\zeta(t))$. Then $\langle \zeta(t), v\rangle = G(w(t), v)$ for $v \in E(\xi_*(t))$. So (MC.b) is equivalent to

$$\text{(MC.c)} \qquad G(w(t), \dot{\xi}_*(t)) + \frac{\nu}{2}G(\dot{\xi}_*(t), \dot{\xi}_*(t)) \leq G(w(t), v) + \frac{\nu}{2}G(v,v) \quad \text{for all } v \in E(\xi_*(t)).$$

If $\nu = 0$, then (MC.c) is true iff $w(t) = 0$ for a.e. t. If $\nu = 1$, then (MC.c) holds iff $w(t) = -\dot{\xi}_*(t))$ for a.e. t. So (MC.c) is equivalent to

$$\text{(MC.d)} \qquad w(t) = -\nu\dot{\xi}_*(t).$$

Now, if $\zeta(t) = 0$ for some t, then $\nu = 1$. Since (MC.d) holds a.e., w is continuous, and $w(t) = 0$, the set $\{s : \|\dot{\xi}_*(t)\|_G \leq \delta\}$ has positive measure for each $\delta > 0$. Since ξ_* is PCAL, we see that $\dot{\xi}_*(t) = 0$ a.e., so ξ_* is a constant arc. (Constant arcs are of course minimizers, and satisfy (AE) and (MC.d) with $\zeta \equiv 0$, $\nu = 1$.) From now on, we assume that ξ_* is a nonconstant arc, so $\zeta(t) \neq 0$ for all t.

We distinguish two cases, depending of whether $\nu = 0$ or $\nu = 1$. A nonzero absolutely continuous field of covectors ζ along ξ_* that satisfies (MC.d) and (AE) with $\nu = 1$ (resp. $\nu = 0$) is called a *normal* (resp. *abnormal*) minimizing adjoint vector along ξ_*, and the curve (ξ_*, ζ) will then be called a *normal* (resp. *abnormal*) *extremal lift* of ξ_*. A trajectory ξ_* that has a normal (resp. abnormal) extremal lift is called a *normal* (resp. *abnormal*) *extremal*.

We begin by analyzing the normal case, i.e. we assume that $\zeta \neq 0$, $\nu = 1$, and (AE) and (MC.d) hold.

Let $H^S : T^*M \to \mathbb{R}$ be the function defined by

$$H^S(x,z) = -\frac{1}{2}||G^\#(x)(z)||_G^2. \tag{14.2}$$

Then H^S will be called the *Hamiltonian* of the sub-Riemannian manifold S. (This should not be confused with the *control theory Hamiltonian*, which is a family of Hamiltonian functions parametrized by the controls $\eta \in \mathcal{U}$.) If $g = (g_1, \ldots, g_m)$ is an orthonormal basis of sections of E on an open subset Ω of M, and we write $\widetilde{\Omega} = \pi_{T^*M}^{-1}(\Omega)$, then $H^S = -\frac{1}{2}\sum_{j=1}^m H_{g_j}^2$ on $\widetilde{\Omega}$. (Recall from Section 3 that if X is a vector field on M then H_X is the function $(x,z) \to z \cdot X(x)$ on T^*M.) If X is any smooth vector field on M, then $\{H^S, H_X\} = -\sum_{j=1}^m \{H_{g_j}, H_X\}H_{g_j} = -\sum_{j=1}^m H_{[g_j,X]}H_{g_j}$ on $\widetilde{\Omega}$. So

$$\begin{aligned}\{H^S, H_X\}(x,z) &= -\sum_{j=1}^m \langle z, [g_j, X](x)\rangle \cdot \langle z, g_j(x)\rangle \\ &= -\Big\langle z, \sum_{j=1}^m G(G^\#(x)(z), g_j(x))[g_j, X](x)\Big\rangle.\end{aligned}$$

If ω is a smooth 1-form on ω, and we write $Y = G^\#\omega$, then

$$\{H^S, H_X\}(x, \omega(x)) = \Big\langle \omega(x), -\sum_{j=1}^m G(Y(x), g_j(x))[g_j, X](x)\Big\rangle.$$

Clearly, $G(Y, g_j)[g_j, X] = [G(Y, g_j)g_j, X] + (XG(Y, g_j))g_j$. So

$$\sum_{j=1}^m G(Y(x), g_j(x))[g_j, X](x) = \Big[\sum_{j=1}^m G(Y, g_j)g_j, X\Big](x)$$

$$+\sum_{j=1}^{m}(XG(Y,g_j))(x)g_j(x)$$
$$=[Y,X](x)+\sum_{j=1}^{m}(XG(Y,g_j))(x)g_j(x).$$

Therefore

$$\begin{aligned}\{H^S,H_X\}(x,\omega(x)) = &-\langle\omega(x),[Y,X](x)\rangle\\ &-\Big\langle\omega(x),\sum_{j=1}^{m}(XG(Y,g_j))(x)g_j(x)\Big\rangle\\ = &-H_{[Y,X]}(x,\omega(x))-\sum_{j=1}^{m}(XG(Y,g_j))(x)\langle\omega(x),g_j(x)\rangle\\ = &\ H_{[X,Y]}(x,\omega(x))-\sum_{j=1}^{m}(XG(Y,g_j))(x)G(Y,g_j)(x)\\ = &\ H_{[X,Y]}(x,\omega(x))-\frac{1}{2}X\Big(\sum_{j=1}^{m}G(Y,g_j)^2\Big)(x)\\ = &\ H_{[X,Y]}(x,\omega(x))-\frac{1}{2}(XG(Y,Y))(x).\end{aligned}$$

So we have shown that

$$\{H^S,H_X\}(x,\omega(x))=H_{[X,Y]}(x,\omega(x))-\frac{1}{2}(XG(Y,Y))(x). \tag{14.3}$$

Returning now to our normal minimizer ξ_*, we can pick for each t a smooth 1-form ω_t such that $\omega_t(\xi_*(t))=\zeta(t)$ and $G^{\#}(\omega_t)=-\eta_*(t)$. (This is clearly possible, because $G^{\#}(\xi_*(t))(\zeta(t))=w(t)=-\eta_*(t)(\xi_*(t))$.) Then (14.3) implies that

$$\begin{aligned}\{H^S,H_X\}(\xi_*(t),\zeta(t)) &= H_{[\eta_*(t),X]}(\xi_*(t),\zeta(t))\\ &\quad-\frac{1}{2}(XG(\eta_*(t),\eta_*(t)))(\xi_*(t))\\ &=\langle\zeta(t),[\eta_*(t),X](\xi_*(t)\rangle-XL_{\eta_*,t}(\xi_*(t))\end{aligned} \tag{14.4}$$

for all $X\in\Gamma^\infty(TM)$. So the right-hand side of (AE) is equal to $\{H^S,H_X\}$

