



Technical communique

Contact geometry of the Pontryagin maximum principle[☆]

Tomoki Ohsawa

Department of Mathematical Sciences, The University of Texas at Dallas, 800 W Campbell Rd, Richardson, TX 75080-3021, United States

ARTICLE INFO

Article history:

Received 27 June 2014

Received in revised form

17 November 2014

Accepted 29 January 2015

Keywords:

Optimal control

Maximum principle

Differential geometric methods

Geometric approaches

ABSTRACT

This paper gives a brief contact-geometric account of the Pontryagin maximum principle. We show that key notions in the Pontryagin maximum principle – such as the separating hyperplanes, costate, necessary condition, and normal/abnormal minimizers – have natural contact-geometric interpretations. We then exploit the contact-geometric formulation to give a simple derivation of the transversality condition for optimal control with terminal cost.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

It is well known that a necessary condition for optimality of the Pontryagin maximum principle may be interpreted as a Hamiltonian system, and so its geometric formulation usually exploits the language of symplectic geometry; see e.g., Agrachev and Sachkov (2004, Chapter 12), Jurdjevic (1997, Chapter 11).

The main focus of this paper is to change this perspective slightly to look at the maximum principle from the point of view of contact geometry, the “odd-dimensional cousin” (Arnold, 1989, Appendix 4) of symplectic geometry. The correspondence between contact and symplectic Hamiltonian systems is elementary and well known (see, e.g., Arnold, 1989, Appendix 4), and thus switching between symplectic and contact views is fairly trivial as far as the mathematical technicality is concerned. Our stress here is rather that the language of contact geometry fits more naturally to a proof of the Pontryagin maximum principle and so we may exploit the contact-geometric view from the outset. It also provides an alternative geometric perspective on applications of the maximum principle.

2. Geometry of optimal control on \mathbb{R}^n

2.1. Extended system

Let \mathbb{R}^n be the state space and \mathcal{U} be a compact subset of \mathbb{R}^m that defines the space of controls. Define a control system by $f: \mathbb{R}^n \times$

$\mathcal{U} \rightarrow \mathbb{R}^n$, and let $L: \mathbb{R}^n \times \mathcal{U} \rightarrow \mathbb{R}$ be the cost function and S_1 be a submanifold of \mathbb{R}^n . Then consider the following optimal control problem:

$$\min_{u(\cdot) \in \mathcal{U}} \int_{t_0}^{t_1} L(x(t), u(t)) dt \quad (1)$$

subject to $\dot{x} = f(x, u)$, $x(t_0) = x_0$, $x(t_1) \in S_1$,

where the initial time $t_0 \in \mathbb{R}$ and initial point $x_0 \in \mathbb{R}^n$ are fixed whereas the terminal time $t_1 \in \mathbb{R}$ is free.

Recall (see, e.g., Lewis (unpublished), Liberzon (2012, Chapter 4) and Pontryagin, Boltyanskii, Gamkrelidze, and Mishchenko (1962)) that the first step in proving the Pontryagin maximum principle is to introduce a new variable (running cost) x^0 by

$$x^0(t) := \int_{t_0}^t L(x(s), u(s)) ds,$$

that is, x^0 may be regarded as a solution of the differential equation $\dot{x}^0 = L(x, u)$

with the initial condition $x^0(t_0) = 0$. One then augments the original control system by the above system: we define the extended state variable

$$\hat{x} := (x^0, x) \in \mathbb{R}^{n+1},$$

and define $\hat{f}: \mathbb{R}^{n+1} \times \mathcal{U} \rightarrow \mathbb{R}^{n+1}$ by

$$\hat{f}(\hat{x}, u) := \begin{bmatrix} L(x, u) \\ f(x, u) \end{bmatrix}.$$

Then the optimal control problem (1) is restated as

$$\min_{u(\cdot) \in \mathcal{U}} x^0(t_1)$$

[☆] The material in this paper was presented at the 51st Annual Allerton Conference on Communication, Control, and Computing, October 2–4, 2013, Monticello, IL, USA. This paper was recommended for publication in revised form by Associate Editor Henrik Sandberg under the direction of Editor André Tits.

E-mail address: tomoki@utdallas.edu.




