Optimal control and almost analytic feedback for some nonholonomic systems*

Héctor J. Sussmann
SYCON - Rutgers Center for Systems and Control
Department of Mathematics, Rutgers University, New Brunswick, NJ 08903

Abstract. One possible approach to do path planning for nonholonomic systems is to use optimal control. It is intuitively clear that optimal control should give rise to solutions of the path planning problem that are in feedback form, and one can expect that the resulting feedback control laws will be "piecewise smooth" or "piecewise analytic" in some sense. Here we focus on the weaker property of *analyticity on a large dense set* (ALDS), and explain how such a property is related to the problem of the structure of the optimal trajectories and, in particular, to the existence of "proper subanalytic sufficient families for optimality" that are made of a bounded number of pieces depending on finitely many parameters. The existence of such sufficient families can be proved in a number of cases using techniques from differential-geometric control theory. This is illustrated by briefly discussing the example of the "Reeds-Shepp car."
Key words: optimal control, nonholonomic systems, path finding

1. Introduction

We consider a system Σ without drift, of the form

$$\dot{x} = F(x,u) = u_1 f_1(x) + \ldots + u_m f_m(x)\,,\; x \in M\,, \quad (1)$$

where M is a connected real analytic manifold, and the vector fields $f_1, \ldots, f_m$ are real analytic. Let L_Σ denote the Lie algebra of vector fields on M generated by $f_1, \ldots, f_m$. Assume that Σ satisfies the *Lie algebra rank condition* (LARC), i.e. that, for each $x \in M$, the set $L_\Sigma(x) = \{X(x) : X \in L_\Sigma\}$ is the whole tangent space $T_x M$ to M at X. Under these conditions, it follows that for every pair (p,q) of points of M there exists an open-loop control $u(\cdot) : [a,b] \to \mathbb{R}^m$ that steers p to q. (This open-loop control can be taken to be in $L^1([a,b],\mathbb{R}^m)$, or in a more restricted class of controls, e.g. functions of class C^∞, or even polynomials. The interval $[a,b]$, with $a < b$, can also be chosen arbitrarily.)

The *motion planning problem* is the problem of finding, for each pair (p,q) of points of M, an open-loop control that steers p to q. In this note we will discuss the question of finding, for a given target point q, a family of open-loop controls $\mathbf{u} = \{u_p(\cdot)\}_{p\in M}$ such that for each p the control u_p steers p to q. Moreover, we want $\mathbf{u}$ to result from a reasonably regular "feedback control law" (FCL). So our first problem is to define what will be meant by an FCL. We could define an FCL to be a locally Lipschitz function $\varphi : M \to \mathbb{R}^m$. However, if we do so, then naturally the resulting vector field F_φ on M —defined by $F_\varphi(x) = F(x,\varphi(x))$— will have the uniqueness property of trajectories. In particular, it will not be possible for the trajectories from all possible initial points p to reach q in finite time. Alternatively, we could just require φ to be defined and locally Lipschitz on $M\backslash\{q\}$, and impose the extra condition that, for each initial p, the maximal trajectory $t \to \gamma_p(t) \in M\backslash\{q\}$ of F_φ that satisfies $\gamma_p(0) = p$ should exist on an interval $[0,T_p)$, with $T_p < \infty$, and satisfy $\lim_{t\to T_p} \gamma_p(t) = q$. However, such a feedback will in general not exist. For example, for a driftless system Σ given by (1), a locally Lipschitz feedback steering to q some neighborhood U of q cannot possibly exist if the vectors $f_1(q), \ldots, f_m(q)$ are linearly independent and $m < \dim M$. (This follows from the well-known Brockett topological necessary condition for stabilizability.) Moreover, one very common procedure for constructing FCL's that steer every point to some given target point q is to use *optimal control*, and we definitely want the laws obtained in this way to be FCL's in the sense of our definition. It is well known that, even in the simplest textbook examples of optimal synthesis, the optimal feedback can be discontinuous. (Consider, e.g., the problem of reaching $(0,0)$ in minimum time for the plane system $\dot{x} = y$, $\dot{y} = u$, with control constraint $|u| \le 1$.)

So in general we have to allow FCL's that are not locally Lipschitz, and possibly not even continuous. If we define an FCL to be a more general $\mathbb{R}^m$-valued function φ on M —or on $M\backslash\{q\}$— then uniqueness of trajectories is not guaranteed. So we choose a different approach, and directly define an FCL to be a *flow* — i.e. a family of open-loop controls and corresponding trajectories— with the special property of being *memoryless*. Precisely: a *feedback control law* (FCL) on a subset S of M is a mapping $\mu : p \to u_p^\mu$, defined for $p \in S$, that assigns to each $p \in S$ an open-loop control (i.e. an integrable $\mathbb{R}^m$-valued function) on a interval $[-T_p^\mu, 0]$, with $T_p^\mu \ge 0$, in such a way that (i) there is a trajectory $[-T_p^\mu, 0] \ni t \to \gamma_p^\mu(t) \in S$ of u_p^μ which goes through p at time $-T_p^\mu$, and (ii) the trajectories γ_p^μ —that are obviously unique— satisfy the *memoryless condition*: if $-T_p^\mu < t \le 0$, and $\gamma_p^\mu(t) = \tilde{p}$, then $-T_{\tilde{p}}^\mu = t$ and $u_{\tilde{p}}^\mu(s) = u_p^\mu(s)$ for almost all $s \in [t,0]$. (This implies in particular that $\gamma_{\tilde{p}}^\mu$ is just the restriction of γ_p^μ to $[-T_{\tilde{p}}^\mu, 0]$.)