$(\xi_*(t), \zeta(t))$, so (AE) is equivalent to

$$\frac{d}{dt} H_X(\xi_*(t), \zeta(t)) = \{H^S, H_X\}(\xi_*(t), \zeta(t)) \quad \text{for all } X \in \Gamma^\infty(TM). \tag{14.5}$$

This is easily seen to be equivalent to the condition that

$$\text{(AE.N)} \qquad \Xi : t \to (\xi_*(t), \zeta(t)) \text{ is an integral curve of } H^S.$$

(It suffices to show that if $z \in T_x^*M \setminus \{0\}$ then every tangent vector $W \in T^*_{(x,z)}(T^*M)$ is of the form $\vec{H}_X(x,z)$ for some X, a fact that is easily verified in coordinates. So every covector $\widehat{W}$ at (x,z) is of the form $dH_X(x,z)$ for some X. Then (14.5) says that $dH_X.\dot{\Xi} = dH_X.\vec{H}^S$ along Ξ for every X, so $\dot{\Xi} = \vec{H}^S$ along Ξ.)

So we have proved that (MC) is equivalent to (MC.d), and that when these hold with $\nu = 1$ then (AE) is equivalent to (AE.N). It is easy to show that if $t \to (\xi(t), \zeta(t))$ is an integral curve of H^S, then ξ is admissible and PCAL and ζ satisfies (MC.d) with $\nu = 1$. (Relative to an orthonormal basis of sections $(g_1, \ldots, g_m)$ as before, Formula (14.2) implies that $H^S(x,z) = -\frac{1}{2}\sum_j \langle z, g_j(x)\rangle^2$. So $\dot{\xi}(t) = -\sum_{j=1}^m \langle \zeta(t), g_j(\xi(t)\rangle g_j(\xi(t))$, showing that ξ is admissible, and $\dot{\xi}(t) = -G^{\#}(\xi_*(t))(\zeta(t))$, so (MC.d) holds with $\nu = 1$. Finally, along (ξ, ζ) the function H^S is itself constant, and in view of (MC.d) the value of this constant is $-|\frac{1}{2}|\dot{\xi}(t)||_G^2$, so ξ is PCAL.)

To study the case when $\nu = 0$, we let $E^\perp$ denote the annihilator of E, so $E^\perp(x) = \{z \in T_x^*M : (\forall v \in E(x))(z \cdot v = 0)\}$. If $\omega \in \Gamma^\infty(E^\perp)$, $V \in \Gamma^\infty(E)$ and $X \in \Gamma^\infty(TM)$, then $\langle \omega, [V, X]\rangle(x)$ only depends, for a given X, on the values of ω and V at x. So each X induces a bilinear map $\theta_X(x) : E^\perp(x) \times E(x) \to \mathbb{R}$, such that $\theta_X(z,v) = \langle \omega, [V,X]\rangle(x)$ whenever $\omega \in \Gamma^\infty(E^\perp)$, $V \in \Gamma^\infty(E)$ are such that $\omega(x) = z$, $V(x) = v$.

In particular, (AE) says —since $\nu = 0$— that

$$\frac{d}{dt}\langle \zeta(t), X(\xi_*(t))\rangle = \langle \zeta(t), [\eta_*(t), X](\xi_*(t))\rangle \quad \text{for every } X \in \Gamma^\infty(TM). \tag{14.6}$$

On the other hand, (MC.d) implies that $\langle \zeta(t), v\rangle = 0$ for all $v \in E(\xi_*(t))$, so that $\zeta(t) \in E(\xi_*(t))^\perp$. Then (AE) is equivalent to

$$\text{(AE.A)} \quad \frac{d}{dt}\langle \zeta(t), X(\xi_*(t))\rangle = \theta_X(\zeta(t), \dot{\xi}_*(t)) \qquad \text{for every } X \in \Gamma^\infty(TM).$$

We summarize the above discussion in the following

Theorem 14.1. *Let $S = (M, E, G)$ be a sub-Riemannian manifold, and let Σ be the corresponding control system. Let $\xi_* : [0,1] \to M$ be a nonconstant S-admissible PCAL arc. Let η_*, be a control in $\mathcal{U}$ that generates ξ_* as a trajectory of Σ. Let ζ be a nonzero absolutely continuous field of covectors along ξ_*. Then*

- **[I]** *ζ satisfies* (MC) *and* (AE) *with $\nu = 1$ iff ζ satisfies* (MC.d) *with $\nu = 1$ and* (AE.N), *and in that case $H^S \neq 0$ along Ξ.*
- **[II]** *ζ satisfies* (MC) *and* (AE) *with $\nu = 0$ iff $\zeta(t) \in E^\perp(\xi_*(t))$ for all t and* (AE.A) *holds.*

More generally, if $(\xi_*, \zeta) : [0,1] \to T^*M$ is any integral curve of H^S, then ξ_* is admissible and PCAL and ζ satisfies (MC) *with $\nu = 1$ and* (AE). *If $H^S \neq 0$ along (ξ_*, ζ), then ξ_* is nonconstant, so ξ_* is a normal extremal.* ∎

Theorem 14.1 implies in particular —both for $\nu = 1$ and $\nu = 0$— that the pair of conditions (MC), (AE) is independent of the choice of η_*. Moreover, the nonconstant normal extremals of Σ are exactly the projections on M of the integral curves of H^S, from which it follows that *all the normal extremals are smooth*.

For a Riemannian metric, the abnormal case cannot arise, because $\zeta \in E^\perp$ implies $\zeta = 0$, which is excluded. However, in the sub-Riemannian case Theorem 14.1 does not rule out the possibility that some minimizers might be abnormal, or even *strictly abnormal* (i.e. abnormal but not normal[6]).

The possibility of the existence of strictly abnormal minimizers was incorrectly excluded in some papers such as [14], where it was asserted that every sub-Riemannian minimizer is a normal extremal, with the obvious

---

[6] A trajectory can be normal and abnormal at once, since it can have two different lifts with different values of the abnormal multiplier.

corollary that every sub-Riemannian minimizer is smooth. Since then, it was recognized, e.g. in [15], that the proof of this theorem had a gap, due precisely to the possibility that $\nu = 0$. However, it was not until the work of R. Montgomery [12] that an example was found of a minimizer that is actually abnormal and not normal. The following much simpler example was presented in Liu-Sussmann [9].

