Differential-Geometric Methods: a Powerful Set of New Tools for Optimal Control

Héctor J. Sussmann[1]

Department of Mathematics, Rutgers University
New Brunswick, NJ 08903

Abstract. Differential-geometric methods can and have been successfully used in control theory, not only to pose and solve new problems, but also to get new insights and prove new results about classical problems of optimal control. In particular, they yield new results on local controllability, continuity properties of value functions, the structure of reachable sets, and the properties of optimal trajectories. For large classes of systems given by controlled differential equations with real-analytic right-hand sides, these methods make it possible —in conjunction with the theory of subanalytic sets—- to prove piecewise analyticity of the value function and the existence of a piecewise smooth optimal feedback.

1 Introduction

Since its early days in the 1950's, the development of finite-dimensional, deterministic control theory has proceeded along essentially two main parallel tracks. The first one is the study of *linear systems*, and of a host of related problems such as static and dynamic stabilization, linear-quadratic control, linear time-optimal control, various "geometric" problems —e.g. tracking, disturbance decoupling, disturbance attenuation, noninteracting control— and linear identification and adaptive control. The second one is the study of *general* nonlinear systems $\dot{x} = f(x, u)$, possibly linear in the control u, posssibly required to satisfy some technical conditions such as differentiability of f, but not subject to any special "algebraic" requirements. Not surprisingly, the first line of inquiry has led, thanks to the algebraic structure, to a very rich body of theory for a very narrow class of systems, while the second one has produced some theorems of great generality, whose ability to yield useful results for specific nonlinear problems often turns out to be disappointingly limited.

However, since the early 1960's a new line of research began to emerge, starting with the pioneering work of R. Hermann [3] in 1963, followed later by that of C. Lobry [6] in 1970, and then by that of Brockett [2], Krener [5], Sussmann-Jurdjevic, and others in the early 1970's, and by a veritable explosion in the late 70's and the 80's, some of whose achievements are described, e.g., in the books [4,8,10]. This research makes systematic use of *differential-geometric methods* (henceforth abbreviated as DGM's). Our purpose here is to provide a brief outline of some of the applications of DGM's, using simple examples and stressing the comparison with classical approaches.

Differential geometry studies *differentiable manifolds endowed with "geometric" structures such as connections, Riemannian metrics, symplectic forms, foliations, group structures (Lie groups), group actions, bundle structures, etc.* For this

[1] This work was supported in part by NSF Grant DMS-8902994.

purpose, it has developed a general language and a distinctive philosophy, referred to here as the "differential geometric point of view," and outlined in Section 2 below.[2]

The distinctive feature of the work using DGM's in control theory is the adoption of this point of view, rather than the use of the specific structures of traditional differential geometry. DGM's are applied in control theory in three major ways:

1. To get new insights into *classical* problems of optimal control theory, such as local and global controllability, and the structure and properties of optimal trajectories, reachable sets, value functions, and optimal synthesis.
2. To study nonlinear control problems that are intrinsically differential-geometric, because (i) their natural state space is a differentiable manifold —such as $\mathbb{R}^3 \times SO(3)$ for the kinematics of a rigid body, $\mathbb{R}^3 \times SO(3) \times \cdots \times SO(3)$ for the kinematics of a system of several rigid bodies linked by spherical joints, or the tangent bundles of the above spaces for more realistic control problems where the controls are accelerations and torques— and/or (ii) the dynamics and controls are naturally related related to the geometry.
3. To develop general theories of "geometric" nonlinear problems such as feedback linearization, feedback equivalence, smooth feedback stabilization, tracking, path finding, observer design, disturbance decoupling, noninteracting control, disturbance attenuation —including nonlinear analogues of H^∞ theory— and other nonlinear versions of various well-known problems of linear control theory, for which the language and tools of differential geometry are essential even at the level of problem formulation (cf. the books [4,8]).

In this note we concentrate mainly on the first of the three types of applications. Our goal is to show that *DGM's are already useful and necessary even to study classical optimal control questions that do not appear to be "differential-geometric,"* and to argue our case we will present several simple examples of classical problems and show how DGM's provide new insights.

2 The Differential-Geometric Point of View

Suppose we want to prove that, if A is an $m \times n$ real matrix, then there are nonsingular square matrices P, Q of the appropriate sizes such that PAQ has, for some k, the 2×2 block form $\begin{pmatrix} I_k & 0 \\ 0 & 0 \end{pmatrix}$, where I_k is the identity matrix. An *algebraist* would do it by showing that A can by transformed into the desired form by elementary row and column operations, and that every such operation corresponds to left or right multiplication by a nonsingular matrix. A *geometer* would regard A as the matrix of a linear map $L : \mathbb{R}^n \ni x \to L(x) = Ax \in \mathbb{R}^m$,[3] and then would forget that $\mathbb{R}^n$ is $\mathbb{R}^b$ and $\mathbb{R}^m$ is $\mathbb{R}^m$, and regard both $\mathbb{R}^n$ and $\mathbb{R}^m$ as abstract linear spaces V, W. He would then let $W_1 = L(V)$, write $V = V_1 \oplus V_2$, where V_2 is the kernel of L, pick a basis $f_1, \ldots, f_k$ of V_1, and a basis $f_{k+1}, \ldots, f_n$ of V_2, combine them to get a basis $\mathbf{f} = (f_1, \ldots, f_n)$ of V, define $g_i = L(f_i)$ for $i = 1, \ldots, k$, show that $(g_1, \ldots, g_k)$ is a basis of W_1, extend this to a basis $\mathbf{g} = (g_1, \ldots, g_m)$ of W, and then observe that the matrix of L with respect to $\mathbf{f}$ and $\mathbf{g}$ has the desired block form.

---

[2] Defining differential geometry as the study of certain geometric objects is, of course, circular, but the circularity is unavoidable, since —for this as for any other area of mathematics— the list of structures commonly regarded by practitioners as part of a subject is determined by its history rather than by a general definition.

[3] Here, and throughout the paper, vectors in $\mathbb{R}^n$ are written as columns.

This elementary example captures the essence of the geometric point of view. Rather than work with spaces of tuples of numbers, geometers like to think about points in abstract spaces and about "geometric objects" such as lines, subspaces, linear maps (if the space has a linear structure that makes these things meaningful), or curves and surfaces and fields of directions in more general settings. Points are not tuples of numbers, even though they can be represented as such in various ways, by choosing coordinates. Geometers intensely dislike coordinates, and use them sparingly. No coordinate system is regarded as "natural" or preferred. This has a cost, namely, that one cannot define geometric objects by referring to a particular coordinate system. A geometric definition should preferably be formulated in coordinate-free language. If this cannot be done, and coordinate systems have to be used, then they should *all* be used and treated equally. A geometric object $\mathcal{O}$ should have a well defined *coordinate representation* $\mathcal{O}^\kappa$ with respect to every coordinate system κ. Naturally, one does not not require that $\mathcal{O}^\kappa$ be the same for all κ, but that the various representations be related in appropriate ways. For example, on an n-dimensional *linear space* V over $\mathbb{R}$, the coordinate systems are linear 1-1 onto maps $\kappa : V \to \mathbb{R}^n$. Then any two such systems $x = \kappa(v) \stackrel{\text{def}}{=} v^\kappa$, $y = \tilde{\kappa}(v) \stackrel{\text{def}}{=} v^{\tilde{\kappa}}$, are linearly related, i.e. the coordinate change map $L : v^{\tilde{\kappa}} \to L(v^\kappa) = \tilde{\kappa} \circ \kappa^{-1}$ is linear. For this reason one can do *linear algebra*, and define linear objects such as lines, planes and segments, and linear concepts such as linear independence. (For example, linear independence of k vectors $v_1, \ldots, v_k$ can either be defined using coordinates —e.g. by the condition that the matrix $(v_1^\kappa, \ldots, v_k^\kappa)$ have rank k, which would then have to be proved to be independent of the choice of κ— or in coordinate-free language, by requiring that $c_1 = \ldots = c_k = 0$ whenever $\sum_{i=1}^k c_i v_i = 0$. Either definition is acceptable, but geometers always prefer coordinate-free formulations.)

Differential geometry operates on abstract spaces M of points that are not thought of as tuples of numbers, but can be represented in various ways by means of tuples of numbers. These *coordinate representations* are called *charts.* (Formally, an n-dimensional chart on a set X is a one-to-one map κ from a subset $D(\kappa)$ of X —the *domain* of κ— onto an open subset $Im(\kappa)$ of $\mathbb{R}^n$.) Contrary to what happens in the linear case, a chart κ need not assign coordinates to all possible points, but one requires that every point be in the domain of some chart. (A set $\mathcal{A}$ of charts with this property is called an *atlas.*) Most importantly, one does not require that the (partially defined) transformations $\tilde{\kappa} \circ \kappa^{-1}$ relating two charts be linear, and so linear concepts such as lines and planes are not well defined. On the other hand, many other classes of objects are well defined. Exactly which objects can be defined depends, of course, on the atlas $\mathcal{A}$ used to endow M with a structure. The smaller $\mathcal{A}$ is, the easier it is for a definition to be invariant under coordinate changes. So reducing an atlas $\mathcal{A}$ —i.e. replacing $\mathcal{A}$ by a smaller atlas $\tilde{\mathcal{A}}$— corresponds to *enriching* a structure, and enlarging the class of geometric objects.