Notice that, if the open-loop controls u_p^μ are continuous, then, if p, t, $\tilde{p}$ are as above, the equality $u_{\tilde{p}}^\mu(s) = u_p^\mu(s)$ will actually hold for *all* $s \in [t,0]$. If we then define $\varphi^\mu(p) = u_p^\mu(-T_p^\mu)$, then it is easy to see that each curve γ_p^μ is a solution of the equation $\dot{x}(t) = F_{\varphi^\mu}(x(t))$. However, there is no reason why the vector field F_{φ^μ} should be continuous, let alone locally Lipschitz, so we cannot characterize γ_p^μ as the unique trajectory of F_{φ^μ} that starts at p. In fact, it

*This research was supported in part by NSF Grant DMS-8902994.

can happen that two *different* FCL's μ_1, μ_2 give rise to the same function φ, so that this function contains less information than the full FCL μ.

If $\mu : p \to u_p^\mu$, $p \in S$, is an FCL on a set S, and $q \in S$, we say that μ *steers* S *to* q if $\gamma_p^\mu(0) = q$ for all $p \in S$. Our goal is to construct, for systems Σ as above, FCL's on M that steer M to a given target point q.

2. Existence of FCL's

Existence of a feedback is guaranteed by the following theorem, in which analyticity is not needed. (We recall that a system is *completely controllable* if for every pair of points p, q there is a trajectory from p to q.)

Theorem 1 *Consider a system Σ given by (1), where M is a connected smooth manifold, and the f_i are locally Lipschitz vector fields. Assume that Σ is completely controllable. Let $q \in M$. Then there exists a feedback control law $\mu : p \to u_p^\mu$ that steers M to q.*

PROOF. We choose a sequence $K_1, K_2, \ldots$ of compact subsets of M, such that $K_j \subseteq \mathrm{Int}\, K_{j+1}$ for each j, and $M = \cup_{j=1}^\infty K_j$. We let B be the unit ball $\{(u_1, \ldots, u_m) : u_1^2 + \ldots + u_m^2 \leq 1\}$ in $\mathbb{R}^m$. We define inductively sets H_j by starting with $H_0 = \{q\}$, and then letting H_j be the set of all those points in K_j from which it is possible to reach some point in H_{j-1} in time ≤ 1 by means of a trajectory of Σ that is entirely contained in K_j and corresponds to a B-valued control. It then follows easily by induction that H_j is compact and $H_j \subseteq H_{j+1}$ for each j, and $M = \cup_{j=0}^\infty H_j$. We define FCL's μ_j on H_j inductively. Assume μ_j has been defined. For $0 < t \leq 1$, let U_t be the space of all measurable B-valued functions on $[-t, 0]$. For $i \in \{1, \ldots, m\}$, $k \in \{0, 1, 2, \ldots\}$, let $I_{t,i,k} : U_t \to \mathbb{R}$ be the functional given by $I_{t,i,k}(u) = \int_{-t}^0 s^k u_i(s) ds$, for $u = (u_1, \ldots, u_m) \in U_t$. Order the set $\mathcal{P} = \{1, \ldots, m\} \times \{0, 1, 2, \ldots\}$ of all pairs (i, k) into a sequence $\pi_1, \pi_2, \ldots$, and write I_{t,π_ℓ} instead of $I_{t,i,k}$ if $\pi_\ell = (i, k)$. For each $p \in H_{j+1}$, let $t(p)$ be the infimum of the times of all trajectories of Σ that start at p, end at some point in H_j, are entirely contained in K_{j+1}, and correspond to a B-valued control. Then $t(p)$ is actually a minimum, i.e. the set $\mathcal{U}_p$ of those controls $u : [-t(p), 0] \to B$ that steer p to some point in H_j by means of a trajectory that is entirely contained in K_{j+1} is nonempty. Moreover, it is clear that $\mathcal{U}_p$ is compact with respect to the weak topology of $L^1([-t(p), 0], \mathbb{R}^m)$. We choose a control $u_p : [-t(p), 0] \to B$ as follows. If the time-minimization problem described above has a unique solution, i.e. if $\mathcal{U}_p$ has only one member, then we choose u_p to be this member. If not, we let $\mathcal{U}_p^1$ denote the set of all members of $\mathcal{U}_p$ that minimize the restriction to $\mathcal{U}_p$ of the functional $I_{t(p),\pi_1}$. The set $\mathcal{U}_p^1$ is nonempty and weakly compact (because $I_{t(p),\pi_1}$ is weakly continuous and $\mathcal{U}_p$ is weakly compact). If $\mathcal{U}_p^1$ has only one member, then we choose u_p to be this member. If not, we let $\mathcal{U}_p^2$ denote the set of all members of $\mathcal{U}_p^1$ that minimize the restriction of $I_{t(p),\pi_2}$ to $\mathcal{U}_p^1$. Again, the set $\mathcal{U}_p^2$ is nonempty and weakly compact. Continuing in this way, we produce a decreasing sequence $\{\mathcal{U}_p^\ell : \ell = 1, 2, \ldots\}$ of nonempty weakly compact subsets of $\mathcal{U}_p$. The intersection $\cap_{\ell=1}^\infty \mathcal{U}_p^\ell$ is therefore nonempty, and it is clear that it consists of exactly one control. This control is our u_p. Let $\gamma_p : [-t(p), 0] \to K_{j+1}$ be the corresponding trajectory. Then $\hat{p} = \gamma_p(0) \in H_j$, and it is easy to see that, if $-t(p) < s < 0$ and $\tilde{p} = \gamma_p(s)$, then $\tilde{p} \in H_{j+1}$, $t(\tilde{p}) = -s$, and $u_{\tilde{p}}$ is just the restriction of u_p to $[s, 0]$. We define $u_p^{\mu_{j+1}}$ to be the concatenation of u_p and $u_{\hat{p}}^{\mu_j}$, i.e. we let $T_p^{\mu_{j+1}} = t(p) + T_{\hat{p}}^{\mu_j}$, and then define $u_p^{\mu_{j+1}}(s) = u_{\hat{p}}^{\mu_j}(s)$ for $-T_{\hat{p}}^{\mu_j} \leq s \leq 0$, $u_p^{\mu_{j+1}}(s) = u_p(s + T_{\hat{p}}^{\mu_j})$ for $-T_{\hat{p}}^{\mu_{j+1}} \leq s < -T_{\hat{p}}^{\mu_j}$.