Example 14.1. We let $M = \mathbb{R}^3$, with the usual coordinates x_1, x_2, x_3, and take E to be the kernel of the 1-form $\omega = x_1^2\, dx_2 - (1 - x_1)\, dx_3$. Since ω never vanishes, E is a smooth two-dimensional subbundle of TM. The vector fields f, g, given by $f = \frac{\partial}{\partial x_1}$, $g = (1 - x_1)\frac{\partial}{\partial x_2} + x_1^2 \frac{\partial}{\partial x_3}$ form a global basis of sections of E. The Lie brackets $[f, g]$, $[f, [f, g]]$ are easily computed, and it turns out that $f(x)$, $g(x)$, $[f, g](x)$ and $[f, [f, g]](x)$ span $\mathbb{R}^3$ at every point. Therefore E is bracket-generating.

We now define a metric G on E by $G = dx_1^2 + h(x_1)(dx_2^2 + dx_3^2)$, where $h(x_1) = \frac{1}{(1-x_1)^2+x_1^4}$. Then the vector fields f and g form an orthonormal basis of sections of E.

It is easy to see that for every a, b such that $a < b$ the restriction to $[a, b]$ of the curve $\xi_* : t \to (0, t, 0)$ is an abnormal extremal which is not normal. The following result was proved in [9]:

Theorem 14.2. *The restriction of ξ_* to any interval of length $\leq \frac{2}{3}$ is optimal.* ∎

A much more detailed study of abnormal extremals for sub-Riemannian manifolds with a rank-two distribution E has been carried out in [10].

Remark 14.1. The above example shows that not all minimizers are normal extremals, so the smoothness of all minimizers cannot be proved by using the fact that the normal extremals are smooth, and it is natural to conjecture that nonsmooth minimizers exist. However, even though it is very easy to produce examples of abnormal sub-Riemannian extremals that are not smooth, all examples found so far of abnormal extremals that can be proved to be optimal have turned out to be smooth. As of this writing, the question whether there exist nonsmooth sub-Riemannian minimizers remains unanswered. ∎

15. Miscellaneous Examples

In this section we present several examples, in order to illustrate some of the technical issues raised in the preceding discussion.

Our first example pertains to the question of the *locality* of the conditions of the MP, and shows that there is a striking contrast between the classical Calculus of Variations and the more general optimal control situation. Indeed, it is rather obvious that the Euler-Lagrange equations (EL) are *local*, in the sense that a trajectory ξ defined on an interval I is a solution if and only if every $t \in I$ has a neighborhood $N(t)$ relative to I such that the restriction $\xi\lceil N(t)$ of ξ to $N(t)$ is a solution. The following elementary example shows that the analogous property is not true in general for optimal control problems.

Example 15.1 (Nonlocality of the MP conditions). Consider the minimum time problem for the control system in $\mathbb{R}^2$ given by $\dot{x}_1 = x_2$, $\dot{x}_2 = u$, $|u| \leq 1$. Notice that for this system a trajectory ξ uniquely determines the control that generates it, so in this case there is no need to distinguish between "trajectories" and "controlled trajectories." The Hamiltonian is given by $H = z_1 x_2 + z_2 u$. The adjoint equations for a covector field $t \to \zeta(t) = (\zeta_1(t), \zeta_2(t))$ along a trajectory $t \to \xi(t) = (\xi_1(t), \xi_2(t))$ say that $\dot{\zeta}_1 = 0$ and $\dot{\zeta}_2 = -\zeta_1$. So ζ_2 is a linear function of t, which is not identically zero. The minimization condition says that $u(t) = 1$ when $\zeta_2(t) < 0$ and $u(t) = -1$ when $\zeta_2(t) > 0$. Since the function ζ_2 is linear, it can have at most one zero. This means that $u(t)$ is "bang-bang" (i.e. with values in the set $\{-1, 1\}$) with at most one switching. Moreover, the constant ζ_1 determines the sense of the switching. Indeed, the condition that $H \leq 0$ implies that at a switching time $\bar{t}$ we have $\zeta_1 \xi_2(\bar{t}) \leq 0$. If $\zeta_1 > 0$ — which can only happen if $\xi_2(\bar{t}) \leq 0$— then $\dot{\zeta}_2(\bar{t}) < 0$, which means that ζ_2 changes sign from positive to negative, and u switches from -1 to 1. Similarly, if $\zeta_1 < 0$ —which is only possible if $\xi_2(\bar{t}) \geq 0$— then $\dot{\zeta}_2(\bar{t}) > 0$, which means that ζ_2 changes sign from negative to positive, so u switches from 1 to -1. (If $\zeta_1 = 0$ then ζ_2 is a nonzero constant, so there are no switchings.) Therefore *if a trajectory ξ is an extremal, then the control*

$u(t)$ is bang-bang with at most one switching. A switching from $u = 1$ to $u = -1$ can only occur in the upper half-plane $x_2 \geq 0$, and a switching from $u = -1$ to $u = 1$ can only occur in the lower half-plane $x_2 \leq 0$. Conversely, it is easy to see that every trajectory ξ that satisfies these conditions is an extremal. Indeed, suppose that ξ corresponds to a control $u(\cdot)$ that has exactly one switching. Suppose that this switching occurs at time $\bar{t}$, and has the "correct" direction, i.e. either (i) u switches from 1 to -1 and $\xi_2(\bar{t}) \geq 0$ or (ii) u switches from -1 to 1 and $\xi_2(\bar{t}) \leq 0$. Let $t \to \zeta_2(t)$ be a linear function that vanishes at $\bar{t}$ and has the correct slope, positive in Case (i) and negative in Case (ii). Let $\zeta_1 = -\dot{\zeta}_2$. Then $\zeta = (\zeta_1, \zeta_2)$ is a minimizing adjoint vector along ξ, along which $H \leq 0$. So ξ is an extremal. Now, if $\xi : [a,b] \to \mathbb{R}^2$ is any trajectory corresponding to a bang-bang control with *two* switchings, and if these two switchings have the correct direction, then it is clear that ξ is locally extremal (i.e. every t has a relative neighborhood $N(t)$ in $[a,b]$ such that $\xi\lceil N(t)$ is an extremal) but ξ is not an extremal. (An example of such a ξ is the curve $\xi : [0,4] \to \mathbb{R}^2$ that starts with $\xi(0) = (0,0)$ and corresponds to a control $u(\cdot)$ such that $u(t) = -1$ for $0 \leq t < 1$, then $u(t) = 1$ for $1 \leq t < 3$ and, finally, $u(t) = -1$ for $3 \leq t \leq 4$.) ■

Our second example illustrates how the formulation based on continuous vector fields —due, as indicated in the introduction, to S. Lojasiewicz Jr.— adds extra strength to the MP.