Differential geometry is specifically interested in *differentiable structures.* A n-dimensional *structure of class* C^k on M is an atlas $\mathcal{A}$ on M such that the domain $\kappa(D(\kappa) \cap D(\tilde{\kappa}))$ of every associated coordinate change $L = \tilde{\kappa} \circ \kappa^{-1}$ is open in $\mathbb{R}^n$, and L is of class C^k there. (As usual, we allow $k = +\infty$ and $k = \omega$, and "C^ω" means "real-analytic.") A set M endowed with a C^k structure of class is a C^k *manifold.* If $k \geq 1$ then M is a *differentiable manifold.*

A simple example of a differentiable manifold is the two-dimensional sphere S^2, defined to be the set of points (x_1, x_2, x_3) such that $x_1^2 + x_2^2 + x_3^2 = 1$. One can define a chart κ whose domain is the open upper hemisphere $H^+ = \{x : x_3 > 0\}$ by observing that, on H^+, $x_3 = +\sqrt{x_1^2 + x_2^2}$, so each point $p = (x_1, x_2, x_3) \in H^+$ is determined by $\kappa(p) = (x_1, x_2)$. This, of course, is only one chart, and its domain does not cover the whole sphere. To get a true atlas, notice that H^+ is a neighborhood of the "north pole" $(0,0,1)$. However (thinking geometrically!) every $p \in S^2$ is the "north pole" for some Cartesian coordinate system in $\mathbb{R}^3$. So what we have done for $(0,0,1)$ works for any $p \in S^2$, and we can get a chart covering a neighborhood of p. So we have produced enough charts to cover all points, i.e. we have constructed an atlas $\mathcal{A}$. If $\kappa, \tilde{\kappa} \in \mathcal{A}$ arise from Cartesian coordinates y and z in $\mathbb{R}^3$, then $z_i = \sum_{j=1}^{3} a_{ij} y_j$, and so if $\kappa(p) = (y_1, y_2)$, $\tilde{\kappa}(p) = (z_1, z_2)$, then $z_i = a_{i1}y_1 + a_{i2}y_2 + a_{i3}\sqrt{y_1^2 + y_2^2}$ if $p \in D(\kappa) \cap D(\tilde{\kappa})$. So all the coordinate changes are C^ω, and S^2 is a C^ω manifold.

The sphere S^2 has a well defined *tangent space* T_pS^2 at each point p, consisting of all those vectors in $\mathbb{R}^3$ that are starting directions of C^1 curves from p that stay in S^2. In other words, $v \in T_pS^2$ iff $v = \dot{\gamma}(0)$, where $\gamma : [0, \varepsilon) \to \mathbb{R}^3$ is C^1, $\gamma(0) = p$, and $\gamma(t) \in S^2$ for all t. Since $||\gamma(t)|| = 1$, it is easy to see that $\gamma(0) \perp p$. Conversely, for any $v \in \mathbb{R}^3$ such that $v \perp p$ we can construct a curve γ, e.g. by letting $\gamma(t) = \nu(p + tv)$ (where $\nu(w) = w/||w||$). Then $\gamma(t) = (p + tv)(\sqrt{1+t^2})$ —using $v \perp p$— so that $\gamma(t) = p + tv + o(t)$ and $\dot{\gamma}(0) = v$). So T_pS^2 is exactly the set of $v \in \mathbb{R}^3$ that are orthogonal to p. Notice that T_pS^2 varies as p changes. In particular, on a sphere it does not make sense to compute, for instance, the sum of two velocity vectors v_1, v_2 at different points p_1, p_2, since v_1 and v_2 live in different linear spaces.

The definition of tangent vectors for S^2 can be extended to arbitrary differentiable manifolds M. A "tangent vector" at $p \in M$ is defined intrinsically by first defining what it means for a curve $\gamma : I \to M$ to be C^1, and for two C^1 curves $\gamma, \delta : [0, \varepsilon) \to M$ such that $\gamma(0) = \delta(0) = p$ to be *tangent* at time 0. (This is done by referring to a chart κ with $p \in D(\kappa)$, and showing that the definition does not depend on the choice of κ.) Tangency is then an equivalence relation among C^1 curves starting at p, and the *tangent vectors at* p are the equivalence classes. If $\gamma : [0, \varepsilon) \to M$ is a C^1 curve from p whose class is v, then we write $v = \dot{\gamma}(0)$.

The set of tangent vectors at p is the *tangent space* at p, denoted by T_pM. Then every $v \in T_pM$ has a well defined representation v^κ with respect to every chart κ with $p \in D(\kappa)$ (obtained by picking any C^1 curve $\gamma : [0, \varepsilon) \to M$ with $\dot{\gamma}(0) = v$ and letting $v^\kappa = \dot{\gamma}^\kappa(0)$, where $\gamma^\kappa = \kappa \circ \gamma$). Of course, v^κ depends on κ. But,

$$\text{if } \kappa_1\,,\ \kappa_2 \text{ are charts and } p \in D(\kappa_1) \cap D(\kappa_2)\,,\ \text{then } v^{\kappa_2} = J(\kappa_2, \kappa_1)(p^{\kappa_1})v^{\kappa_1}\ , \quad (1)$$

where $J(\kappa_2, \kappa_1)$ is the Jacobian matrix of the map $\kappa_2 \circ \kappa_1^{-1}$. Conversely, if we specify a collection $\{v^\kappa\}$ of vectors in $\mathbb{R}^n$, for all charts κ such that $p \in D(\kappa)$, and these vectors transform according to (1), then the v^κ are the representations of a unique $v \in T_pM$, so we could equally well have defined a tangent vector to be such a collection $\{v^\kappa\}$. As explained before, both definitions are acceptable in differential geometry, but geometers prefer the first one because it is more "cordinate-free".

The dual space T_p^*M of T_pM is the *cotangent space* of M at p, and its elements are called *covectors* at p. A covector $\lambda \in T_p^*M$ has coordinate representations λ^κ, that are taken to be *row vectors*. The λ^κ obey the following transformation law:

$$\text{if } \kappa_1\,,\ \kappa_2 \text{ are charts and } p \in D(\kappa_1) \cap D(\kappa_2)\,,\ \text{then } \lambda^{\kappa_2} = \lambda^{\kappa_1} J(\kappa_1, \kappa_2)(p^{\kappa_2})\ . \quad (2)$$

Notice that *vectors and covectors are intrinsically different objects*, because their transformation laws (1) and (2) are different, If φ is a real-valued C^1 function on some neighborhood of p —the concepts of "neighborhhod" and "C^1 function" are well defined, using charts— then we can define its *gradient* $\nabla\varphi(p)$ at p to be the functional $\lambda \in T_p^*M$ given by $\langle \lambda, v \rangle = \frac{d}{dt}\Big|_{t=0} \varphi(\gamma(t))$, where γ is any C^1 curve from p such that $\dot{\gamma}(0) = v$. Then *the gradient of a function is a covector, not a vector.*

Having defined tangent vectors, it is obvious how to define vector fields and integral curves. A *vector field* is a function f that to every $p \in M$ assigns a tangent vector $f(p) \in T_pM$. If $I \subseteq \mathbb{R}$ is an interval, we call a C^1 curve $\gamma : I \to M$ an *integral curve* of f if $\dot{\gamma}(t) = f(\gamma(t))$ for all $t \in I$, where $\dot{\gamma}(t)$ is defined to be $\dot{\delta}_t(0)$, and $\delta_t(s) = \gamma(t+s)$. The concept of a C^k vector field is well defined on any C^{k+1} manifold. In particular, on a C^1 manifold we can talk about C^0 vector fields, and the usual local existence theorem for solutions of initial value problems $\dot{x} = f(x)$, $x(t_0) = x_0$ holds. If uniqueness also holds (for example, if f is C^1) then we say that f is a $C^{0,u}$ vector field, and use exponential notation for the flow of f, i.e. we write xe^{tf} for $\gamma_{x,f}(t)$, where $\gamma_{x,f}$ is the maximal integral curve of f which goes through x at time 0.

In control theory we study differential equations $\dot{x} = f(x,u)$ involving a parameter u, and for there to be true control u must take at least two values. So we have to work with *systems*, involving *at least two vector fields.* Clearly, then, one should look for operations that combine vector fields to get new ones. Besides the linear operations, there is an important binary operation that, to any two C^1 vector fields f, g, associates their *Lie bracket* $[f,g]$, defined in coordinates by $[f,g]^\kappa = \frac{\partial g^\kappa}{\partial x}.f^\kappa - \frac{\partial f^\kappa}{\partial x}.g^\kappa$ (writing both f^κ and g^κ as columns of functions, so that $\frac{\partial f^\kappa}{\partial x}$ and $\frac{\partial g^\kappa}{\partial x}$ are square matrices of functions). An easy computation shows that the above representations transform like a vector field, so there is a well defined vector field $[f,g]$. According to the basic differential-geometric philosophy, the Lie bracket ought to have an important control-theoretic meaning for non-linear systems, and should be definable in an intrinsic way. Remarkably, both facts are true, and the intrinsic definition actually explains why the Lie bracket matters in control theory. In fact:

$$[f,g](p) = \lim_{t\to 0} t^{-2}(pe^{tf}e^{tg}e^{-tf}e^{-tg} - p) \quad . \tag{3}$$

The control-theoretic importance of this will be discussed below in Section 3.1.