Finally, if $p \in M$, then $p \in H_j$ for some j, and it is easy to see that the control $u_p^{\mu_j}$ does not depend on j. We let $u_p^\mu = u_p^{\mu_j}$. It then follows easily that this prescription defines an FCL. $\square$

Naturally, we want FCL's with extra regularity properties. For instance, we can call an FCL μ *open-loop continuous* if each control u_p^μ is continuous, and we can define in a similar way other open-loop properties, such as open-loop analyticity. Under the real analyticity assumptions made here, the main result of [4] applies, so we can conclude that there exists a "subanalytic feedback" steering M to q. Such a feedback however, will typically involve discontinuous open-loop controls. It is natural to ask whether one can always guarantee the existence of an open-loop analytic FCL.

It turns out that the existence of such a feedback can be proved for systems Σ that satisfy the "strong bracket generating condition" (SBGC), introduced by R. Strichartz in [3]. (We do not know if the conclusion is true for more general real-analytic systems.)

Precisely, we say that the m-tuple $(f_1, \ldots, f_m)$ satisfies the SBGC if, whenever $\alpha = (\alpha_1, \ldots, \alpha_m)$ is an m-tuple of real numbers such that $\alpha \neq (0, \ldots, 0)$, and $x \in M$, then it follows that the vectors $f_1(x), \ldots, f_m(x), [\alpha f, f_1](x), \ldots, [\alpha f, f_m](x)$ span the tangent space $T_x M$. (Here αf denotes the vector field $\alpha_1 f_1 + \ldots + \alpha_m f_m$.) Equivalently, the SBGC holds if and only if, whenever $x \in M$, and λ is a nonzero covector at x such that $\langle \lambda, f_i(x) \rangle = 0$ for $i = 1, \ldots, m$, it follows that the matrix $\{\langle \lambda, [f_i, f_j](x) \rangle\}_{i,j=1,\ldots,m}$ is nonsingular. We then have:

Theorem 2 *Consider a system Σ given by (1), where M is a connected real analytic manifold, and the f_i are real analytic vector fields that satisfy the SBGC. Let $q \in M$. Then there exists an FCL $\mu : p \to u_p^\mu$ that steers M to q and is such that each open-loop control u_p^μ is real analytic.*

PROOF. We use the well known fact that a C^ω manifold admits a complete C^ω Riemannian metric $\mathbf{g}$. We let $h(x) = (1 + ||f_1(x)||^2 + \ldots + ||f_m(x)||^2)^{-\frac{1}{2}}$, where the length $||v||$ of a tangent vector is measured with respect to the metric $\mathbf{g}$. Then let $g_i(x) = h(x) f_i(x)$, so that the g_i are real analytic vector fields that satisfy $||g_i(x)|| \leq 1$ for all i, x.

We now consider instead of (1) the system

$$\tilde{\Sigma} : \quad \dot{x} = u_1 g_1(x) + \ldots + u_m g_m(x) \,. \qquad (2)$$

It is clear that any FCL μ with the desired property for $\tilde{\Sigma}$ gives rise to a feedback μ for Σ that also satisfies the desired conclusion. (It suffices to define $u_p^\mu(t) = h(\gamma_p^\nu(t))u_p^\nu(t)$, for $t \in [-T_p^\nu, 0]$.) Moreover, the vector fields g_i also satisfy the SBGC. This means that it suffices to prove our conclusion for (1) under the extra hypothesis that there exists a complete Riemannian metric with respect to which $||f_i(x)|| \leq 1$ for all i, x.

If we let B denote the closed unit ball in $\mathbb{R}^m$, and consider the problem $\mathbf{P}_p$ of steering p to q in *minimum time* (recall that q is fixed) by means of a B-valued control, it now follows that $\mathbf{P}_p$ has at least one solution. Choose, in any fashion whatsoever, a solution u_p^μ of $\mathbf{P}_p$. It will turn out that this defines a feedback. But, in order to show this, we first prove that all the solutions of $\mathbf{P}_p$ are real analytic.