Example 15.2. Consider the system $\dot{x} = u(1+x^{1/3})$, $\dot{y} = x$, $|u| \leq 1$, in two dimensions. Let $\xi(t) = (0,0)$ for $0 \leq t \leq 1$. Let us study, using the MP, whether ξ is a boundary trajectory. (It is quite easy to verify directly that ξ is not a boundary trajectory, but our goal here is to see how various forms of the MP would deal with the problem.) Our trajectory corresponds to the vector field $(0,x)$, which is smooth. On the other hand, all the other vector fields of the system, corresponding to nonzero values of the control, are continuous but fail to be locally Lipschitz. So all standard formulations of the MP, including the "nonsmooth" ones, fail to be applicable. The version given here does apply, however, since it only requires the reference vector field to be LIL (and our case this vector field is actually of class C^1), while the other vector fields are allowed to be just continuous. The Hamiltonian is $H = z_1 u(1 + x^{1/3}) + z_2 x$, and the Hamiltonian minimization condition along ξ requires $\zeta_1 \equiv 0$. The adjoint equation along ξ says $\dot{\zeta}_1 = -\zeta_2$, so $\zeta_2 \equiv 0$, contradicting the nontriviality condition. So ξ is not a boundary trajectory. ∎

The next two examples illustrate the importance of the distinction between "trajectories" and "controlled trajectories," and the usefulness of formulating the MP in a way that allows for large classes of control vector fields. The first example illustrates the possibility that a trajectory ξ may arise from two different controls η and $\widetilde{\eta}$, in such a way that the controlled trajectory (ξ, η) is a controlled extremal —for, say, the controllability problem— but $(\xi, \widetilde{\eta})$ is not. Naturally, since the property of being a boundary trajectory only depends on the trajectory ξ itself, but not on the choice of a control that generates it, the fact that there is one control $\widetilde{\eta}$ generating ξ in such a way that $(\xi, \widetilde{\eta})$ is not a controlled extremal already suffices to imply that ξ is not a boundary trajectory, but this would not have been detected by just looking at the controlled trajectory (ξ, η). The second example —borrowed from the work of Kaskosz and Lojasiewicz [6]— illustrates an even more dramatic possibility. It can happen that a trajectory ξ of a system Σ is such that (ξ, η) satisfies the necessary conditions of the MP —applied to Σ— for all possible choices of a control η that generates ξ, but that there is a larger system $\widehat{\Sigma}$ that gives rise to the same trajectories as Σ, for which there is a control $\widehat{\eta}$ such that $(\xi, \widehat{\eta})$ is a LIC^1-controlled trajectory that violates the conditions of the MP for $\widehat{\Sigma}$. In other words,

the power of the MP can be enhanced by applying it to a suitably enlarged system rather than to the original system. This is possible in our setting, because our formulation of the MP is stated for essentially arbitrary collections of vector fields, so the addition of many new vector fields does not take us out of the range of applicability of the MP.

Example 15.3. Consider the system $\dot{x}_1 = u_1$, $\dot{x}_2 = u_2 x_1$, where the controls u_1, u_2, take values in the interval $[-1,1]$, so that the control space U is the two-dimensional square $[-1,1]^2$. The curve $t \to \xi(t) = (0,0)$, $0 \le t \le 1$, is obviously a trajectory, corresponding to the control η given by $\eta(t) \equiv (0,0)$, and also to the control $\widetilde{\eta}$ given by $\widetilde{\eta}(t) = (0,1)$. We want to know if ξ is a boundary trajectory.

The Hamiltonian is $H = z_1u_1 + z_2u_2x_1$. If $t \to \zeta(t) = (\zeta_1(t), \zeta_2(t))$ is an adjoint vector along (ξ,η) or $(\xi,\widetilde{\eta})$, then ζ must satisfy $\dot{\zeta}_1 = -u_2\zeta_2$, $\dot{\zeta}_2 = 0$. Moreover, the minimization of the Hamiltonian requires that $\zeta_1 \equiv 0$, since $u_1 = 0$ minimizes $\zeta_1(t)u_1$ for $u_1 \in [-1,1]$. For the controlled trajectory (ξ,η), the control u_2 vanishes, so we can satisfy the adjoint equations by taking $\zeta_1 = 0$, $\zeta_2 = -1$, and then it is easy to verify that the minimization condition is satisfied. On the other hand, for $(\xi,\widetilde{\eta})$, u_2 is equal to 1, and therefore the only way to satisfy the adjoint equations with $\zeta_1 = 0$ is to let $\zeta_2 = 0$ as well, contradicting the requirement that $\zeta \ne 0$. Therefore the controlled trajectory (ξ,η) is a controlled extremal, but $(\xi,\widetilde{\eta})$ is not. Since $(\xi,\widetilde{\eta})$ is not a controlled extremal, we conclude that $(0,0)$ is an interior point of the reachable set from $(0,0)$, a fact which is also easily verified directly. ∎

Example 15.4. Let Σ be the control system $\dot{x}_1 = u_1$, $\dot{x}_2 = |x_1| + u_2 - 1$, where the controls u_1, u_2, take values in the interval $[-1,1]$, so that the control space U is the square $[-1,1]^2$. (The controls are all measurable U-valued functions.) The curve $t \to \xi(t) = (0,0)$, $0 \le t \le 1$ is obviously a trajectory, corresponding to the control η given by $\eta(t) \equiv (0,1)$. Once again, we want to find out whether ξ is a boundary trajectory.

Notice, to begin with, that η is the only control giving rise to ξ. The corresponding reference vector field is $(0, |x_1|)$, which is Lipschitz. The version of the MP stated in this paper therefore applies to Σ, and one might suspect that, if ξ is not a boundary trajectory, then it might be possible to detect this by applying Theorem 8.2. We show that this is not the case, by proving:

(A) (ξ,η) *satisfies the necessary condition of Theorem* 8.2 *applied to* Σ.

(B) *There is a larger system* $\widehat{\Sigma}$, *having exactly the same trajectories as* Σ, *and a control $\widehat{\eta}$ for $\widehat{\Sigma}$, such that $(\xi, \widehat{\eta})$ does not satisfy the necessary condition of Theorem* 8.2 *applied to $\widehat{\Sigma}$.*

To prove (A), we first seek to apply Theorem 8.2 to Σ. The Hamiltonian is $H = z_1 u_1 + z_2(|x_1| + u_2 - 1)$. The adjoint differential inclusion says that $\dot{\zeta}_1 \in [-|\zeta_2|, |\zeta_2|]$ and $\dot{\zeta}_2 = 0$. For our reference controlled trajectory (ξ, η), these conditions can be satisfied by taking $\zeta_1 \equiv 0$ and $\zeta_2 \equiv -1$. With this choice, the Hamiltonian minimization condition is satisfied as well, and the resulting value of the Hamiltonian is zero. So ξ passes the test of Theorem 8.2 applied to Σ.