If M is a C^∞ manifold, we use $\Gamma^\infty(TM)$ to denote the set of all C^∞ vector fields on M. Then $[f,g] \in \Gamma^\infty(TM)$ whenever f, g are in $\Gamma^\infty(TM)$. A *Lie subalgebra* of $\Gamma^\infty(TM)$ is a linear subspace L of $\Gamma^\infty(TM)$ such that $f, g \in L \Rightarrow [f,g] \in L$. If $S \subseteq \Gamma^\infty(TM)$, then there is a smallest Lie subalgebra $\mathrm{Lie}(S)$ of $\Gamma^\infty(TM)$ containing S, called the *Lie algebra generated by S.*

3 Some Examples

We now present some simple examples, to illustrate how DGM's work, and to show the difference between their power and that of classical control theory methods such as the Pontryagin Maximum Principle (PMP). Since lack of space prevents us from dicussing the general theory, we will work with a fairly limited class of problems, namely, with *minimum time problems* for C^∞ control-linear systems of the form

$$\Sigma : \ \dot{x} = f_0(x) + u_1 f_1(x) + \ldots + u_m f_m(x) \ , \ \ x \in M \ , \ \ u = (u_1, \ldots, u_m) \in U \ , \tag{4}$$

where the state space M is a C^∞ manifold, and $f_0, \ldots, f_m$ are C^∞ vector fields on M. The *control space* U is assumed to be a compact, convex subset of $\mathbb{R}^m$ with a nonempty interior. If $f_0 \equiv 0$ and $0 \in \mathrm{Int}\, U$, then we say that Σ is *driftless.* (If M and the f_i are C^ω, then we call Σ a C^ω *system.*) For $a \leq b$, we let $\mathcal{U}(a,b)$ denote the set of all measurable functions $\eta : [a,b] \to U$, and write $\mathcal{U}$ for the set $\cup\{\mathcal{U}(a,b) : -\infty < a \leq b < \infty\}$. A *control* for Σ is a member of $\mathcal{U}$. We write $f(x,u)$ for $f_0(x) + u_1 f_1(x) + \ldots + u_m f_m(x)$. If $\eta \in \mathcal{U}(a,b)$, an *$\eta$-trajectory* of Σ is an absolutely continuous curve $\gamma : [a,b] \to M$ such that $\dot{\gamma}(t) = f(\gamma(t), \eta(t))$ for almost all $t \in [a,b]$. We use $\mathrm{Tr}(\Sigma, \eta)$ to denote the set of all η-trajectories of Σ, and define $\mathrm{Tr}(\Sigma)(a,b) = \cup\{\mathrm{Tr}(\Sigma,\eta) : \eta \in \mathcal{U}(a,b)\}$, $\mathrm{Tr}(\Sigma) = \cup\{\mathrm{Tr}(\Sigma,\eta) : \eta \in \mathcal{U}\}$. If

$\gamma \in \mathrm{Tr}(\Sigma, \eta)$, $\eta \in \mathcal{U}(a, b)$, then (i) a, b, $b-a$, $\gamma(a)$ and $\gamma(b)$ are, respectively, the *initial time*, the *terminal time*, the *duration*, the *initial point* and the *terminal point* of γ, and are denoted by $\tau_-(\gamma)$, $\tau_+(\gamma)$, $T(\gamma)$, $x_-(\gamma)$, $x_+(\gamma)$; (ii) (γ, η) is called an *admissible pair*; (iii) if $p = x_-(\gamma)$, $q = x_+(\gamma)$, $T = T(\gamma)$, then γ *goes from p to q in time T*, and η *steers p to q in time $T(\gamma)$*, and write $\gamma : p \overset{T}{\leadsto} q$ and $\eta : p \overset{T}{\leadsto} q$. We use $\mathrm{Adm}(\Sigma)$ to denote the set of all admissible pairs of Σ. We write $\gamma : p \leadsto q$ if $(\exists T \geq 0)(\gamma : p \overset{T}{\leadsto} q)$, and similarly for $\eta : p \leadsto q$. If there is an $\eta \in \mathcal{U}$ such that $\eta : p \leadsto q$ (resp. $\eta : p \overset{T}{\leadsto} q$), then we say that q is Σ-*reachable* (resp. Σ-*reachable in time T*) from p, and that p is Σ-*controllable* (resp. Σ-*controllable in time T*) to q.

If $S \subseteq M$, we use $\mathcal{R}^\Sigma(S)$ (resp. $\mathcal{R}^\Sigma(\leq T; S)$, $\mathcal{R}^\Sigma(T; S)$, $\mathcal{C}^\Sigma(S)$, $\mathcal{C}^\Sigma(\leq T; S)$, $\mathcal{C}^\Sigma(T; S)$) to denote the set of all points q such that q is Σ-reachable from p (resp. Σ-reachable from p in time $\leq T$, Σ-reachable from p in time T, Σ-controllable to p, Σ-controllable to p in time $\leq T$, Σ-controllable to p in time T) for some $p \in S$. (If $S = \{p\}$, $p \in M$, then we write $\mathcal{R}^\Sigma(p)$, etc. instead of $\mathcal{R}^\Sigma(\{p\})$, etc.) We call Σ *controllable* if $\mathcal{R}^\Sigma(p) = M$ for all $p \in M$ (i.e. if $\mathcal{C}^\Sigma(p) = M$ for all $p \in M$).

Our hypotheses clearly imply the uniqueness and local existence of trajectories. We call Σ *complete* if it has no explosions, i.e. if for every a, b, $\eta \in \mathcal{U}(a, b)$, $\bar{t} \in [a, b]$, $\bar{p} \in M$ there is a $\gamma \in \mathrm{Tr}(\Sigma, \eta)$ such that $\gamma(\bar{t}) = \bar{p}$. (This will happen, e.g., if M is compact, or if $M = \mathbb{R}^n$ and the f_i satisfy a linear growth bound $||f_i(x)|| \leq C(1 + ||x||)$.) If Σ is complete and S is compact, then *all the sets $\mathcal{R}^\Sigma(\leq T; S)$, $\mathcal{R}^\Sigma(T; S)$, $\mathcal{C}^\Sigma(\leq T; S)$, $\mathcal{C}^\Sigma(T; S)$, are compact.*

Reachable sets are related to controllable sets as follows. Define the *reversed* system Σ^- to be the system $\dot{x} = -f(x, u)$, $x \in \Omega$, $u \in U$. Then the controllable sets for Σ are the reachable sets for Σ^-, and viceversa.

A system Σ is *small-time locally controllable (STLC) to* (resp. *from*) a set S if for every $T > 0$ the set S is contained in the interior of $\mathcal{C}^\Sigma(\leq T; S)$ (resp. $\mathcal{R}^\Sigma(\leq T; S)$). Clearly, Σ is STLC to S if and only if Σ^- is STLC from S.

A *boundary trajectory* is a $\gamma \in \mathrm{Tr}(\Sigma)$ such that $x_+(\gamma) \in \partial\mathcal{R}^\Sigma(\leq T, x_-(\gamma))$ for some $T > T(\gamma)$. (Here $\partial X \overset{\mathrm{def}}{=} \mathrm{Clos}\, X - \mathrm{Int}\, X$.) The *value function* $V^\Sigma : M \times M \to \mathbb{R} \cup \{+\infty\}$ is defined by letting $V^\Sigma(p, q) = \inf\{T(\gamma) : \gamma \in \mathrm{Tr}(\Sigma), \gamma : p \leadsto q\}$. (In particular, $V^\Sigma(p, q) = +\infty$ if $q \notin \mathcal{R}^\Sigma(p)$.) If we specify a *target set* S — i.e. a nonempty closed subset S of M— we then define the *value function* V_S^Σ by $V_S^\Sigma(p) = \inf\{V^\Sigma(p, q) : q \in S\}$. A trajectory γ such that $T(\gamma) = V^\Sigma(x_-(\gamma), x_+(\gamma))$ (resp. $x_+(\gamma) \in S$ and $T(\gamma) = V_S^\Sigma(x_-(\gamma)))$ is called a *time-minimizer* (resp. an *S-time-minimizer*) for Σ. We let $\mathcal{M}(\Sigma)$ (resp. $\mathcal{M}_S(\Sigma)$, $\mathcal{B}(\Sigma)$) denote the set of all time-minimizers (resp. S-time-minimizers, boundary trajectories) for Σ. (If $S = \{p\}$, $p \in M$, we write V_p^Σ, $\mathcal{M}_p(\Sigma)$.) Clearly, $\mathcal{M}_S(\Sigma) \subseteq \mathcal{M}(\Sigma)$. Our assumptions then imply that, if Σ is complete, then (i) V^Σ *is lower semicontinuous* (LSC); (ii) *for every $p, q \in M$ such that $V^\Sigma(p, q) < +\infty$ there is a $\gamma \in \mathcal{M}(\Sigma)$ such that $\gamma : p \leadsto q$;* (iii) V_S^Σ is LSC and $(\forall p \in M)((V_S^\Sigma(p) < +\infty) \Rightarrow (\exists \gamma \in \mathcal{M}_S(\Sigma))(x_-(\gamma) = p)$ for every closed S. Moreover, $V_S^\Sigma(p) \geq 0$ for all $p \in M$, and $V_S^\Sigma(p) = 0$ if and only if $p \in S$. Finally, V_S^Σ is continuous if and only if Σ is STLC to S.