Fig. 1. Terminal cone $C_{\hat{x}_1^*}$, separating hyperplane $\mathcal{H}_{\hat{x}_1^*}$, and costate $\hat{v}^*(t_1^*)$.

subject to the *extended system* (sometimes called the Mayer form; see, e.g., Liberzon (2012, Chapter 4)) on \mathbb{R}^{n+1} defined by

$$\dot{\hat{x}} = \hat{f}(\hat{x}, u) \quad (2)$$

along with the end points $\hat{x}(t_0) = (0, x_0) =: \hat{x}_0$ and $\hat{x}(t_1) = (x^0(t_1), x_1) =: \hat{x}_1$.

2.2. Costate lives in a projective space

Now, let $u^*: [t_0, t_1^*] \rightarrow \mathcal{U}$ be an optimal control and $\hat{x}^*: [t_0, t_1^*] \rightarrow \mathbb{R}^{n+1}$ be the corresponding optimal trajectory of the extended system (2). Combining needle variations and temporal variations at the terminal time t_1^* of the optimal control u^* gives the terminal cone $C_{\hat{x}_1^*} \subset \mathbb{R}^{n+1}$ that approximates the reachable set near the terminal point $\hat{x}_1^* = (x_1^0, x_1^*) := \hat{x}^*(t_1^*)$ (see, e.g., Liberzon (2012, Section 4.2) and Lewis (unpublished, Chapter 5) for details of the construction of the terminal cone).

One then argues that the interior of the cone $C_{\hat{x}_1^*}$ does not intersect $\mathbb{R}_{\le 0} \times T_{\hat{x}_1^*} S_1$, where $\mathbb{R}_{\le 0}$ is the set of non-positive real numbers and $T_{\hat{x}_1^*} S_1$ is the tangent space to S_1 at \hat{x}_1^* ; because if it did then that implies that there exists a variation of the optimal control u^* with the terminal point still in S_1 but with a lower total cost; so $\mathbb{R}_{\le 0} \times T_{\hat{x}_1^*} S_1$ defines “forbidden” directions. As a result, one concludes that there exists a hyperplane $\mathcal{H}_{\hat{x}_1^*} \subset \mathbb{R}^{n+1}$ that separates the interior of $C_{\hat{x}_1^*}$ and $\mathbb{R}_{\le 0} \times T_{\hat{x}_1^*} S_1$ in the sense that they sit on different sides from each other (see Fig. 1).

One then introduces the costate vector $\hat{v}^*(t_1^*) \in (\mathbb{R}^{n+1})^* \setminus \{0\} \cong \mathbb{R}^{n+1} \setminus \{0\}$ as an element such that $\ker \hat{v}^*(t_1^*) = \mathcal{H}_{\hat{x}_1^*}$. However, we observe that $\hat{v}^*(t_1^*)$ is not uniquely defined: we may multiply $\hat{v}^*(t_1^*)$ by any $k \in \mathbb{R} \setminus \{0\}$ and have $\ker(k \hat{v}^*(t_1^*)) = \mathcal{H}_{\hat{x}_1^*}$, i.e., two costate vectors \hat{v}_1 and \hat{v}_2 in \mathbb{R}^{n+1} are equivalent if one is a nonzero constant multiple of the other:

$$\hat{v}_1 \sim \hat{v}_2 \iff \hat{v}_2 = k \hat{v}_1 \quad \text{for some } k \in \mathbb{R} \setminus \{0\}.$$

This defines an equivalence relation and thus we may define the equivalence class $[\hat{v}]$ of an element $\hat{v} \in \mathbb{R}^{n+1}$ by

$$[\hat{v}] := \{\hat{\mu} \in \mathbb{R}^{n+1} \setminus \{0\} \mid \hat{\mu} = k \hat{v} \text{ for some } k \in \mathbb{R} \setminus \{0\}\}.$$

Geometrically, the equivalence class $[\hat{v}]$ corresponds to the straight line along the vector \hat{v} , and the collection of these equivalence classes (straight lines passing through the origin) $[\hat{v}]$ defines the projective space $\mathbb{P}(\mathbb{R}^{n+1}) := (\mathbb{R}^{n+1} \setminus \{0\}) / \sim$. Therefore, the costate is better defined as the straight line along the vector $\hat{v}^*(t_1^*)$ than the vector itself (see Fig. 2), or more mathematically speaking, *the costate is most naturally defined as an element in the projective space $\mathbb{P}(\mathbb{R}^{n+1})$.*