Now consider the vector field f given by $f(x_1, x_2) = (0, x_1)$ for $x_1 \geq -1$, $f(x_1, x_2) = (0, |x_1| - 2)$ for $x_1 \leq -1$. Then $f(x_1, x_2)$ belongs to the set $\{(u_1, |x_1| + u_2 - 1) : (u_1, u_2) \in [-1, 1]^2\}$ for every (x_1, x_2). So, if we add this new vector field to our system, the resulting extended system has exactly the same trajectories.

Precisely, define $\widehat{U} = U \cup \{\#\}$, where $\#$ is a new object. Let $\widehat{\mathcal{U}}$ be the set of measurable $\widehat{U}$-valued functions defined on compact intervals. For $\widehat{\eta} \in \widehat{\mathcal{U}}$, define $\widehat{F}_{\widehat{\eta}}(x, t) = f(x, t)$ if $\widehat{\eta}(t) = \#$, $\widehat{F}_{\widehat{\eta}}(x, t) = (u_1, |x_1| + u_2 - 1)$ if $\widehat{\eta}(t) = (u_1, u_2) \in U$. Then define $\widehat{\Sigma} = (\mathbb{R}^2, \widehat{\mathcal{U}}, \widehat{F})$.

It is clear that ξ is an integral curve of f. Therefore $(\xi, \widehat{\eta}) \in \mathrm{Ctraj}(\widehat{\Sigma})$, if we define $\widehat{\eta}$ by letting $\widehat{\eta}(t) \equiv \#$. Moreover, the adjoint equation along the controlled arc $(\xi, \widehat{\eta})$ says $\dot{\zeta}_1 = -\zeta_2$ and $\dot{\zeta}_2 = 0$. The Hamiltonian minimization condition still requires ζ_1 to be $\equiv 0$. But this forces ζ_2 to be $\equiv 0$ as well, violating the nontriviality condition. So ξ fails to satisfy the necessary condition of Theorem 8.2 applied to $\widehat{\Sigma}$. ∎

Example 15.5. The purpose of this example is to show that the integral bounds on partial derivatives or Lipschitz constants of the reference CVF are crucial for the validity of the MP. We will do this by exhibiting an example of a system $\dot{x} = f(x, u, t)$, $x \in \mathbb{R}$, where f is jointly continuous in (x, u, t), globally bounded, and smooth in x, u for each t, for which there exists a trajectory $\xi_* : [0, 1] \to \mathbb{R}$ that goes from 1 to 0 and corresponds to a control u_*, in such a way that (a) ξ_* is smooth, (b) u_* is constant, (c) the function $t \to \frac{\partial f}{\partial x}(\xi_*(t), u_*(t), t)$ is bounded (actually, $\equiv 0$), and (d) 0 is on the boundary of the reachable set from 1, but (e) the conclusion of the MP is not true.

The only hypothesis of the MP that fails to hold in this case is the integral bound $\left|\frac{\partial f}{\partial x}(x, u_*(t), t)\right| \leq h(t)$, with h integrable. In the MP, this is required to hold for the reference control for x near the the reference trajectory. (In the example, the bound does hold along the trajectory itself, because $\frac{\partial f}{\partial x}$ vanishes there. So in particular we are showing that it is important that the bound hold not just on ξ_* but on a whole neighborhood.)

We let $\varphi : \mathbb{R}^2 \to \mathbb{R}$ be a continuous function such that (i) $\varphi(x,t)$ is smooth in x for each fixed t, (ii) φ is bounded, (iii) $\varphi((1-t)^3, t) = 3(1-t)^2$ for $0 \leq t \leq 1$, (iv) $\frac{\partial \varphi}{\partial x}((1-t)^3, t) = 0$ for $0 \leq t \leq 1$, (v) $\varphi(x,t) \leq 0$ for $x \leq 0$. (For example, we can construct such a φ by starting with the function $\psi : \mathbb{R} \to \mathbb{R}$ given by $\psi(s) = \frac{2}{1+(s-1)^2} - 1$, which is smooth and such that $\psi(1) = 1$, $\psi(s) \leq 0$ for $s \leq 0$, and $|\psi(s)| \leq 1$ for all s. Then plug in $s = \frac{x}{(1-t)^3}$, for $t < 1$. The resulting function $(x,t) \to \theta(x,t) = \psi\big(\frac{x}{(1-t)^3}\big)$, defined for $t < 1$, is still smooth in both variables, and such that $\theta(x,t) \leq 0$ for $x \leq 0$, $\theta((1-t)^3, t) = 1$, and $|\theta(x,t)| \leq 1$. Now multiply this by $3(1-t)^2$. The resulting function $(x,t) \to \sigma(x,t) = 3(1-t)^2\psi\big(\frac{x}{(1-t)^3}\big)$ is still defined for $t < 1$ and smooth in both variables. Moreover, $\sigma(x,t) \leq 0$ for $x \leq 0$, $\sigma((1-t)^3, t) = 3(1-t)^2$, and $|\sigma(x,t)| \leq 3(1-t)^2$. Extend σ to $\mathbb{R}^2$ by letting $\sigma(x,t) = 0$ for $t \geq 0$, and multiply by a smooth nonnegative function of t that equals 1 for $t \geq 0$ and 0 for $t \leq -1$. Then φ satisfies all our conditions.)

Consider the control system $\dot{x} = -u^2\varphi(x,t)$, $u \in U$, where $U = \mathbb{R}$, or $U = [-K, K]$, where $K > 1$. The class of control functions can be fairly arbitrary, e.g. L^1 or L^∞ or the class of all piecewise continuous controls.