Let $v : [-T, 0] \to C$ be a solution of $\mathbf{P}_p$, and let $\gamma : [-T, 0] \to M$ be the corresponding trajectory, so $\gamma(-T) = p$, $\gamma(0) = q$. We apply the Pontryagin Maximum Principle, and conclude that there exists a covector-valued function $t \to \lambda(t) \in T^\#_{\gamma(t)}M$ on $[-T, 0]$ (where $T^\#_x M$ denotes the cotangent space of M at x with the origin removed) that satisfies — in local coordinates— the *adjoint equation* and the *Hamiltonian minimization condition* (HMC). The adjoint equation says that $\dot{\lambda}(t) = -\frac{\partial H}{\partial x}(\lambda(t), \gamma(t), v(t))$, where the *Hamiltonian* H is the function defined by $H(\lambda, x, u) = \sum_i u_i\langle\lambda, f_i(x)\rangle$. The HMC says that there is a nonnegative constant λ^0 such that, for almost every t, the control $v(t)$ minimizes the function $B \ni w \to H(\lambda(t), \gamma(t), w) \in \mathbb{R}$, and $H(\lambda(t), \gamma(t), v(t)) = -\lambda^0$. If $\lambda^0 \neq 0$ then, after dividing $\lambda(t)$ by λ^0, we may assume that $\lambda^0 = 1$. In that case, the HMC implies that the components $v_i(t)$ of $v(t)$ are given by $v_i(t) = -\langle\lambda(t), f_i(\gamma(t))\rangle$. It then follows easily that the curve $\Gamma : t \to (\gamma(t), \lambda(t))$ is a solution of the Hamiltonian system $\dot{x} = \frac{\partial \mathcal{H}}{\partial \lambda}$, $\dot{\lambda} = -\frac{\partial \mathcal{H}}{\partial x}$, where $\mathcal{H}(\lambda, x) = -\frac{1}{2}\sum\langle\lambda, f_i(x)\rangle^2$. So Γ is real analytic. Since the components v_i are analytic functions of λ and x, it follows that they are analytic as well.

So the analyticity of the optimal controls is proved, provided that we exclude the possibility that $\lambda^0 = 0$, i.e. provided that we show that there do not exist time-optimal "exceptional extremals." The SBGC does precisely that, as can be verified by direct calculation (cf. Strichartz [3]).

Now we are ready to prove that the map $p \to u_p^\mu$ is an FCL. If $p \in M$, $-T_p^\mu < t < 0$, and $\tilde{p} = \gamma_p^\mu(t)$, then we can define a control $v : [-T_p^\mu, 0] \to C$ that is equal to u_p^μ on $[-T_p^\mu, t]$, and to $u_{\tilde{p}}^\mu$ on $(t, 0]$. Clearly, v is C-valued, and steers p to q in time T_p^μ. So v is also time-optimal, and therefore analytic. This implies that $u_{\tilde{p}}^\mu \equiv u_p^\mu$ on $[t, 0]$, so the memoryless property holds. □

3. Almost analytic FCL's

Theorems 1 and 2 show that optimal control can be used to produce FCL's. We now discuss how one can actually use optimal control to get FCL's with some extra regularity properties *as functions of the state*. In particular, we know that in many cases there is a "very large" set —typically, the complement of a "switching locus" which is a union of lower-dimensional submanifolds— where the optimal feedback is smooth, and even analytic. So we will seek FCL's that are "almost analytic," in the sense that they satisfy the condition of *analyticity on a dense set* (ADS), or have the stronger *analyticity on a large dense set* (ALDS) property.

To define these properties, we first agree to call a subset Ω of M a *set of analyticity* of an FCL μ if Ω is open in M and satisfies: (i) for each $p \in \Omega$ the function u_p^μ is continuous on some interval $[-T_p^\mu, -T_p^\mu + \varepsilon]$, for some ($p$-dependent) $\varepsilon > 0$, and (ii) the associated function φ^μ —defined on Ω by $\varphi^\mu(p) = u_p^\mu(-T_p^\mu)$ — is real analytic on Ω. We will call Ω *large* if $M \backslash \Omega$ is a finite or countable union of connected embedded real analytic submanifolds of positive codimension.

We will say that an FCL μ *has the ADS property* if there is an open dense subset Ω which is a set of analyticity for μ. If Ω can be chosen to be large, then we will say that μ *has the ALDS property.*

If μ has the ADS property, then the trajectories γ_p^μ, as long as they are contained in Ω, are the integral curves of the real analytic vector field F_{φ^μ}, defined by $F_{\varphi^\mu}(p) = F(p, \varphi^\mu(p))$.

We will consider the use of *time-optimal control* to obtain FCL's with the ADS property. We let C be a compact, convex, subanalytic subset of $\mathbb{R}^m$ such that $0 \in \operatorname{Int} C$. (Cf. [6] for the definition and basic properties of subanalytic sets.)

We then consider the control system Σ_C obtained from Σ by imposing the restriction that the controls take values in C. We fix the target point q and consider, for each $p \in M$, the problem $\mathcal{P}_p^C$ of finding a minimum time trajectory of Σ_C from p to q. Such a problem will have a solution for all p, if Σ_C satisfies the *nonexplosion condition*:

(NE) For every measurable $u : (-\infty, 0] \to C$, the trajectory γ of Σ that corresponds to u and satisfies $\gamma(0) = q$ is defined for all $t \in (-\infty, 0]$.

We define the *value function* $V : M \to \mathbb{R}$ by letting $V(p)$ be the infimum of the times of all the trajectories of Σ_C that go from p to q. If (NE) holds, then this infimum is actually a minimum.

The following theorem shows that we can deduce the existence of a feedback with the ADS property if the value function is analytic on an open dense set.

Theorem 3 *Let M be a connected real analytic manifold, and let $f_1, \ldots, f_m$ be real analytic vector fields on M that satisfy the Lie algebra rank condition. Let C be a compact, convex, subanalytic subset of $\mathbb{R}^m$ such that $0 \in \operatorname{Int} C$. Let $q \in M$. Assume that the nonexplosion condition (NE) holds for Σ_C and q. Let $V : M \to \mathbb{R}$ be the value function for the minimum time problem for Σ_C with target q. Assume that V is real analytic on an open subset Ω of M. Then there is an open subset $\tilde{\Omega}$ which is large as a subset of Ω, and a time-optimal feedback control law μ that steers M to q, such that $\tilde{\Omega}$ is a set of analyticity of μ. In particular, if Ω is dense in M then $\tilde{\Omega}$ is dense as well, so μ has the ADS property. If Ω is large as a subset of M, then the same is true of $\tilde{\Omega}$, so that μ has the ALDS property.*