If we need to choose a representative element $\hat{v} = (v_0, v)$ of $[\hat{v}]$, we choose by convention an element $\hat{v} \in \mathbb{R}^{n+1}$ such that $\langle \hat{v}, \hat{w} \rangle \geq 0$ for any $\hat{w} \in \mathbb{R}_{\le 0} \times T_{\hat{x}_1^*} S_1$, where $\langle \cdot, \cdot \rangle$ stands for the standard pairing




Fig. 2. Hyperplane \mathcal{H} and costate $[\hat{v}]$.

of vectors in \mathbb{R}^{n+1} ; that is, \hat{v} and $\mathbb{R}_{\le 0} \times T_{\hat{x}_1^*} S_1$ are on the same side as shown in Fig. 1. In particular, choosing $\hat{w} = -\hat{e}_0 = (-1, 0, \dots, 0)$ gives $v_0 \leq 0$ and $\hat{w} = (0, w)$ with arbitrary $w \in T_{\hat{x}_1^*} S_1$ gives the transversality condition:

$$v^*(t_1^*) \in (T_{\hat{x}_1^*} S_1)^\perp. \quad (3)$$

2.3. Adjoint equation and control Hamiltonian

The costate vector $\hat{v}^*(t_1^*)$ defined above encodes a necessary condition for optimality at the terminal time t_1^* . We then propagate the costate $\hat{v}^*(t_1^*)$ back along the optimal trajectory $\hat{x}^*(t)$ to formulate a necessary condition *along* the trajectory.

This is done by defining the costate vectors $\hat{v}^*: [t_0, t_1^*] \rightarrow \mathbb{R}^{n+1}$ as the solution to the adjoint equation

$$\dot{v}_0^* = 0, \quad \dot{v}_i^* = -v_j^* \frac{\partial f^j}{\partial x^i}(x^*, u^*) - v_0^* \frac{\partial L}{\partial x^i}(x^*, u^*) \quad (4)$$

with the terminal condition $\hat{v}^*(t_1^*)$; where the indices i and j run from 1 to n . The motivation for doing so is that one can relate the costate $\hat{v}^*(t)$ and perturbation $\delta \hat{x}(t)$ at $\hat{x}^*(t)$ with those at the terminal point \hat{x}_1^* thanks to the conservation of the pairing of them, i.e.,

$$\langle \hat{v}^*(t), \delta \hat{x}(t) \rangle = \langle \hat{v}^*(t_1^*), \delta \hat{x}(t_1^*) \rangle. \quad (5)$$

Remark 1. The above conservation is due to the fact that the propagation of perturbation $\delta \hat{x}(t) \mapsto \delta \hat{x}(t_1^*)$ is the tangent lift $T\phi_{t_1^*-t}$ of the flow $\phi_{t_1^*-t}: \hat{x}^*(t) \mapsto \hat{x}^*(t_1^*)$ defined by the optimal solution $\dot{\hat{x}}^* = \hat{f}(\hat{x}^*, u^*)$ whereas the time-reversed propagation of costate $\hat{v}^*(t_1^*) \mapsto \hat{v}^*(t)$ is the cotangent lift $T^*\phi_{t_1^*-t}$ of $\phi_{t_1^*-t}$. In fact, the adjoint equation (4) is nothing but the time derivative of the cotangent lift $T^*\phi_{-t}$ (see, e.g., Agrachev & Sachkov, 2004, Chapter 12).