Let $\xi_*(t) = (1-t)^3$, $0 \leq t \leq 1$. Then $\varphi(\xi_*(t), t) = 3(1-t)^2 = -\dot{\xi}_*(t)$. So ξ_* is a trajectory of our system, corresponding to the control $u_*(t) \equiv 1$. Clearly, $\xi_*(0) = 1$, $\xi_*(1) = 0$. So 0 is reachable from 1 in time 1. On the other hand, it is easy to see that no point $x \in (\infty, 0)$ can be reached from 1, because every trajectory ξ satisfies $\dot{\xi}(t) \geq 0$ whenever $\xi(t) < 0$. So $\xi_*(1)$ belongs to the boundary of the reachable set from 1. If the MP applied, it would follow that there exists a nowhere vanishing AC function

$\zeta : [0,1] \to \mathbb{R}$ that satisfies the adjoint equation and is such that the Hamiltonian $H(x,z,u,t) = -u^2 z\varphi(x,t)$, evaluated for $x = \xi_*(t)$, $z = \zeta(t)$, is minimized as a function of $u \in U$ by $u = u_*(t)$, i.e. $u = 1$. On the other hand, the function $U \ni u \to -u^2\zeta(t)\varphi(\xi_*(t),t)$ can only be minimized by $u = 1$ if $\zeta(t)\varphi(\xi_*(t),t) = 0$. Since $\varphi(\xi_*(t),t) = 3(1-t)^2 \neq 0$ for $0 \leq t < 1$, we conclude that $\zeta(t) \equiv 0$, contradicting the nontriviality condition.

Clearly, $\frac{\partial \varphi}{\partial x} = 0$ along ξ_*, so our example has all the promised features.

Notice that *a function φ that satisfies our five conditions cannot possibly obey an integral bound $|\frac{\partial \varphi}{\partial x}(x,t)| \leq h(t)$, $h \in L^1$, on a neighborhood of* $(0,0)$. Indeed, assume there was such a bound. Let P_t denote the point $((1-t)^3, t)$, and let S_t denote the segment in the (x,t)-plane joining $(0,t)$ to P_t. Since $\varphi(0,t) \leq 0$, S_t must contain a point $Q_t = (q_t, t)$ where $\varphi = 0$. But then the Mean Value Theorem implies that there is an R_t in the segment from Q_t to P_t such that $\varphi(P_t) = ((1-t)^3 - q_t)\frac{\partial \varphi}{\partial x}(R_t)$. Since $0 \leq q_t < (1-t)^3$, and $\varphi(P_t) = -3(1-t)^2$, we conclude that $\left|\frac{\partial \varphi}{\partial x}(R_t)\right| \geq \frac{3}{1-t}$. Since $R_t \to (0,0)$ as $t \to 1-$, we see that h cannot be integrable near $t = 1$. ∎

16. Conclusion: Beyond the Maximum Principle

We close this introduction to the Maximum Principle by pointing out that this very powerful result is only the beginning of a very rich theory. It can be shown by means of very simple examples that the Maximum Principle may fail to detect controllability or to exclude optimality in situations where other standard mathematical tools would easily settle the question.

Example 16.1. Consider the system $\dot{x}_1 = u_1$, $\dot{x}_2 = u_2$, $\dot{x}_3 = u_1 x_2$, with controls u_1, u_2 such that $|u_1| \leq 1$, $|u_2| \leq 1$. Consider the trajectory ξ given by $\xi(t) = (0,0,0)$, and corresponding to the control $\eta \equiv (0,0)$. Once again, we ask if ξ is a boundary trajectory. It is easy to show that the necessary conditions of the Maximum Principle are satisfied along (ξ, η). Moreover, η is the only control giving rise to ξ. In fact, it can be shown that the necessary conditions of the Maximum Principle will still be satisfied for any enlarged system that has the same trajectories, so the enlargement procedure considered in the previous section will fail to detect the fact that ξ is not a boundary trajectory.

However, this fact can be detected by elementary means using standard tools from other areas of mathematics. Precisely, let us write our system as $\dot{x} = u_1 f_1(x) + u_2 f_2(x)$, where f_1 is the vector field $(1, 0, 0)$ and f_2 has components $(0, 1, x_1)$. The system dynamics directly tells us that motion in the directions of f_1 and f_2 and their negatives is possible. To find out if a full neighborhood of the system can be reached, we compute the Lie bracket of these two vector fields, and find that $[f_1, f_2] = (0, 0, 1)$, So f_1, f_2 and $[f_1, f_2]$ are linearly independent. This implies that motion in a third direction is possible, so ξ is not a boundary trajectory. ■

The previous example shows that, as far as controllability is concerned (and similar examples can easily be given for optimality problems) the Maximum Principle only captures one aspect of a much more complex situation, while simple tools from differential geometry can deal with cases that are beyond the power of the various versions and extensions of Pontryagin's result. The controllability of the system of Example 16.1 is a particular case of a general theorem relating controllability to Lie algebraic conditions, known as "Chow's Theorem," whose use was advocated by R. Hermann as far back as 1963 (cf. [5]), i.e. shortly after the publication of the book [13] that first presented the Maximum Principle. The natural conclusion to draw is that to realize its full power, the Maximum Principle has to be combined with differential geometric tools. This, of course, is what the burgeoning field of "differential geometric optimal control theory" is about, and the many successes of this line of research in the last twenty years confirm the validity of this conclusion.

REFERENCES

[1] F. Albrecht, *Topics in Control Theory*, Springer-Verlag, Berlin-Heidelberg-New York, 1968.

[2] L. D. Berkovitz, *Optimal Control Theory*, Springer-Verlag, New York, 1974.

[3] F. H. Clarke, *The Maximum Principle under minimal hypotheses*, S.I.A.M. J. Control Optim. **14** (1976), 1078–1091.

[4] F. H. Clarke, *Optimization and Nonsmooth Analysis*, Wiley Interscience, New York, 1983.

[5] R. Hermann, *On the accessibility problem in control theory*, in: International Symposium on Nonlinear Differential Equations and Nonlinear Mechanics, J. P. LaSalle and S. Lefschetz Eds., Academic Press, New York (1963).

[6] B. Kaskosz and S. Lojasiewicz Jr., *A Maximum Principle for generalized control systems*, Nonlinear Anal. TMA **9** (1985), 109–130.

[7] E. B. Lee and L. Markus, *Foundations of Optimal Control Theory*, J. Wiley, New York, 1967.

[8] G. Leitmann, *An Introduction to Optimal Control*, McGraw-Hill, New York, 1966.

[9] W.S. Liu and H.J. Sussmann, *Abnormal sub-Riemannian minimizers*, in: Differential Equations, Dynamical Systems and Control Science, K. D. Elworthy, W. N. Everitt, and E. B. Lee Eds., Lect. Notes Pure Appl. Math., Vol. 152. M. Dekker, New York (1993), 705–716.

[10] W. S. Liu and H. J. Sussmann, *Shortest paths for sub-riemannian metrics on rank-*2 *distributions*, to appear in the Memoirs of the American Math. Society.