PROOF. It is easy to see that, if V is continuously differentiable on a neighborhood of p, then the gradient $\nabla V(p)$ satisfies $\min\{\langle\nabla V(p), F(p,u)\rangle : u \in C\} = -1$. For $p \in \Omega$, let $U(p)$ be the set of those $u \in C$ such that $\langle\nabla V(p), F(p,u)\rangle = -1$, so $U(p)$ is compact and nonempty. It follows from general properties of subanalytic sets that the set $G(U) = \{(p,u) : p \in \Omega \text{ and } u \in U(p)\}$ (i.e. the graph of the multivalued function U) is subanalytic in $\Omega \times \mathbb{R}^m$. From this one can get a single-valued subanalytic selection, i.e. a subanalytic function $\eta : \Omega \to C$ such that $\eta(p) \in U(p)$ for each p. (A subanalytic function is a function whose graph is subanalytic. The existence of a subanalytic selection follows, for example, by choosing $\eta(p)$ to be the first element of the compact set $U(p)$ with respect to the lexicographical ordering of $\mathbb{R}^m$.) Finally, a subanalytic function on Ω is actually analytic on a large open dense subset $\tilde{\Omega}$ of Ω. (Cf., e.g., [6].)

We now let $\tilde{\eta}$ be the restriction of η to a large open dense subset $\tilde{\Omega}$ of Ω where η is analytic. The "bad set" $B = M \backslash \tilde{\Omega}$ is then closed. We let X be the vector field on $\tilde{\Omega}$ defined by $X(x) = F(x, \tilde{\eta}(x))$. It then follows that the directional derivative $D_X V$ of V in the direction of X satisfies $D_X V \equiv -1$ on $\tilde{\Omega}$.

We then define, for each p, a class $\mathcal{U}_p$ of controls, as follows: we let $\mathcal{U}_p$ consist of all the measurable C-valued controls u, defined on the interval $[-V(p), 0]$, that steer p to q and have the property that, if γ^u denotes the trajectory of u that starts at p at time $-V(p)$ (so that $\gamma^u(0) = q$), then, whenever I is a relatively open subinterval of $[-V(p), 0]$ such that $\gamma^u(t) \in \tilde{\Omega}$ for all $t \in I$, it follows that $u(t) = \tilde{\eta}(\gamma^u(t))$ for $t \in I$. It is easy to see that, for each p, $\mathcal{U}_p$ is weakly compact. If we could also show that $\mathcal{U}_p$ is always nonempty, then we could define a feedback μ by selecting, from each $\mathcal{U}_p$, a member u_p^μ as in the proof of Theorem 1. The resulting feedback clearly has $\tilde{\Omega}$ as a set of analyticity.

So all we need is to prove that $\mathcal{U}_p$ is always nonempty. We do this by fixing p and constructing a member u of $\mathcal{U}_p$, by means of "Euler approximations." For an integer $N > 0$, we let $\pi^N = (t_0, t_1, \ldots, t_N)$ be a partition of $[-V(p), 0]$ into N equal subintervals, so that $t_j = \frac{j-N}{N} V(p)$ for $j = 0, 1, \ldots, N$. We now define u^N and the corresponding trajectory γ^N on $[t_{j-1}, t_j]$ by induction on j. Assuming that u^N and γ^N have already been defined on $[t_0, t_{j-1}]$, and that $V(p_{j-1}) = -t_{j-1}$, where $\gamma^N(t_{j-1}) = p_{j-1}$, we define u^N on $[t_{j-1}, t_j]$ by distinguishing two cases. The first case is when $p_{j-1} \in \tilde{\Omega}$ and the integral curve θ of X that goes through p_{j-1} at time t_{j-1} is defined —and hence takes values in $\tilde{\Omega}$— on the whole interval $[t_{j-1}, t_j]$. In this case we let $\gamma^N(t) = \theta(t)$ and $u^N(t) = \tilde{\eta}(\theta(t))$ for $t \in [t_{j-1}, t_j]$. The second case is when either $p_{j-1} \notin \tilde{\Omega}$ or $p_{j-1} \in \tilde{\Omega}$ but θ fails to be defined on the whole interval $[t_{j-1}, t_j]$. In this case we pick an arbitrary time-optimal control for Σ_C that steers p_{j-1} to q and is defined on $[t_{j-1}, 0]$, and then let u^N be the restriction of this control to $[t_{j-1}, t_j]$. In either case, it is clear that $V(p_j) = -t_j$. So the inductive definition can be continued all the way. Since $V(p_j) = -t_j$ for each j, we get in particular $V(\gamma^N(0)) = 0$, so that $\gamma^N(0) = q$. We now use the weak compactness of the space of C-valued measurable functions on $[-V(p), 0]$ and find a sequence $\{N(k) : k = 1, 2, \ldots\}$ such that $N(k) \to \infty$ and $u^{N(k)}$ converges weakly to a control v as $k \to \infty$. Then the trajectories $\gamma^{N(k)}$ converge to the trajectory $\gamma : [-V(p), 0] \to M$ of v that goes through p at time $-V(p)$. Clearly, $\gamma(0) = q$, so v is time-optimal. If I is a relatively open subinterval of $[-V(p), 0]$ such that $\gamma(t) \in \tilde{\Omega}$ for $t \in I$, and $\alpha < \beta$, $[\alpha, \beta] \subseteq I$, then we can pick compact subsets J_1, J_2 of $\tilde{\Omega}$ such that $\gamma([\alpha, \beta]) \subset \operatorname{Int} J_1$ and $J_1 \subset \operatorname{Int} J_2$. Then there is a $\delta > 0$ such that every trajectory of X that starts at a point in J_1 at time 0 exists on $[0, \delta]$. If k is sufficiently large, then $\gamma^{N(k)}([\alpha, \beta]) \subset \operatorname{Int} J_1$ as well. If, in addition, we require that $\frac{V(p)}{N(k)} < \sigma < \delta$, and $t_{j_1}, t_{j_1+1}, \ldots, t_{j_2}$ are the partition points of $\pi^{N(k)}$ that lie in $[\alpha, \beta]$, then $\gamma^{N(k)}(t_j) \in J_1$ for $j_1 \le j \le j_2$. Therefore $\gamma^{N(k)}$ is an integral curve of X on each interval $[t_j, t_{j+1}]$ for $j_1 \le j \le j_2$. So γ is an integral curve of X on $[\alpha + \sigma, \beta]$. Since σ, α, β are arbitrary, we see that γ is an integral curve of X on I. This proves that v belongs to $\mathcal{U}_p$. □