Recall (see, e.g., Lewis, unpublished; Liberzon, 2012; Pontryagin et al., 1962) also that setting $\delta \hat{x}(t_1^*)$ equal to the variation resulting from a needle variation ending at t in (5) yields the following essential result of the maximum principle: for any $u \in \mathcal{U}$,

$$H_c(\hat{x}^*(t), \hat{v}^*(t), u) \leq H_c(\hat{x}^*(t), \hat{v}^*(t), u^*(t)),$$

with the control Hamiltonian $H_c: \mathbb{R}^{n+1} \times \mathbb{R}^{n+1} \times \mathcal{U} \rightarrow \mathbb{R}$ defined by

$$H_c(\hat{x}, \hat{v}, u) := \langle \hat{v}, \hat{f}(x, u) \rangle = v \cdot f(x, u) + v_0 L(x, u). \quad (6)$$

Therefore, we may define the optimal Hamiltonian $H: \mathbb{R}^{n+1} \times \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ as follows:

$$H(\hat{x}, \hat{v}) := \max_{u \in \mathcal{U}} H_c(\hat{x}, \hat{v}, u) = H_c(\hat{x}, \hat{v}, u^*(\hat{x}, \hat{v})).$$

2.4. Contact control Hamiltonian and normal and abnormal extremals

Notice that the control Hamiltonian (6) is homogeneous of degree 1 in the costate vector $\hat{v} = (v_0, v) \in \mathbb{R}^{n+1}$, i.e., for any $k \in \mathbb{R} \setminus \{0\}$,

$$H_c(\hat{x}, k \hat{v}, u) = k H_c(\hat{x}, \hat{v}, u),$$




Fig. 3. The costates for normal and abnormal minimizers.

and hence, taking the quotient by $\mathbb{R} \setminus \{0\}$, it projects to the *contact control Hamiltonian*

$$h_c: \mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1}) \times \mathcal{U} \rightarrow \mathbb{R}. \quad (7)$$

Likewise, we define the *optimal contact Hamiltonian* $h: \mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1}) \rightarrow \mathbb{R}$ as follows:

$$h(\hat{x}, [\hat{v}]) := \max_{u \in \mathcal{U}} h_c(\hat{x}, [\hat{v}], u). \quad (8)$$

The above definitions of control Hamiltonians are more natural for us because, as shown in Section 2.2, the costate essentially lives in the projective space $\mathbb{P}(\mathbb{R}^{n+1})$.

The contact control Hamiltonian h_c is well-defined globally for any costate in $\mathbb{P}(\mathbb{R}^{n+1})$ regardless of whether the extremal in question is normal or abnormal. However, it turns out that abnormal extremals fall into the coordinate singularity of the natural coordinates for normal extremals: for normal extremals, i.e., if $\nu_0 \neq 0$, then we set $\lambda := -\nu/\nu_0 \in \mathbb{R}^n$. Note that λ is nothing but *homogeneous coordinates* for $\mathbb{P}(\mathbb{R}^{n+1})$: a common way of giving coordinates to $\mathbb{P}(\mathbb{R}^{n+1})$ is to identify $[\hat{v}] = [\nu_0 : \nu_1 : \dots : \nu_n] \in \mathbb{P}(\mathbb{R}^{n+1})$ with $[-\hat{v}/\nu_0]$ and write

$$[\hat{v}] = [-\hat{v}/\nu_0] = [(-1, \lambda)] = [-1 : \lambda_1 : \dots : \lambda_n], \quad (9)$$

where $\{\lambda_i := -\nu_i/\nu_0\}_{i=1}^n$ are coordinates for $[\hat{v}]$. Recall from Section 2.2 that we have $\nu_0 \leq 0$ by convention and so $\nu_0 < 0$ here: the negative sign in the definition of λ makes the above representative element $(-1, \lambda) \in \mathbb{R}^{n+1}$ meet this convention. As a result, we may define the contact control Hamiltonian (7) as follows:

$$h_c(\hat{x}, [\hat{v}], u) = H_c(\hat{x}, -\hat{v}/\nu_0, u) = \lambda \cdot f(x, u) - L(x, u),$$

and the optimal contact Hamiltonian (8) becomes

$$h(\hat{x}, [\hat{v}]) = \lambda \cdot f(x, u^*(x, \lambda)) - L(x, u^*(x, \lambda)). \quad (10)$$

Therefore, *the conventional practice of getting rid of the redundancy in the costate vector $\hat{v} \in \mathbb{R}^{n+1}$ for normal extremals by setting $\nu_0 = -1$ is equivalent to regarding the costate as an element $[\hat{v}]$ in the projective space $\mathbb{P}(\mathbb{R}^{n+1})$ and expressing it in terms of appropriate homogeneous coordinates for $\mathbb{P}(\mathbb{R}^{n+1})$.*