The *Pontryagin Maximum Principle* (PMP) gives necessary conditions for a trajectory γ to be in $\mathcal{M}(\Sigma)$ and for γ to be in $\mathcal{B}(\Sigma)$. To state the PMP, we first define the *Hamiltonian lift* Σ^* of Σ. Assume first that M is open in $\mathbb{R}^n$. The *Hamiltonian* H^Σ of Σ is the function $H^\Sigma : M \times \mathbb{R}^n \times U \to \mathbb{R}$ given by $H^\Sigma(x, z, u) = \langle z, f(x, u)\rangle$. Then Σ^* is the control system $\dot{x} = \frac{\partial H^\Sigma}{\partial z}(x, z, u)$, $\dot{z} = -\frac{\partial H^\Sigma}{\partial x}(x, z, u)$, with state space $M \times (\mathbb{R}^n - \{0\})$ and control space U. (The first equation just says that $\dot{x} = f(x, u)$, so every $\xi \in \mathrm{Tr}(\Sigma^*)$ is of the form $t \to \xi(t) = (\gamma(t), \zeta(t))$, where $\gamma \in \mathrm{Tr}(\Sigma)$ and $\zeta(t) \neq 0$.) Now, if M is a general manifold, then Σ^* makes parfect sense, provided that $\zeta(t)$ is interpreted as a *covector* at $\gamma(t)$. The state space of Σ^* is now $T^\# M$, where $T^\# M = \{(x,z) : x \in M\,,\, z \in T_x^* M\,,\, z \neq 0\}$ (so $T^\# M$ is the "cotangent bundle T^*M with the 0 section removed"). An admissible pair (ξ, η) for Σ^* is *H-minimizing* if there is a constant $\zeta_0 \geq 0$ such that $H^\Sigma(\xi(t), \eta(t)) = \min\{H^\Sigma(\xi(t), u) : u \in U\} = -\zeta_0$ for almost all $t \in [\tau_-(\gamma), \tau_+(\gamma)]$. If $\zeta_0 = 0$, then (ξ, η) is *zero-H-minimizing.* A *Pontryagin extremal* (resp. an *abnormal Pontryagin extremal*) is an admissible pair (γ, η) for which there is a ζ such that $(\gamma, \zeta) \in \mathrm{Tr}(\Sigma^*, \eta)$ and (γ, ζ, η) is H-minimizing (resp. zero-H-minimizing). The PMP says that *if $\gamma \in \mathrm{Tr}(\Sigma, \eta)$, then: (i) if $\gamma \in \mathcal{M}(\Sigma)$, then (γ, η) is a Pontryagin extremal, and (ii) if $\gamma \in \mathcal{B}(\Sigma)$, then (γ, η) is an abnormal Pontryagin extremal.*

3.1 Local controllability of the "Reeds-Shepp car" and Lie brackets

Consider the three-dimensional system Σ_{RS} given by $\dot{x}_1 = u\cos\theta$, $\dot{x}_2 = u\sin\theta$, $\dot{\theta} = v$, with state variables x_1, x_2, θ, and controls u, v, subject to $|u| \leq 1$, $|v| \leq 1$. Think of Σ_{RS} as describing the motion in a plane of a vehicle $\mathbf{V}$ that has, at each time, a position (x_1, x_2) and an orientation θ. Then $\mathbf{V}$ can move forwards and backwards but not sideways, and its orientation can be controlled as well. The state space M is not really $\mathbb{R}^3$, but $\mathbb{R}^2 \times S^1$, since θ is an angle. (This will be important later.) Our objective is to understand how to drive $\mathbf{V}$ from any given initial configuration (x_1^-, x_2^-, θ^-) to any other terminal configuration (x_1^+, x_2^+, θ^+) in minimum time. (This problem was studied by Reeds and Shepp in [9] using *ad hoc* techniques, and by Sussmann and Tang in [11] with control theory methods.)

A first natural question is to investigate the *controllability* and *local controllability* of Σ_{RS}. It is easy to verify directly, using combinations of small forwards and backwards motions and small rotations, that Σ_{RS} is STLC to and from any point. Let us see whether we can recover this using classical optimal control techniques, e.g. the PMP. Clearly, a driftless system Σ is not STLC from p if and only if, for arbitrary $T > 0$, the *constant trajectory* $t \to \gamma(t) = p$, $0 \leq t \leq T$ is in $\mathcal{B}(\Sigma)$. Now, the PMP gives a necessary condition for $\gamma \in \mathcal{B}(\Sigma)$, and this yields a sufficient condition for STLC from p. In our case, this simply amounts to the condition that the linearization of our equations around $\bar{p}$ be controllable, and it is easy to see that this condition fails. So *the sufficient condition arising from the PMP does not cover this case.*

Now contrast this with the differential geometric approach. Notice first that by using constant controls $(u, v) = (\pm 1, 0)$, $(u, v) = (0, \pm 1)$ we can move freely along the directions of the vector fields f, g with components $(\cos\theta, \sin\theta, 0)$ and $(0, 0, 1)$. So we have two "directly controllable" directions. To have STLC we must be able to move our vehicle forwards and backwards in three independent directions, so we are missing just one direction. At this point, a differential geometer would remind us that there is a binary operation on vector fields called the *Lie bracket*, which does precisely what we need, namely, associate to any two vector fields f and g a new vector field $h = [f, g]$ such that, even if one cannot directly steer the system along trajectories of h, one can approximately achieve the same effect by combining f- and g-motions. Identity (3) tells us how to produce a curve $t \to \delta(t)$ whose direction to second order is $[f, g](p)$ (i.e. $\delta(t) = t^2[f, g](p) + o(t^2)$), all whose points can be reached from p by combining forwards and backwards f- and g-motions, even though δ itself need not be a permissible motion. So, if $[f, g](p)$, $f(p)$, $g(p)$ are linearly independent, then it is as though we really had three linearly independent directions of motion, and STLC from p should follow, if we are in three dimensions. (For a rigorous proof, apply the Inverse Function Theorem to the map $(t, s, \tau) \to$

$\sigma(t)e^{sf}e^{\tau g}$, where σ is the C^1 curve given by $\sigma(t) = pe^{\sqrt{t}f}e^{\sqrt{t}g}e^{-\sqrt{t}f}e^{-\sqrt{t}g}$ for $t \geq 0$ and $\sigma(t) = pe^{\sqrt{-t}g}e^{\sqrt{-t}f}e^{-\sqrt{-t}g}e^{-\sqrt{-t}f}$ for $t \leq 0$.)

For Σ_{RS}, h is easily seen to have components $(\sin\theta, -\cos\theta, 0)$. So h is precisely the "missing" third direction of sideways motion. Then $f(p)$, $g(p)$, $h(p)$ are linearly independent for each p. So STLC holds. So *by using Lie brackets we have been able to derive a result that was beyond the power of the Mamimum Principle.*

3.2 Lie Algebras and Chow's Theorem

The local controllability proof given above can in fact be carried out under much more general situations. The key idea is to iterate the procedure and compute all possible higher-order brackets such as $[f_i, [f_j, f_k]]$ or $[[f_i, f_j], [f_k, [f_\ell, f_m]]]$. Then STLC from and to p should follow, provided that the directions at p of all these brackets linearly span the space of all possible directions. This turns out to be a true, rigorous theorem, known as "Chow's Theorem," which in additon is remarkably easy to prove. (Cf. Krener [5]. The systematic use of Chow's Theorem in control theory was first advocated by R. Hermann in [3]). We associate to a system Σ its *accessibility Lie algebra* L^Σ, defined by $L^\Sigma = \mathrm{Lie}(\{f_0, f_1, \ldots, f_m\})$. We say that Σ has the *Lie algebra rank property* (LARP) a $p \in M$ if $L^\Sigma(p) = T_pM$, where $L^\Sigma(p) = \{X(p) : X \in L^\Sigma\}$. Chow's Theorem then says that, *if Σ has the LARP at p and is driftless, then it is STLC to and from p.*

3.3 Lie Brackets, Normal Forms and the Lobry Picture

The fact that the Lie bracket is intrinsic is often used as follows: suppose we know that certain numbers E_j, defined in terms of brackets, are nonzero at a point p. Using this, it may be possible to find a chart κ with respect to which our problem is in a simpler "canonical form," and the E_j will still be $\neq 0$ in the κ-representation. We illustrate this method by outlining the rigorous theory of the "Lobry picture," which appeared in C. Lobry's path-breaking 1970 paper [6].[4]

We study the *local structure of the reachable set* from a point p, for a three-dimensional control system of the form $\dot{x} = f(x) + ug(x)$, $-1 \leq u \leq 1$, where f, g are C^2 vector fields on a open subset Ω of $\mathbb{R}^3$. Let $X = f - g$, $Y = f + g$. Let $S_{XY}(p)$, $S_{YX}(p)$ be, respectively, the XY- and YX-surfaces from p, i.e. $S_{XY}(p) = \{pe^{tX}e^{sY} : t \geq 0\,,\, s \geq 0\}$, $S_{YX}(p) = \{pe^{sY}e^{tX} : t \geq 0\,,\, s \geq 0\}$, so that each surface looks like a —smoothly deformed— two-dimensional closed quarter-plane, and the two surfaces are glued at their boundary, which is the union of the two half-curves $L_X(p) = \{pe^{tX} : t \geq 0\}$, $L_Y(p) = \{pe^{tY} : t \geq 0\}$. Our goal is to outline a rigorous formulation and proof of the assertion that, *if $f(p)$, $g(p)$, $[f,g](p)$ are linearly independent. then, locally, the reachable set from p is the set that lies between $S_{XY}(p)$ and $S_{YX}(p)$.* It turns out that, using DGM's, we can provide both a correct mathematical formulation and a rigorous proof.