Theorem 3 shows that one can obtain conclusions about existence of "almost analytic" optimal FCL's if one can prove that the value function V is analytic on a reasonably large set. On way to do that is to prove that V is *subanalytic*. Precisely, we recall (cf., e.g., [6]) that for a subanalytic locally bounded function on M there is a real analytic stratification of M such that V is analytic on each stratum. If we let Ω be the union of all the open strata, we see that Ω is open, dense and large, and V is analytic on Ω. In view of Theorem 3, this implies:

Theorem 4 *Let M, $f_1, \ldots, f_m$, C, $q \in M$ satisfy the hypotheses of Theorem 3. Let $V : M \to \mathbb{R}$ be the value function for the minimum time problem for Σ_C with target q. Assume that V is subanalytic. Then there is a large open subset Ω on which V is analytic. In particular, there exists a time-optimal FCL μ that steers M to q and has the ALDS property.*

4. Trajectory analysis and almost analyticity

Theorem 4 naturally leads us to the question of how to prove, for various problems Σ_C and targets q, that the value function is subanalytic. This issue has been discussed in several papers —cf., e.g., [5], [7], and the references given there— so here we will just outline the main facts.

The main point in proving subanalyticity of V is to show that one can put together an optimal synthesis using a small "subanalytically parametrized" family of trajectories. The precise definitions are as follows. (For simplicity, we will assume from now on that the vectors $f_1(x), \ldots, f_m(x)$ are linearly independent for every x. This implies that a trajectory $\gamma : [a, b] \to M$ uniquely determines the corresponding control $u : [a, b] \to C$, so we will not have to talk about "admissible pairs" consisting of a control and corresponding trajectory.) A family $\mathcal{F}$ of trajectories of a system Σ_C will be called *sufficient for time-*

optimality with target q if, for every p from which q can be reached, $\mathcal{F}$ contains a time-optimal trajectory from p to q. Notice that the definition makes perfect sense for systems with drift. In the driftless case, if the Lie algebra rank condition holds, then the words "from which q can be reached" are redundant, since q can be reached from every p. In that case, if the non-explosion property (NE) hold, then the set $\mathrm{Traj}(\Sigma_C)$ of *all* trajectories of Σ_C is sufficient. On the other hand, it is clear that one can find sufficient subsets of $\mathrm{Traj}(\Sigma_C)$ that are much smaller. For instance, it suffices to choose in some way a time-optimal trajectory γ^p from p to q for each p. In this way, we obtain a sufficient family that is "finite-dimensionally parametrized," in the sense that the γ^p depend on the finite-dimensional parameter p.

This, however, does not imply that the dependence on p is "nice." We are interested is singling out a particular class of parametrized families, in which the dependence on the parameter is subanalytic.

Precisely, we consider a *parametrized family* $\Gamma = \{\Gamma^\alpha : \alpha \in A\}$ of trajectories, depending on a parameter α that takes values in an *ambient parameter space* A. We will allow the possibility that Γ^α is only defined for α in some subset of A, called the *domain of* Γ, and denoted by $\mathrm{Dom}(\Gamma)$. Each Γ^α is defined on an interval $[a_\Gamma^\alpha, b_\Gamma^\alpha]$.

Associated to a parametrized family Γ as above, there are several sets and maps, and we can define various concepts of subanalyticity depending on which sets or maps we require to be subanalytic. For instance, the graph $G(\Gamma)$, is defined to be the set $\{(\alpha, t, x) : \alpha \in \mathrm{Dom}(\Gamma), a_\Gamma^\alpha \le t \le b_\Gamma^\alpha, x = \Gamma^\alpha(t)\}$. The *initial point map* IP_Γ is the map $\alpha \to \Gamma(a_\Gamma^\alpha)$. The *initial time map* IT_Γ is the map $\alpha \to a_\Gamma^\alpha$. We can combine IP_Γ and IT_Γ into one map, denoted I_Γ, given by $I_\Gamma(\alpha) = (a_\Gamma^\alpha, \Gamma(a_\Gamma^\alpha))$. Similarly, we can define the *terminal point map* TP_Γ, the *terminal time map* TT_Γ and the combined map T_Γ.