Those lines or costates $[\hat{v}]$ with $\nu_0 = 0$ are at the *coordinate singularity* of the above homogeneous coordinates λ and so one needs to employ a different coordinate chart for such costate $[\hat{v}]$:

$$[\hat{v}] = [0 : \alpha_1 : \dots : \alpha_n],$$

where $\alpha_i = \nu_i/\nu_a$ for some fixed $a \in \{1, \dots, n\}$ such that $\nu_a \neq 0$. So we may now define the contact control Hamiltonian (7) as

$$h_c(\hat{x}, [\hat{v}], u) := H_c(\hat{x}, \hat{v}/\nu_a, u) = \alpha \cdot f(x, u).$$

Therefore, *an abnormal extremal is identified as a costate $[\hat{v}]$ at the coordinate singularity of the standard homogeneous coordinates (9) (see Fig. 3); so the use of the projective space $\mathbb{P}(\mathbb{R}^{n+1})$ gives rise to a differential-geometric classification of normal and abnormal extremals.*




Fig. 4. Hyperplane field propagated along optimal solution.

2.5. Hyperplane field and manifold of contact elements

Recall from Section 2.2 that we introduced the following identification (see Fig. 2):

separating hyperplane $\mathcal{H}_{\hat{x}^*}$ in \mathbb{R}^{n+1}

\Downarrow

costate $[\hat{v}^*(t_1^*)]$ in $\mathbb{P}(\mathbb{R}^{n+1})$ s.t. $\mathcal{H}_{\hat{x}^*} = \ker \hat{v}^*(t_1^*)$.

In fact, *the idea of identifying a hyperplane in \mathbb{R}^{n+1} with an element in the projective space $\mathbb{P}(\mathbb{R}^{n+1})$ is essential in contact geometry.* One may define the set \mathbb{H} of hyperplanes in \mathbb{R}^{n+1} , where we identify those hyperplanes that are translations of one another and so a single representative element in \mathbb{H} would be a hyperplane \mathcal{H} passing through the origin. Then one easily sees that \mathbb{H} is identified with $\mathbb{P}(\mathbb{R}^{n+1})$; therefore,

$$\mathbb{H} := \{\text{hyperplanes in } \mathbb{R}^{n+1}\} \cong \{\text{costates}\} = \mathbb{P}(\mathbb{R}^{n+1}).$$

Now, recall that we employed the adjoint equation (4) to propagate the costate vector $\hat{v}^*(t)$ along the optimal solution $\hat{x}^*(t)$. We may now see the pair $(\hat{x}^*(t), [\hat{v}^*(t)])$ as a curve in $\mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1})$; alternatively, we may define the *hyperplane field* $\mathcal{H}(t) := \ker \hat{v}^*(t)$ along $\hat{x}^*(t)$ and see the pair $(\hat{x}^*(t), \mathcal{H}(t))$ as a curve in $\mathbb{R}^{n+1} \times \mathbb{H}$ (see Fig. 4). They are two different pictures of the same thing. In fact, the identification $\mathbb{R}^{n+1} \times \mathbb{H} \cong \mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1})$ is standard in contact geometry, and they are a simple example of *manifold of contact elements* and is a basic example of *contact manifold* as well.

2.6. Digression on contact geometry

Saving a more general treatment for later (see Section 3.2), this subsection briefly explains what makes the spaces $\mathbb{R}^{n+1} \times \mathbb{H}$ and $\mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1})$ introduced above contact manifolds.

First we define some shorthand notation:

$$\mathcal{C} := \mathbb{R}^{n+1} \times \mathbb{H}, \quad \mathbb{P}(T^*\mathbb{R}^{n+1}) := \mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1}).$$