Theorem 1. *Assume that Ω is an open subset of $\mathbb{R}^3$, $k \geq 2$, f and g are C^k vector fields on Ω, Σ is the system $\dot{x} = f(x) + ug(x)$, $x \in \Omega$, $|u| \leq 1$, and $p \in \Omega$ is such that $f(p)$, $g(p)$, $[f,g](p)$ are linearly independent. Let $X = f - g$, $Y = f + g$. Then there exist a $T > 0$, a $\rho > 0$, a cubic C^k chart κ of radius ρ centered at p, and a C^k function $\varphi : (-\rho, \rho)^2 \to (-\rho, \rho)$, such that $\varphi(x_1, x_2) = x_1x_2(1 + \psi(x_1, x_2))$ —where $\psi \in C^{k-2}$, $\psi(0,0) = 0$, and $|\psi(x_1, x_2)| < 1$ for $|x_1| < \rho$, $|x_2| < \rho$— with the property that, if $\mathcal{R} = \mathcal{R}^\Sigma(\leq T; p)$, then*

[4] I am indebted to H. Schättler for suggesting the approach presented here.

$$\kappa(\mathcal{R}\cap D(\kappa)) = \{(x_1,x_2,x_3) : x_1 \geq 0\,,\ x_2 \geq 0\,,\ 0 \leq x_3 \leq \varphi(x_1,x_2)\}\ . \qquad (5)$$

Moreover, the sets $S^- = \kappa^{-1}(\{(x_1,x_2,x_3) : x_1 \geq 0\,,\ x_2 \geq 0\,,\ x_3 = 0\})$*,* $S^+ = \kappa^{-1}(\{(x_1,x_2,x_3) : x_1 \geq 0\,,\ x_2 \geq 0\,,\ x_3 = \varphi(x_1,x_2)\})$ *are precisely the intersections with* $D(\kappa)$ *of the sets* $S^T_{XY}(p)$*,* $S^T_{YX}(p)$*, defined by* $S^T_{XY}(p) = \{pe^{tX}e^{sY} : 0 \leq t \leq T\,,\ 0 \leq s \leq T\}$*,* $S_{YX}(p) = \{pe^{sY}e^{tX} : 0 \leq t \leq T\,,\ 0 \leq s \leq T\}$*.*

Remark. A *cubic chart of radius* ρ *centered at* p is a chart κ such that $\kappa(p) = 0$ and $Im(\kappa)$ is the cube $(-\rho,\rho)^n$.

OUTLINE OF THE PROOF. We let $Z = [X,Y]$ (so Z is a C^{k-1} vector field) and define an Ω-valued map F by $F(x_1,x_2,x_3) = pe^{x_3Z}e^{x_2Y}e^{x_1X}$. Then F is well defined and of class C^k near $(0,0,0)$ (because the curve $t \to pe^{tZ}$ is C^k, since Z is C^{k-1}) and its Jacobian at $(0,0,0)$ is nonsingular. Therefore, there is a $\bar\rho > 0$ such that, if $0 < \rho \leq \bar\rho$, then F is 1-1 on the cube $C(\rho) = (-\rho,\rho)^3$, the image $N(\rho) = F(C(\rho))$ is open, and the restriction $F\lceil C(\rho)$ of F to $C(\rho)$ has a C^k inverse $\kappa_\rho : N(\rho) \to C(\rho)$. Then κ_ρ is a cubic C^k chart centered at p, $D(\kappa_\rho) = N(\rho)$, $X^{\kappa_\rho} \equiv (1,0,0)$, and $Y^{\kappa_\rho} \equiv (\alpha_1, 1+\alpha_2, \alpha_3)$, where the α_i are C^{k-1} functions on $C(\rho)$ that vanish when $x_1 = 0$. Then $\alpha_i = x_1\tilde\alpha_i$, where the $\tilde\alpha_i$ are C^{k-2}. Also, $Z^{\kappa_\rho} \equiv (\beta_1,\beta_2,\beta_3)$, where the β_i are of class C^{k-1}, and $\beta_1 \equiv \beta_2 \equiv \beta_3 - 1 \equiv 0$ wherever $x_1 = x_2 = 0$, so that $\beta_i = x_1\hat\beta_i + x_2\check\beta_i$ for $i = 1,2$, and $\beta_3 = 1 + x_1\hat\beta_3 + x_2\check\beta_3$, where the $\hat\beta_i$, $\check\beta_i$ are C^{k-2} functions. Then $Z = [X,Y]$ implies $\beta_i = \frac{\partial\alpha_i}{\partial x_1}$. So $\frac{\partial\alpha_i}{\partial x_1} \in C^{k-1}$, so the $\tilde\alpha_i$ are C^{k-1}, since $\tilde\alpha_i(x_1,x_2,x_3) = \int_0^1 \frac{\partial\alpha_i}{\partial x_1}(sx_1,x_2,x_3)ds$. But then $\tilde\alpha_i = \beta_i - x_1\frac{\partial\tilde\alpha_i}{\partial x_1}$, and therefore $\tilde\alpha_i = x_1\hat\alpha_i + x_2\check\alpha_i$ for $i = 1,2$, $\tilde\alpha_3 = 1 + x_1\hat\alpha_3 + x_2\check\alpha_3$, where $\hat\alpha_i = \hat\beta_i - \frac{\partial\tilde\alpha_i}{\partial x_1}$, $\check\alpha_i = \check\beta_i$, so that the $\hat\alpha_i$ and $\check\alpha_i$ are C^{k-2}. This shows that $Y^{\kappa_\rho} = Y_0 + x_1^2\hat Y + x_1x_2\check Y$, where $Y_0 \equiv (0,1,x_1)$, $\hat Y \equiv (\hat\alpha_1,\hat\alpha_2,\hat\alpha_3)$, $\check Y \equiv (\check\alpha_1,\check\alpha_2,\check\alpha_3)$. So, by changing coordinates via $\kappa_{\bar\rho}$, we have put (X,Y) in the "normal form" $X = (1,0,0), Y = (0,1,x_1) + o(||x||^2)$. From now on, we identify $N(\bar\rho)$ with $C(\bar\rho)$ via κ_ρ.

Now, the YX-surface is easy to compute, since the Y-trajectory from p is just the positive part of the x_2 axis, and the X-trajectories are just lines $\dot x_1 = 1$, $\dot x_2 = \dot x_3 = 0$. To compute the XY-surface, we let $\Phi(t,s) = pe^{tX}e^{sY}$, and observe that Φ is well defined and C^k for (t,s) near $(0,0)$. Moreover, the derivatives $\frac{\partial\Phi}{\partial t}(0,0)$ and $\frac{\partial\Phi}{\partial s}(0,0)$ are equal, respectively, to $X(p)$ and $Y(p)$, so they are linearly independent. This means, by the Implicit Function Theorem, that there is an open square $U = (-T,T)^2$ in $\mathbb{R}^2$ such that Φ is defined on U, $\Phi(U)$ is a two-dimensional surface of class C^k, $\Phi(U) \subseteq N(\bar\rho)$, the restriction $\Phi\lceil U$ of Φ to U is one-to one, and its inverse $(\Phi\lceil U)^{-1} : \Phi(U) \to U$ is C^k. From now on, we restrict the domain of Φ to U, so Φ is a mapping into $N(\bar\rho)$, i.e. into $C(\bar\rho)$. Moreover, if we let $\pi(x_1,x_2,x_3) = (x_1,x_2)$, then $\Psi = \pi\circ\Phi$ maps $(0,0)$ to $(0,0)$ and its Jacobian matrix at $(0,0)$ is the identity matrix. Therefore, given any sufficiently small $\rho > 0$, there is an open neighborhood $V_\rho \subseteq U$ of $(0,0)$ in $\mathbb{R}^2$ which is mapped by Ψ in a one-to-one and onto fashion, and with a C^k inverse, onto the square $(-\rho,\rho)^2$. Then $\Phi(V_\rho)$ is precisely the graph of a C^k function $\varphi_\rho : (-\rho,\rho)^2 \to \mathbb{R}$. The function φ_ρ vanishes when $x_1 = 0$ or $x_2 = 0$ (since $\Phi(t,0) = (t,0,0)$ and $\Phi(0,s) = (0,s,0)$). So $\varphi_\rho(x_1,x_2) = x_1x_2\tilde\varphi_\rho(x_1,x_2)$, where $\tilde\varphi_\rho \in C^{k-2}$. In particular, by picking ρ sufficiently small, we can guarantee that φ_ρ takes values in $(-\rho,\rho)$, so $\Phi(V_\rho) \subseteq C(\rho)$. Since $\alpha_i = x_1(1 + x_1\hat\alpha_i + x_2\check\alpha_i)$ for $i = 1,2$, and $\alpha_3 = x_1(1 + x_1\hat\alpha_3 + x_2\check\alpha_3)$, we can make sure, by making ρ even smaller if needed, that $\alpha_3 > 0$ on C_ρ whenever $x_1 > 0$, and also that $\alpha_2 > 1/2$ and $|\alpha_1| < 1/2$ on C_ρ. Then, with such a choice of ρ, we have $S^T_{XY}(p) \cap N(\rho) = S^+(\rho)$, and $S^T_{YX}(p) \cap n(\rho) = S^-(\rho)$, where $S^+(\rho)$, $S^-(\rho)$ are, respectively, the intersections of $\Phi(V_\rho)$ and the square $(-\rho,\rho)^2 \times \{0\}$ with the set $\{(x_1,x_2,x_3) : x_1 \geq 0\,,\, x_2 \geq 0\}$.