We will call Γ *subanalytic* if A is a real analytic NNPD manifold, and the graph $G(\Gamma)$ is subanalytic as a subset of $A \times \mathbb{R} \times M$. Here the abbreviation "NNPD" stands for "not necessarily of pure dimension." An NNPD manifold is a manifold in the usual sense, with a finite or countably infinite collection of connected components, that are allowed to be of different dimensions. (In particular, it is not even required that the dimensions be uniformly bounded. For instance, the union of all the spaces $\mathbb{R}^n$, for all $n \in \{0, 1, 2, \ldots\}$, endowed with a manifold structure such that each $\mathbb{R}^n$ is a component, is an NNPD manifold.) Similarly, we will call Γ *I-subanalytic* if the map I_Γ is subanalytic (i.e. if the graph $G(I_\Gamma) = \{(\alpha, a_\Gamma^\alpha, \Gamma(a_\Gamma^\alpha)) : \alpha \in \mathrm{Dom}(\Gamma)\}$ is subanalytic as a subset of $A \times \mathbb{R} \times M$). Similarly, one can define *IP-subanalyticity, IT-subanalyticity,* and analogous notions involving the T maps.

It is easy to see that subanalyticity of Γ implies I-subanalyticity. On the other hand, I-subanalyticity implies IP-subanalyticity *provided that the map* IT_Γ *is* (Identity $\times\, IP_\Gamma$)*-bounded.* (This means that for every compact subset K of $A \times M$ there exists $R > 0$ such that $|IT_\Gamma(\alpha)| \le R$ whenever $\alpha \in \mathrm{Dom}(\Gamma)$ and $(\alpha, IP_\Gamma(\alpha)) \in K$. In general, if ψ, θ are (possibly partially defined) maps from a set S to topological spaces X, Y, we say that ψ is *θ-bounded* if for every compact subset K of Y there is a compact subset J of X such that $\psi(s) \in J$ whenever $s \in \mathrm{Dom}(\psi) \cap \mathrm{Dom}(\theta)$ and $\theta(s) \in K$. Similarly, I-subanalyticity implies IT-subanalyticity if IP_Γ is (Identity $\times\, IT_\Gamma$)-bounded.

A parametrized family Γ as above will be called *I-proper* if the map I_Γ is proper, i.e. if the identity map of $\mathrm{Dom}(\Gamma)$, regarded as a partial map from A to A, is I_Γ-bounded. In order words, Γ is I-proper if we can give "a priori bounds" for α in terms of the initial condition $(IT_\Gamma(\alpha), IP_\Gamma(\alpha))$.

We illustrate all these definitions with a simple example. Consider a linear control system $\dot{x} = AX + Bu$, $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, with a polyhedral control constraint $u \in C$, where C is the convex hull of a finite subset C_0 of $\mathbb{R}^m$. (This example has drift, but our definitions do not require that the system be driftless.) Suppose we are interested in reaching the origin in minimum time. We can construct a sufficient family $\mathcal{F}$ by taking all the *bang-bang trajectories* that end at 0 at time 0, i.e. all trajectories $\gamma : [-T, 0] \to \mathbb{R}^n$ that satisfy $\gamma(0) = 0$ and correspond to piecewise constant C_0-valued controls. Not every trajectory in $\mathcal{F}$ is optimal, and not every optimal trajectory is in $\mathcal{F}$, but the bang-bang theorem says that for every point p from which 0 can be reached there is a time-optimal trajectory in $\mathcal{F}$ that goes from p to 0. So $\mathcal{F}$ is sufficient in our sense. Moreover, it is easy to exhibit a parametrization that turns $\mathcal{F}$ into a subanalytic family Γ. For each finite sequence σ of elements of C_0, we let $|\sigma|$ denote the length of σ, and let A_σ be a copy of $\mathbb{R}^\sigma$. We take the ambient parameter space A to be the disjoint union of the A_σ's for all finite sequences σ. (Notice that A is a real analytic NNPD manifold, but has components of all possible dimensions, so it is not a "manifold" in the usual sense.) The actual domain of Γ will be the union of the nonnegative orthant A_σ^+ of A_σ, for all σ's. If $\alpha \in \mathrm{Dom}(\Gamma)$, let $\sigma = (u^1, \ldots, u^N)$ be such that $\alpha \in A_\sigma^+$. Let $\alpha = (\tau^1, \ldots, \tau^N)$. We let Γ^α be the trajectory that ends at 0 at time 0 and corresponds to the bang-bang control u^α that equals u^1 during time τ^1, then u^2 during time τ^2, etc. Therefore $IT_\Gamma = -(\tau^1 + \ldots + \tau^N)$. It is clear that Γ is subanalytic, I-subanalytic, IP-subanalytic and IT-subanalytic. (The last two follow because a bound on α automatically gives bounds for $IT_\Gamma(\alpha)$ and $IP_\Gamma(\alpha)$.) However, Γ as defined is *not* proper. Indeed, properness of Γ would mean that, if we specify $R > 0$, then the bound $|IT_\Gamma(\alpha)| \le R$ forces α to be in a compact subset of A. This is clearly false, since a compact subset of A must be contained in the union of finitely many connected components of A, whereas a bang-bang trajectory whose total time is bounded by R can still have an arbitrarily large number of switchings.

We are rescued from this situation by the *bang-bang theorem with bounds on the number of switchings*, according to which there exists, for every $R > 0$, an integer $N(R)$ such that, if q can be reached from p in time $\le R$, then there is a control that time-optimally steers p to q and is bang-bang with at most $N(R)$ switchings. (We will choose $N(R)$ to be the smallest integer with this property, so $N(R)$ is a monotonically nondecreasing function of R.) Using this, we construct a new family $\tilde{\Gamma}$ by restricting Γ to a subset $\tilde{A}$ of Dom(Γ). We choose $\tilde{A}$ to be the union of the sets $\tilde{A}^+_\sigma$, where $\tilde{A}^+_\sigma$ is the set of those $|\sigma|$-tuples $(\tau^1, \ldots, \tau^{|\sigma|}) \in A^+_\sigma$ such that $\tau^1 + \ldots + \tau^{|\sigma|} \geq R_{|\sigma|}$, and, for an integer $\nu > 0$, we let $R_\nu = \inf\{R : N(R) \geq \nu - 1\}$. The family $\tilde{\Gamma}$ is subanalytic, I-subanalytic, IP-subanalytic, IT-subanalytic, and sufficient for time optimality with target 0. Moreover, $\tilde{\Gamma}$ *is I-proper.* Indeed, if $R > 0$, and $|IT_{\tilde{\Gamma}}(\alpha)| \leq R$, then α must lie in $\tilde{A}^+_\sigma$ for a σ such that $|\sigma| \leq N(R) + 1$. Since the components of α are ≥ 0, and their sum is $\leq R$, we see that α lies in a compact subset $K(R)$ of A.