[11] J. Macki and A. Strauss, *Introduction to Optimal Control Theory*, Springer-Verlag, New York, 1982.

[12] R. Montgomery, *Geodesics which do not satisfy the geodesic equations*, 1991 preprint.

[13] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, E. F. Mischenko, *The Mathematical Theory of Optimal Processes* (translated by K. N. Trigoroff, L. W. Neustadt, editor), J. Wiley, 1962.

[14] R. Strichartz, *Sub-Riemannian Geometry*, J. Diff. Geom. **24** (1986), 221–263.

[15] R. Strichartz, *Corrections to 'Sub-Riemannian Geometry'*, J. Diff. Geom. 30, **2**, (1989), 595–596.

[16] H. J. Sussmann, *A bang-bang theorem with bounds on the number of switchings*, S.I.A.M. J. Control and Optimization **17** (1979), 629–651.

[17] H. J. Sussmann, *The structure of time-optimal trajectories for single-input systems in the plane: the C^∞ nonsingular case*, S.I.A.M. J. Control and Optimization **25** (1987), 433–465.

[18] H. J. Sussmann, *The structure of time-optimal trajectories for single-input systems in the plane: the general real-analytic case*, S.I.A.M. J. Control and Optimization **25** (1987), 868–904.

[19] H. J. Sussmann, *Regular synthesis for time-optimal control of single-input real-analytic systems in the plane*, S.I.A.M. J. Control and Optimization 25, No. **5** (1987), 1145–1162.

[20] H. J. Sussmann, *A weak regularity theorem for real analytic optimal control problems*, Revista Matemática Iberoamericana **2**, No. 3 (1986), 307–317.

[21] H. J. Sussmann, *Trajectory regularity and real analyticity: some recent results*, Proc. 25th I.E.E.E. Conference on Decision and Control, Athens, Greece (1986), 592–595.

[22] H. J. Sussmann, *A strong version of the Lojasiewicz Maximum Principle*, in: Optimal Control of Differential Equations, Nicolai H. Pavel Ed., M. Dekker Inc., New York (1994).

[23] H. J. Sussmann, *A strong version of the Maximum Principle under weak hypotheses*, Proc. 33rd IEEE Conf. Dec. Control, Orlando, FL 1994, 1950–1956.

Appendix: List of Abbreviations

AC: absolutely continuous

Adj(γ): the set of all adjoint vectors along the controlled curve $\gamma = (\xi, F)$, i.e. Adj(γ, 0)

Adj(γ, L): the set of all L-adjoint vectors along the controlled curve γ

$\mathrm{Adj}^{\Sigma}(\gamma)$, $\mathrm{Adj}^{\Sigma}(\gamma, L)$: the set of all adjoint (resp. L-adjoint) vectors along the controlled trajectory γ of a control system Σ

$\mathrm{Adj}^{\Sigma}_{smin}(\gamma)$: $\mathrm{Adj}^{\Sigma}_{smin}(\gamma, 0)$

$\mathrm{Adj}^{\Sigma}_{smin}(\gamma, L)$: the set of all strongly minimizing $\zeta \in \mathrm{Adj}^{\Sigma}(\gamma, L)$

$\mathrm{Adj}^{\Sigma}_{wmin}(\gamma)$: $\mathrm{Adj}^{\Sigma}_{wmin}(\gamma, 0)$

$\mathrm{Adj}^{\Sigma}_{wmin}(\gamma, L)$: the set of all weakly minimizing $\zeta \in \mathrm{Adj}^{\Sigma}(\gamma, L)$

AE: adjoint equation

AI: adjoint inclusion

$ARC(M)$: the set of all arcs in M

$ARC(I, M)$: the set of all arcs in M with domain I

Carc(Σ): the set of all controlled arcs of a system $\Sigma = (M, \mathcal{U}, F)$

$\mathrm{Carc}_{\mathcal{P}}(\Sigma)$: the set of $\gamma \in \mathrm{Carc}(\Sigma)$ that have Property $\mathcal{P}$

$\mathrm{Carc}_{wa}(\Sigma, L)$: the set of all $\gamma \in \mathrm{Carc}(\Sigma)$ that are weakly acceptable for L

$C^{\dagger}$: the polar of a cone C

Ccrv($\mathcal{F}$): the set of all controlled curves $\gamma = (\xi, F)$ such that $F \in \mathcal{F}$

$\mathrm{Ccrv}_{\mathcal{P}}(\mathcal{F})$: the set of $\gamma \in \mathrm{Ccrv}(\mathcal{F})$ that have Property $\mathcal{P}$

CF: Carathéodory function

CF(M): the set of all CF's on M

$CF_{\mathcal{P}}(M)$: the set of all CF's on M that have Property $\mathcal{P}$

$CF_{\mathcal{P}}(I,M)$: the set of all $\varphi \in CF_{\mathcal{P}}(M)$ with time domain I

$CFV_{\mathcal{P}}(M)$: the set of all CVF's on M that have Property $\mathcal{P}$

$CFV_{\mathcal{P}}(I,M)$: the set of all $f \in CFV_{\mathcal{P}}(M)$ with time domain I

$C^k(M)$: the space of real-valued functions of class C^k on M

C. of V.: Calculus of Variations

$CRV(M)$, $CRV_{LAC}(M)$: the set of all curves (resp. LAC curves) in M

$CRV(I,M)$, $CRV_{LAC}(I,M)$: the set of $\xi \in CRV(M)$ (resp.$CRV_{LAC}(M)$) with domain I

Ctraj(Σ): the set of all controlled trajectories of a control system Σ

Ctraj$_{\mathcal{P}}(\Sigma)$: the set of all $\gamma \in$ Ctraj(Σ) that have Property $\mathcal{P}$

CVF: control vector field

CVF(M): the set of all CVF's on M

$\mathcal{D}(\cdot)$: domain of

$DEL(E,F)$: the CVF obtained from F by removing $F\lceil E$

F^L: $M \times I \ni (x,t) \to F^L(x,t) = (F(x,t), L(x,t)) \in TM \times \mathbb{R}$

$\mathcal{F}(\Sigma)$, $\widehat{\mathcal{F}}(\Sigma)$: the set $\{F_\eta : \eta \in \mathcal{U}\}$

$\widehat{\mathcal{F}}(\Sigma)$: the set of CVF's that occur in Σ, cf. Section 5

$G(\cdot)$: graph of

$\vec{H}$: the Hamilton vector field corresponding to the function H

H_F: $H_{F,0}$, i.e. $H_F(x,z,t) = z.F(x,t)$

$H_{F,L}$: (for $F \in CVF(M)$ and $L \in CF(M)$) the time-varying function on T^*M given by $H_{F,L}(x,z,t) = H_F(x,z,t) + L(x,t)$