From now on we just write $\kappa = \kappa_\rho$, $\varphi = \varphi_\rho$, $\tilde{\varphi} = \tilde{\varphi}_\rho$, $S^+ = S^+(\rho)$, $S^- = S^-(\rho)$.

Let us now show that, possibly after shrinking ρ, we can assume that $\frac{\partial\varphi}{\partial x_1} > 0$ when $0 < x_1 < \rho$, $0 < x_2 < \rho$. Along a Y-trajectory $t \to x(t)$ we have $\dot{x}_1 = \alpha_1 = x_1^2\hat{\alpha}_1 + x_1x_2\check{\alpha}_1$, $\dot{x}_2 = 1+\alpha_2 = 1+x_1^2\hat{\alpha}_2+x_1x_2\check{\alpha}_2$, and $\dot{x}_3 = \alpha_3 = x_1(1+x_1\hat{\alpha}_3+x_2\check{\alpha}_3)$. On the other hand, $x_3 = \varphi(x_1,x_2)$, so $\dot{x}_3 = \frac{\partial\varphi}{\partial x_1}\dot{x}_1 + \frac{\partial\varphi}{\partial x_2}\dot{x}_2$. Plugging in the values of the $\dot{x}_i$, we see that $\frac{\partial\varphi}{\partial x_2} = x_1(1+o(1))$. Hence $\frac{\partial^2\varphi}{\partial x_1 x_2}(0,0) = 1$. Therefore $\tilde{\varphi}(0,0) = 1$. Now, $\frac{\partial\varphi}{\partial x_1} = x_2(\tilde{\varphi} + x_1\frac{\partial\tilde{\varphi}}{\partial x_1})$. By shrinking ρ, we can assume that $\tilde{\varphi} + x_1\frac{\partial\tilde{\varphi}}{\partial x_1} > 0$ on $(-\rho,\rho)^2$, and then $\frac{\partial\varphi}{\partial x_1} > 0$ when $x_2 \geq 0$, proving our claim. Also, if we define $\psi = \tilde{\varphi} - 1$, then we have shown that $\psi(0,0) = 0$, and then we can assume that $|\psi| < 1/2$ throughout $(-\rho,\rho)^2$ by shrinking ρ.

If we now define $\mathbf{R}$ to be the right-hand side of (5), then all we need is to show that $\mathcal{R}^* = \mathbf{R}$, where $\mathcal{R}^* = \mathcal{R} \cap N(\rho)$.

The set $\mathbf{R}$ is clearly closed in $C(\rho)$, and has the property that, if $q \in \mathbf{R}$, then qe^{tX} and qe^{tY} are in $\mathbf{R}$ for $t \leq 0$, t sufficiently small. (The conclusion about qe^{tX} is trivial unless $q \in S^+$. But in that case the function $\zeta = \varphi - x_3$ vanishes at q and is increasing along the X-trajectory from q, and then $\zeta(qe^{tX}) \geq 0$ for $t > 0$, t small. But this means that qe^{tX} lies below the graph of φ, as desired. The conclusion for qe^{tY} is trivial unless $q \in S^-$. But in that case it follows from the fact that α_3 is ≥ 0 for $x_1 \geq 0$, so $Y(q)$ points up if $q \in S^-$.) From this it follows easily that every trajectory γ that starts at p, stays in $N(\rho)$, and corresponds to a bang-bang control (i.e. a piecewise constant control with values ± 1), is entirely contained in $\mathbf{R}$. Then the same conclusion is true for an arbitrary measurable control $u(\cdot)$ with values between -1 and 1, since every trajectory is a limit of trajectories for bang-bang controls. On the other hand, there is a time $\tilde{T} > 0$ such that every trajectory that starts at p and is defined on an interval of length $\leq \tilde{T}$ stays in $N(\rho)$. Now replace T by $\min(T,\tilde{T})$, and shrink ρ further so that $5\rho < T$. Then $\mathcal{R} \subseteq \mathbf{R}$. Moreover, if $x = (x_1,x_2,x_3) \in \mathbf{R}$ then $x_1 \geq 0$, $x_2 \geq 0$, and $\varphi(x_1,x_2) - x_3 \geq 0$, whereas $\varphi(0,x_2) - x_3 \leq 0$, since $\varphi(0,x_2) = 0$ and $x_3 \geq 0$. Therefore there is a t such that $0 \leq t \leq x_1 < \rho$ for which $\varphi(x_1 - t, x_2) = x_3$. This means that $\tilde{x} = (x_1 - t, x_2, x_3)$ is in S^+. Therefore $\tilde{x} = p^\kappa e^{\tau X^\kappa} e^{\sigma Y^\kappa}$, for some $(\tau,\sigma) \in V_\rho$. The bounds $|\alpha_1| < 1/2$, $\alpha_2 > 1/2$ imply that $\sigma < 2\rho$ and $\tau < 2\rho$. So x is reachable from p in time $< 5\rho$. Since $5\rho < T$, we see that $x \in \mathcal{R}$. Hence $\mathbf{R} \subseteq \mathcal{R}$, and our proof is complete. $\square$

3.4 Geometric Selections and Optimal Control of the "Reeds-Shepp Car"

Let us return to the minimum-time control problem for the system Σ_{RS} of 3.1. We know that Σ_{RS} is STLC from every point, and this easily implies that, given any pair (p,q) of points, there is a $\gamma \in \mathcal{M}(\Sigma_{RS})$ that goes from p to q. We would now like to understand the structure of the time-minimizers.

We would also like to find nice "sufficient classes" $\mathcal{S}$. Call a class $\mathcal{S}$ of trajectories *sufficient for time-optimal control* (STOC) for a system Σ if for every pair p, q such that $q \in \mathcal{R}^\Sigma(p)$ there is a $\gamma \in \mathcal{S} \cap \mathcal{M}(\Sigma)$ such that $\gamma : p \rightsquigarrow q$. (Notice that we do not demand that $\mathcal{S} \subseteq \mathcal{M}(\Sigma)$ or that $\mathcal{M}(\Sigma) \subseteq \mathcal{S}$.) The reason for studying such families is that, if $\mathcal{S}$ is STOC, then the problem of finding a time-minimizer from p to q can be reduced to the problem of finding a $\bar{\gamma} \in \mathcal{S}$ that minimizes time among all $\gamma \in \mathcal{S}$ that go from p to q. If $\mathcal{S}$ is very small, then this is likely to be a much more tractable problem. For a general complete Σ, two obvious examples of sufficient families are: (i) $\mathrm{Tr}(\Sigma)$ and (ii) the class $\mathcal{P}(\Sigma)$ of all γ such that (γ, η) is a Pontryagin extremal for some η. Then $\mathcal{M}(\Sigma) \subseteq \mathcal{P}(\Sigma)$, but in general $\mathcal{M}(\Sigma) \neq \mathcal{P}(\Sigma)$, since the PMP is necessary but not sufficient for optimality. This suggests that to find nice sufficient classes one might start by applying the PMP, which enables us to replace $\mathrm{Tr}(\Sigma)$ by $\mathcal{P}(\Sigma)$, and then try to eliminate those elements of $\mathcal{P}(\Sigma)$ that are not optimal, perhaps using more sophisticated necessary conditions for optimality. There is, however, an intrinsic limitation to this strategy, namely, *the possibility that optimality by itself might not suffice as a criterion for selecting a small class.* It may happen that $\mathcal{M}(\Sigma)$ *is too large,* and to get a truly interesting sufficient class $\mathcal{S}$ one has to make a further selection, and not insist on the inclusion $\mathcal{M}(\Sigma) \subseteq \mathcal{S}$.