The main reason why subanalyticity and properness are important is the following result, whose proof is given, for instance, in [7].

Theorem 5 *Assume that there exists an I-proper, I-subanalytic family Γ of trajectories, all of which have terminal point q and terminal time 0, such that Γ is sufficient for time optimality with target q. Then the value function for the problem of reaching q in minimum time is subanalytic.*

Theorem 5 applies, in particular, to the linear systems example described above. It is clear, however, that the theorem also applies to many other optimal control problems. What is needed is a construction of "nice" sufficient families, where the precise meaning of "nice" is "I-subanalytic and I-proper." To do this for a concrete problem, one has to carry out an appropriate *trajectory analysis.* One has to analyze the structure of the optimal trajectories, with the aim of proving some kind of piecewise regularity of the optimal open-loop controls, *bounds on the number of switchings.* It is not necessary, however, to show that *all* the optimal trajectories have finitely many switchings. All one needs is to be able to *select* a class of trajectories which is rich enough to contain at least one optimal trajectory for each initial point, and on the other hand satisfies a bound on the number of switchings. That is, instead of insisting on results that give some kind of "piecewise regularity" for all possible optimal trajectories, we seek to prove weaker PRS ("piecewise regularity with selection") theorems, i.e. results that guarantee the existence of sufficiently many piecewise regular optimal trajectories.

In recent years, a number of techniques have been developed to prove PRS theorems for large classes of problems. (Cf. the references given in [7].)

Here we will briefly describe the results for one example, namely, the Reeds-Shepp car. (We refer the reader to [8] for the —rather long— analysis.)

The "Reeds-Shepp car" is the system RS given by

$$\dot{x}_1 = u\cos\theta\,,\ \dot{x}_2 = u\sin\theta\,,\ \dot{\theta} = v\,, \qquad (3)$$

with state space $R^2 \times S^1$ and control constraints $|u| = 1$, $|v| \leq 1$. We assume throughout that we have fixed a target point q. The Lie algebra rank condition holds, so RS is completely controllable. However, the set $U = \{-1,1\} \times [-1,1]$ of control values is compact but not convex, so it is not *a priori* obvious that any two points can be joined by an optimal trajectory. This turns to be true, but to see it one needs to consider first the *convexified Reeds-Shepp* (CRS) system, in which the control set is the whole square $[-1,1]$. We remark that, for CRS, it would be totally hopeless to try to prove that all time-optimal trajectories are piecewise smooth. Indeed, it is easy to see that every trajectory that corresponds to a control for which $v(t) \equiv 1$, and to a completely arbitrary measurable function $t \to u(t) \in [-1,1]$, is time-optimal, provided only that its time interval has length $\leq \pi$. So the set of all optimal trajectories is too large. However, in order to put together an optimal synthesis, one does not need all the trajectories, but only a sufficiently rich subfamily. The detailed analysis given in [1] and [8] shows that a sufficient family $\mathcal{F}$ for CRS is obtained by considering all the trajectories for piecewise constant controls with values in the set $C_0 = \{(1,1),(1,0),(1,-1),(-1,1),(-1,0),(-1,-1)\}$ and at most four switchings. (The points $(1,0)$ and $(-1,0)$ correspond to *singular controls.*) The members of $\mathcal{F}$ turn out to be trajectories of the original, nonconvexified RS problem, so $\mathcal{F}$ is sufficient for RS, and existence of time-optimal trajectories for RS follows. Theorem 5 then tells us that the value function for RS is subanalytic, so that, in particular, RS admits a time-optimal FCL with the ALDS property.

References

[1] Reeds, J.A. and L.A. Shepp, "Optimal paths for a car that goes both forwards and backwards," *Pacific J. Math.*, **145** (1990), pp. 367-393.

[2] R. Strichartz, "Sub-Riemannian Geometry," *J. Diff. Geom.* 24, 221-263, (1983).

[3] Strichartz, R., "Corrections to 'Sub-Riemannian Geometry'," *J. Diff. Geom.* 30, no. 2, 595-596, (1989).

[4] Sussmann, H.J., "Subanalytic sets and feedback control," *J. Diff. Equations* **31**, No.1 (1979), pp. 31-52.

[5] Sussmann, H.J., "Lie brackets, real analyticity and geometric control theory," in *Differential Geometric Control Theory*, R.W. Brockett, R.S. Millman and H.J. Sussmann Eds., Birkhäuser Boston Inc. (1983), pp. 1-115.

[6] Sussmann, H.J., "Real analytic desingularization and subanalytic sets: an elementary approach," *Trans. Amer. Math. Soc.* **317** (1990), pp. 417-461.

[7] Sussmann, H.J., "Synthesis, presynthesis, sufficient conditions for optimality, and subanalytic sets," in *Nonlinear Controllability and Optimal Control,* H. J. Sussmann Ed., M. Dekker, Inc., New York, 1990, pp. 1-20.

[8] Sussmann, H.J., 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 (Sycon) Report 91-10, September 1991,