$\mathcal{H}^{\Sigma,L}$ $(\stackrel{\text{def}}{=}\{H_\eta^{\Sigma,L}\}_{\eta\in\mathcal{U}})$: the Hamiltonian function associated with a control system $\Sigma = (M,\mathcal{U},F)$ and a Lagrangian L for Σ

$\mathcal{I}$: the set of all nontrivial intervals

IAE: inhomogeneous AE

IAI: inhomogeneous AI

IC: Integral curve

IC(F): the set of integral curves of the CVF F

IC($\mathcal{F}$): $\cup_{F\in\mathcal{F}}$ IC(F)

$INS(E,F,G)$: the CVF obtained from F by inserting G over E

IVE: inhomogeneous VE

$IVV(\xi,\mathcal{L},y)$, $IAV(\xi,\mathcal{L},z)$: the sets of solutions of $\mathcal{L}^{vec}v = y$ and $\mathcal{L}^{cov}\zeta = z$, respectively

$J^k(E)$: the bundle of k-jets of sections of the bundle E

$J^k_x(E)$: the set of k-jets at x of sections of the bundle E

LAC: locally absolutely continuous

LB: locally bounded

$\mathbf{LD}^{LI}(\xi)$: the class of LI connections along ξ

$\mathbf{LD}[M]$: the bundle $\cup_{x\in M}\mathbf{LD}(x)$

$\mathbf{LD}[[M]]$: the bundle over TM with fiber $\mathbf{LD}(x,v)$ at (x,v)

$\mathbf{LD}(x)$: $\cup_{v\in T_xM}\mathbf{LD}(x,v)$

$\mathbf{LD}(x,v)$: the set of all Lie differentiations at x in the direction of v

$L_{x,X}$: the element of $\mathbf{LD}(x,X(x))$ defined by $L_{x,X}(Y) = [X,Y](x)$

$\mathbf{LD}(\xi)$: the set of all Lie differentiations along a curve ξ

LI: locally integrable

LIB: locally integrably bounded

LIBU: LIB with uniqueness

LIC^k: locally integrably of class C^k

LIL: locally integrably Lipschitz

LL: locally Lipschitz

$\mathcal{L}_X$: the map $M \ni x \to L_{x,X} \in \mathbf{LD}(x)$

$\mathcal{L}^\gamma$: the connection canonically associated with a LIC^1-controlled curve $\gamma = (\xi, F)$

$\boldsymbol{\mathcal{L}}^\gamma$: the set of connections canonically associated with a LIL-controlled curve $\gamma = (\xi, F)$

$\boldsymbol{\mathcal{L}}^{\gamma,L}$: the set of all measurable selections of $\partial F^L \circ \gamma$

$\mathcal{L}^{vec}$, $\mathcal{L}^{cov}$: the maps defined in Section 4

$L^1_{\mathrm{loc},+}(I)$: the set of all nonnegative locally Lebesgue integrable functions on I

MP: Maximum Principle

$\mathbb{N}$: the set of integers ≥ 0

$\overline{\mathbb{N}}$: the set $\mathbb{N} \cup \{\infty\}$

ODE: ordinary differential equation

PCAL: parametrized by constant times arc-length

$\mathbb{R}_+$: the half-line $[0, \infty)$

$\overline{\mathbb{R}}$: $\mathbb{R} \cup \{+\infty\}$

$\overline{\mathbb{R}}_+$: $\mathbb{R} \cup \{+\infty, -\infty\}$

$\mathbb{R}^n$: the space of real n-component *column* vectors

$\mathcal{R}^\Sigma(x)$: The Σ-reachable set from x

$\mathcal{R}^\Sigma_{[a,b]}(x)$: The Σ-reachable set from x over $[a, b]$

$\mathcal{R}^\Sigma_{[a,b\pm\alpha]}(p)$: $\bigcup\Big\{\mathcal{R}^\Sigma_{[a,c]}(p) : b - \alpha \leq c \leq b + \alpha\Big\}$

$SUB(E, F, G)$: the CVF obtained from F by substituting G over E

$\mathcal{TD}(F)$: time domain of F

$\mathcal{TD}(F_1, \ldots, F_m)$: $\mathcal{TD}(F_1) \cap \ldots \cap \mathcal{TD}(F_m)$

TM: the tangent bundle of M
T^*M: the cotangent bundle of M
$T^{\#}M$: T^*M with the zero section removed
T_xM: the tangent space of M at x
T_x^*M: the cotangent space of M at x
$\mathcal{U}_{\mathcal{P},F}$, or $\mathcal{U}_{\mathcal{P},\Sigma}$: the class of all $\mathcal{P}$-admissible $\eta \in \mathcal{U}$
VE: variational equation
$VV(\xi,\mathcal{L})$, $AV(\xi,\mathcal{L})$: the sets $IVV(\xi,\mathcal{L},0)$, $IAV(\xi,\mathcal{L},0)$, respectively
$\Gamma(E)$: the set of sections of the bundle E
$\Gamma(E,\xi)$: the set of sections of the bundle E along the curve ξ
$\Gamma^{\mathcal{P}}(E)$, $\Gamma^{\mathcal{P}}(E,\xi)$: the set of $\sigma \in \Gamma(E)$ (resp. $\sigma \in \Gamma(E,\xi)$) that have Property $\mathcal{P}$
$\Gamma^k(E)$, $\Gamma^k(E,\xi)$: $\Gamma^{C^k}(E)$, $\Gamma^{C^k}(E,\xi)$
π_E: the canonical projection of a bundle E onto its base
$\pi_{\mathbf{LD},x}$: the natural projection $\mathbf{LD}(x) \to T_xM$
χ_S: indicator function of the set S
$\Sigma\lceil\widehat{\mathcal{U}}$: the restriction of the control system $\Sigma = (M,\mathcal{U},F)$ to $\widehat{\mathcal{U}}$, if $\widehat{\mathcal{U}}$ is a subset of $\mathcal{U}$
$\Sigma\lceil\mathcal{P}$: $\Sigma\lceil\mathcal{U}_{\mathcal{P},\Sigma}$
Σ^L: the augmented system associated with a control system $\Sigma = (M,\mathcal{U},F)$ and a Lagrangian $L = \{L_\eta : \eta \in \mathcal{U}\}$
Ω_M: the canonical symplectic form on T^*M
$\{\cdot,\cdot\}$: the Poisson bracket
$[\cdot,\cdot]$: the Lie bracket
$\lceil$: restriction
$\dagger$: transpose