The Reeds-Shepp example provides a striking illustration of this phenomenon. Consider the "left turn" system Σ_{LT} given by $\dot{x}_1 = u\cos\theta$, $\dot{x}_2 = u\sin\theta$, $\dot{\theta} = 1$, $|u| \leq 1$. Clearly, $\mathrm{Tr}(\Sigma_{LT}) \subseteq \mathrm{Tr}(\Sigma_{RS})$. Moreover, if $\gamma \in \mathrm{Tr}(\Sigma_{LT})$ and $T(\gamma) \leq \pi$, then $\gamma \in \mathcal{M}(\Sigma_{RS})$. So we see that $\mathcal{M}(\Sigma_{RS})$ is too large, because it contains —if $T \leq \pi$— all the η-trajectories for all $\eta \in \mathcal{U}(0,T)$ that are of the form $\eta(t) = (u(t), 1)$, *for an arbitrary measurable function* $u(\cdot) : [0,T] \to [-1,1]$. However, since the state space of Σ_{RS} is 3-dimensional, it would be nice to find a family $\mathcal{U}_0$ of controls depending on 3 continuous parameters, such that the class $\mathcal{S} = \cup\{\mathrm{Tr}(\Sigma_{RS}, \eta) : \eta \in \mathcal{U}_0\}$ is STOC. This requires that we make a selection within $\mathcal{M}(\Sigma_{RS})$.

The following *geometric selection method* turns out to work. Call a system Σ *degenerate* if whenever $\gamma_1, \gamma_2 \in \mathrm{Tr}(\Sigma)$ are such that $x_-(\gamma_1) = x_-(\gamma_2)$ and $x_+(\gamma_1) = x_+(\gamma_2)$ then $T(\gamma_1) = T(\gamma_2)$. Then, if Σ is degenerate and complete, one can easily show that all the sets $\mathcal{R}^\Sigma(p)$ are closed. Now fix some value $\hat{u} \in U$ of u. Assume that $\eta \in \mathcal{U}(0,T)$, $\gamma \in \mathrm{Tr}(\Sigma, \eta)$, $\gamma : p \rightsquigarrow q$, and find $\hat{\gamma} \in \mathrm{Tr}(\Sigma, \hat{\eta})$ such that $\hat{\gamma}(T) = q$, where $\hat{\eta} : [0,T] \to U$, $\hat{\eta}(t) \equiv \hat{u}$. Let $\tau = \inf\{t \in [0,T] : \hat{\gamma}(t) \in \mathcal{R}^\Sigma(p)\}$. Then $\hat{\gamma}(\tau) \in \mathcal{R}^\Sigma(p)$, because $\mathcal{R}^\Sigma(p)$ is closed. So there exist $\tilde{\tau} \geq 0$, $\tilde{\eta} \in \mathcal{U}(0, \tilde{\tau})$, $\tilde{\gamma} \in \mathrm{Tr}(\Sigma, \tilde{\eta})$, such that $\tilde{\gamma} : p \rightsquigarrow \hat{\gamma}(\tau)$. Then $q \in \mathcal{R}^\Sigma(\tilde{\tau} + T - \tau; p) \cap \mathcal{R}^\Sigma(T; p)$. So $\tilde{\tau} + T - \tau = T$, i.e. $\tilde{\tau} = \tau$. It is easy to see that $\tilde{\gamma} \in \mathcal{B}(\Sigma)$. So $\tilde{\gamma}$ *is an abnormal extremal.* Therefore, if we let $\mathbf{T}(\Sigma, \hat{u})$ denote the class of all trajectories that are concatenations of an abnormal extremal and a trajectory of for the constant control $\equiv \hat{u}$, we see that $\mathbf{T}(\Sigma, \hat{u})$ is sufficient for time-optimal control for Σ.

The above cannot be applied to Σ_{LT} because Σ_{LT} is not degenerate (since one can go from any p back to p in time 2π). However, if $\check{\Sigma}_{LT}$ is the *lifted system* —with state space $\mathbb{R}^3$— obtained by regarding θ as a true real number rather than as a real number modulo 2π, then $\check{\Sigma}_{LT}$ is degenerate. Suppose $\gamma \in \mathrm{Tr}(\Sigma_{LT}) \cap \mathcal{M}(\Sigma_{RS})$, $\tau_-(\gamma) = 0$, $\tau_+(\gamma) = T$. Then we can write $\gamma = \pi \circ \check{\gamma}$, where $\pi : \mathbb{R}^3 \to \mathbb{R}^2 \times S^1$ is the projection, and $\check{\gamma} \in \mathrm{Tr}(\check{\Sigma}_{LT})$. Then we can find $\check{\delta} \in \mathbf{T}(\Sigma_{LT}, 1)$ such that $\check{\delta}(0) = \check{\gamma}(0)$ and $\check{\delta}(T) = \check{\gamma}(T)$. If $\delta = \pi \circ \check{\delta}$, then $\delta : p \overset{T}{\rightsquigarrow} q$, so $\delta \in \mathcal{M}(\Sigma_{RS})$.

When we apply the PMP to Σ_{RS} to study $\mathcal{M}(\Sigma_{RS})$, we find out (cf. [11]) that every $\gamma \in \mathcal{M}(\Sigma_{RS})$ is of one of two kinds, depending on whether a certain nonnegative constant κ does or does not vanish. When $\kappa \neq 0$, the PMP conditions, plus some extra geometric arguments, imply that γ is —or can be replaced by— a concatenation of at most five pieces, each of which is either a circle or a segment. However, when $\kappa = 0$, the PMP implies that $v \equiv 1$ or $v \equiv -1$, but yields no more information. On the other hand, if $v \equiv 1$ —the case $v \equiv -1$ is similar— then γ is actually in $\mathrm{Tr}(\Sigma_{LT})$, and so the geometric selection argument applies, and yields another $\delta \in \mathcal{M}(\Sigma_{RS})$, with the same initial and terminal points as γ, which is now a concatenation of an abnormal extremal δ_1 of Σ_{LT} and a $u = 1$ trajectory δ_2. The fact that δ_1 is an abnormal extremal gives very strong information, and implies that δ_1 is a concatenation of at most two circles. Since δ_2 is a circle, we see that, when $\kappa = 0$, γ can be replaced by a concatenation of three circles.

Combining all these considerations, one ends up proving that a certain family $\mathcal{F}$, consisting of concatenations of at most five C's (circles) and S's (segments), is sufficient. The elements of $\mathcal{F}$ actually satisfy further restrictions, e.g. CCCCC is not permitted, more than one S is not allowed, for a CCCC the times along the two inner C's must be equal, for a CCCSC the times along the second and third C must equal $\frac{\pi}{2}$, etc. (All this is done in detail in [11].) In particular, the elements of $\mathcal{F}$ that start at a fixed point p are parametrized by three continuous parameters plus a discrete parameter specifying the combinatorial type of the trajectory. (The class given in [11] consists of 46 three-parameter families.)

3.5 Subanalytic sets and Regularity of the Value Function

Now assume that M is a C^ω manifold. Then there are at least two important classes of subsets M that make intrinsic sense, and have nice "piecewise analyticity" properties, namely, the sets $SM(M)$, $SB(M)$ of *semianalytic* and *subanalytic* subsets of M. (By definition, a subset S of M is in $SM(N)$ if every $p \in M$ has a neighborhood U such that $S \cap U$ is a finite union of sets, each of which is defined by finitely many conditions $f(x) = 0$ or $f(x) > 0$, involving C^ω functions $f : U \to \mathbb{R}$. Then $S \in SB(M)$ if $S = f(\tilde{S})$ for some C^ω manifold N and some C^ω map $f : N \to M$ which is proper on the closure of $\tilde{S}$. In other words: semianalytic sets are locally defined by finite collections of analytic equalities and inequalities, and subanalytic sets are images of semianalytic sets under proper analytic maps. [5]

Following our general geometric philosophy, these classes of sets ought to be useful in control theory. To show that this is indeed so, observe that, for a complete system, the reachable sets are *images* of certain maps. (Indeed, let $\mathcal{U}_- = \cup\{\mathcal{U}(-T,0) : T \geq 0\}$, $\mathcal{U}_-(T) = \cup\{\mathcal{U}(-t,0) : 0 \leq t \leq T\}$. Define a map $\Phi_S^\Sigma : \mathcal{U}_- \times S \to M$ by $\Phi_S^\Sigma(\eta, p) = x_-(\gamma)$, where $\gamma \in \mathrm{Tr}(\Sigma, \eta)$ is such that $x_+(\gamma) = p$. Then $\mathcal{C}^\Sigma(S) = \Phi_S^\Sigma(\mathcal{U} \times S)$, $\mathcal{C}^\Sigma(T; S) = \Phi_S^\Sigma(\mathcal{U}(-T,0) \times S)$, and $\mathcal{C}^\Sigma(\leq T; S) = \Phi_S^\Sigma(\mathcal{U}_-(T) \times S)$. Similar identities —with Σ^- instead of Σ— hold for the reachable sets.) In addition, *the graph of the value function is a piece of the boundary of a controllable set.* Precisely, let Σ^a be the system with state space $M \times \mathbb{R}$, control space U, and dynamical equations $\dot{x} = f(x,u)$, $\dot{y} = -1$. (Think of y as "time to go.") Then (q,t) *is Σ^a-controllable to $(p,0)$ if and only if q is controllable to p in time t.* Define the *lower boundary $\partial_- A$* of a subset A of $M \times \mathbb{R}$ to be the set $\partial_- A = \{(q,t) \in A : (\forall s < t)(q,s) \notin A\}$. Then, if we let $S^a = S \times \{0\}$, and Σ is complete, *the set $\mathcal{C}^{\Sigma^a}(S^a)$ is closed in $M \times \mathbb{R}$, and its lower boundary $\partial_- \mathcal{C}^{\Sigma^a}(S^a)$ is precisely the graph $G(V_S^\Sigma)$ of the value function V_S^Σ.*

On the other hand, subanalytic sets are also, by definition, images. And one can prove that, if $S \in SB(M)$ and S is closed, then ∂S is subanalytic. Moreover, if M is actually a product $N \times \mathbb{R}$, and $S \subseteq N \times [0, \infty)$, then $\partial_- S$ is subanalytic in M. Now let us quote a property of subanalytic sets.

Theorem 2. *Let M be a C^ω manifold, and let S be a closed subanalytic subset of M. Let $\rho : S \to \mathbb{R}$ be a function whose graph $G(\rho)$ is a subanalytic subset of $M \times \mathbb{R}$. Then ρ is piecewise analytic, in the sense that there exists a ρ-locally finite partition $\mathcal{P}$ of S into connected, subanalytic embedded C^ω submanifolds of M such that the restriction $\rho\lceil P$ of ρ to each member P of $\mathcal{P}$ is of class C^ω.*

---

[5] If A and B are topological spaces, then a map f from A to B is *proper* if the inverse image of every compact subset of B is compact in A.

(A set $\mathcal{X}$ of subsets of M is *ρ-locally finite* if every compact subset of M on which ρ is bounded intersects finitely many members of $\mathcal{X}$.)

So controllable sets are are analogous to subanalytic sets in the sense that both are images. The lower boundary of a subanalytic set is subanalytic, and the lower boundary of a certain controllable set is the graph of the value function. And if the graph of a function is subanalytic then the function itself is piecewise analytic. One may hope to be able to combine all these facts, and prove that V_S^Σ is indeed piecewise analytic for large classes of problems. However, subanalytic sets are images of semianalytic subsets of finite-dimensional manifolds under analytic maps, whereas $\mathcal{C}^\Sigma(S) = \Phi_S^\Sigma(\mathcal{U} \times S)$, and $\mathcal{U}$ is an infinite-dimensional function space. *Suppose, however, that we could replace $\mathcal{U}$ by a finite-dimensional object.* Then we would be much closer to proving subanalyticity of the graph of V_S^Σ. What we need is an "analytic finite-dimensional reduction" (AFDR), i.e. a collection $\mathcal{G} = \{\gamma_\alpha : \alpha \in A\}$ of trajectories, parametrized by points α in a closed subanalytic subset A of a finite-dimensional C^ω manifold N, and satisfying $\tau_+(\gamma_\alpha) = 0$, such that (i) Γ is STOC for Σ, (ii) the map $\alpha \to (x_-(\gamma_\alpha), x_+(\gamma_\alpha), \tau_-(\gamma_\alpha))$ is C^ω. We call Γ *proper* if the map $\mu_\Gamma : \alpha \to (x_-(\gamma_\alpha), \tau_-(\gamma_\alpha))$ is proper. If a proper AFDR Γ exists, and S is subanalytic, then it is a rigorous theorem that $G(V_S^\Sigma)$ is subanalytic, and then V_S^Σ is piecewise analytic. (Proof: let $A_S = \{\alpha : x_+(\gamma_\alpha) \in S\}$. Then $A_S \in SB(N)$. Now let $\mathcal{C} = \mu_\Gamma(A_S)$. Then $\mathcal{C}$ is subanalytic, because μ_Γ is C^ω and proper on A. And one can easily verify that $G(V_S^\Sigma) = \partial_-\mathcal{C}$. So $G(V_S^\Sigma) \in SB(M \times \mathrm{I\!R})$. Q.E.D.)

The above discussion shows that one can prove piecewise analyticity of the value function if one can find a proper AFDR Γ. It is now natural to ask whether the existence of such an AFDR is going to be an exceptional situation, or whether it is going to be true for large classes of C^ω systems. Examples have been given of systems where the value function is not subanalytic (cf. [7]), so we know that an AFDR need not always exist. On the other hand, it has been proved to exist for large classes of problems. We have already described one example in our analysis of the "Reeds-Shepp car." (The parameter manifold N is $\mathrm{I\!R}^2 \times S^1 \times \hat{N}$, where $\hat{N}$ has 46 connected components $\hat{N}_j$, each one being a copy of $\mathrm{I\!R}^3$. Each $\hat{N}_j$ corresponds to one of the combinatorial types of trajectories in the family $\mathcal{F}$. If we write $N_j = \mathrm{I\!R}^2 \times S^1 \times \hat{N}_j$, then $A \cap N_j$ is the product of $\mathrm{I\!R}^2 \times S^1$ times the nonnegative orthant $\mathrm{I\!R}^3_+$ of $\mathrm{I\!R}^3$. If $\alpha = (x_1, x_2, \theta, t) \in A_j$, so $t = (t_1, t_2, t_3) \in \mathrm{I\!R}^3_+$, then γ_α is the trajectory of the corresponding type whose time parameters are t_1, t_2, t_2, with domain $[-T, 0]$, where $T = t_1 + t_2 + t_3$, and terminal condition $\gamma_\alpha(0) = (x_1, x_2, \theta)$.)

In general, one can find an AFDR whenever one can identify a finite family of analytic vector fields X_i, $i \in I$, possibly defined on other manifolds N_i endowed with analytic maps $\psi_i : N_i \to M$, such that there is a sufficient family $\mathcal{F}$ for our problem consisting of concatenations γ of a finite number $\sigma(\gamma)$ of pieces of the form $\psi_{i(j)} \circ \delta_j$, where δ_j is an integral curve of $X_{i(j)}$. The properness condition then amounts to requiring that the maps ψ_i be proper, and that there be local bounds on the number of switchings (i.e. for every compact $K \subseteq M$ and every $T > 0$ there is a ν such that, if $\gamma \in \mathcal{F}$, $T(\gamma) \leq T$, and $\gamma \subseteq K$, then $\sigma(\gamma) \leq \nu$).

An example of this that may be familiar to geometers is the case of minimum time control on a C^ω Riemannian manifold M. (The trajectories are arbitrary curves γ such that $||\dot{\gamma}(t)|| \leq 1$. Of course, the Pontryagin extremals in this case are the geodesics, and they are projections of the integral curves of the geodesic flow, which is a vector field on the tangent bundle TM.) In this case, the properness condition follows *if the metric is complete.*

In many other cases (e.g. linear time-optimal control with polyhedral control constraints, general C^ω problems $\dot{x} = f(x) + ug(x)$, $|u| \leq 1$, in dimension two, general problems where the optimal controls are bang-bang) the structure of the optimal trajectories has been investigated, leading to the construction of a proper AFDR. Moreover, the existence of an AFDR also implies a stronger conclusion, namely, the existence of a *regular synthesis*, i.e., roughly, a piecewise analytic feedback. (The precise definition is analogous, but not identical, to that given by Boltyanski in [1].) At the moment, the study of properties of trajectories, from the point of view of trying to find nice sufficient families, and if possible AFDR's, is an active area in which many challenging questions (e.g.: is it always true for an analytic complete system with, say, a control set $U = [-1, 1]^m$, that every nondegerate minimizer[6] has at most a countable set of points of nonanalyticity?) still remain open.

References

1. V.G. Boltyanski: Sufficient conditions for optimality and the justification of the Dynamic Programming Principle. SIAM J. Control **4**, pp. 326-361, 1966
2. R.W. Brockett: System Theory on group manifolds and coset spaces. SIAM J. Control **10**, pp. 265-284, 1972
3. R. Hermann: On the accessibility problem in control theory. In: Int. Symp. Nonlinear Diff. Equations and Nonlinear Mechanics, J.P. LaSalle, S. Lefschetz Eds., , pp. 325-332. New York: Academic Press, 1963
4. A. Isidori: Nonlinear Control Systems: An Introduction. Berlin: Springer-Verlag, second ed., 1989
5. A.J. Krener: A generalization of Chow's Theorem and the bang-bang theorem to nonlinear control problems. SIAM J. Control**12**(1), pp. 43-51, 1974
6. C. Lobry: Contrôlabilité des systèmes non-linéaires. SIAM J. Control **8**, pp. 573-605, 1978
7. S. Lojasiewicz Jr. and H.J. Sussmann: Some examples of reachable sets and optimal cost functions that fail to be subanalytic. S.I.A.M. J. Control and Optimization **23**, No. 4, pp. 584-598, 1985
8. H. Nijmeijer and A.J. Van der Schaft: Nonlinear dynamical control systems. New York: Springer-Verlag, 1990
9. J.A. Reeds and L.A. Shepp: Optimal paths for a car that goes both forwards and backwards. Pacific J. Math. **145**, pp. 367-393, 1990
10. H.J. Sussmann, ed.: Nonlinear Controllability and Optimal Control. New York: Marcel Dekker, 1990
11. H.J. Sussmann and G. Tang: "Shortest paths for the Reeds-Shepp car: a worked out example of the use of geometric techniques in nonlinear optimal control. Rutgers Center for Systems and Control Technical Report 91-10, September 1991

[6] A minimizer is *nondegenerate* if there is no other minimizer with the same initial and terminal points.