Since the space \mathbb{H} is the collection of hyperplanes (passing through the origin) in \mathbb{R}^{n+1} , each element in \mathcal{C} may be regarded as the pair $\mathcal{H}_{\hat{x}} := (\hat{x}, \mathcal{H})$ of a base point $\hat{x} \in \mathbb{R}^{n+1}$ and a hyperplane $\mathcal{H} \subset \mathbb{R}^{n+1}$ attached to \hat{x} ; thus the space \mathcal{C} is the collection of all hyperplane fields on \mathbb{R}^{n+1} . Likewise, each element $[\hat{v}]_k := (\hat{x}, [\hat{v}])$ in $\mathbb{P}(T^*\mathbb{R}^{n+1})$ may be regarded as an assignment of a straight line in \mathbb{R}^{n+1} to a base point $\hat{x} \in \mathbb{R}^{n+1}$. We may identify \mathcal{C} with $\mathbb{P}(T^*\mathbb{R}^{n+1})$ by the relation $\mathcal{H} = \ker \hat{v}$ (recall Fig. 2). Using the homogeneous coordinates $\lambda := (\lambda_1, \dots, \lambda_n)$ for $\mathbb{P}(\mathbb{R}^{n+1})$ from (9), we assign coordinates (x^0, x, λ) to an element $(\hat{x}, \mathcal{H}) \in \mathcal{C}$ or $(\hat{x}, [\hat{v}]) \in \mathbb{P}(T^*\mathbb{R}^{n+1})$; thus $\mathcal{C} \cong \mathbb{P}(T^*\mathbb{R}^{n+1})$ is a $(2n + 1)$ -dimensional manifold.

What makes $\mathcal{C} \cong \mathbb{P}(T^*\mathbb{R}^{n+1})$ a *contact manifold* is the one-form θ on it called a *contact form* that is written as

$$\theta_{[\hat{v}]_k} = -dx^0 + \lambda_i dx^i,$$

where \mathbf{d} is the exterior differential. The defining characteristics of the contact form θ is that it defines a hyperplane $\ker \theta|_{[\hat{v}]_{\hat{x}}}$ at each point $[\hat{v}]_{\hat{x}}$ on \mathcal{C} so that the two-form

$$\omega|_{[\hat{v}]_{\hat{x}}} := -\mathbf{d}\theta|_{[\hat{v}]_{\hat{x}}} = \mathbf{d}\mathbf{x}^i \wedge \mathbf{d}\lambda_i$$

is non-degenerate on the hyperplane $\ker \theta|_{[\hat{v}]_{\hat{x}}}$; such a hyperplane assignment is called a *contact structure* which, by definition, makes \mathcal{C} a *contact manifold*; see, e.g., Arnold (1989, Appendix 4), Cannas da Silva (2008, Chapters 10 & 11), Geiges (2008, Chapters 1 & 2) and Kushner, Lychagin, and Rubtsov (2007, Chapter 10) for details.

2.7. Contact Hamiltonian system and necessary condition for optimality

Given a function (contact Hamiltonian) $h: \mathcal{C} \cong \mathbb{P}(T^*\mathbb{R}^{n+1}) \rightarrow \mathbb{R}$, one may define the corresponding *contact Hamiltonian vector field* X_h on $\mathcal{C} \cong \mathbb{P}(T^*\mathbb{R}^{n+1})$ as follows:

$$\theta(X_h) = h, \quad \mathbf{i}_{X_h}\omega = \mathbf{d}h - (\mathbf{d}h \cdot R_\theta)\theta, \quad (11)$$

where R_θ is the *Reeb vector field* associated with the contact form θ , i.e., the vector field R_θ on \mathcal{C} that is uniquely defined by

$$\theta(R_\theta) = 1, \quad \mathbf{i}_{R_\theta}\omega = 0,$$

which gives

$$R_\theta(\hat{x}, \lambda) = -\frac{\partial}{\partial x^0} = -\hat{\mathbf{e}}_0.$$

Hence the Reeb vector field R_θ defines one of the “forbidden” directions in $\mathbb{R}_{\leq 0} \times T_{x^*}S_1$ (see Fig. 1). For a normal extremal, we may use the coordinates (x^0, x, λ) to write (11) as follows:

$$\dot{x}^0 = h - \lambda_i \dot{x}^i, \quad \dot{x}^i = \frac{\partial h}{\partial \lambda_i}, \quad \dot{\lambda}_i = \lambda_i \frac{\partial h}{\partial x^0} - \frac{\partial h}{\partial x^i}.$$

In particular, with the optimal contact Hamiltonian (10), we have

$$\dot{x}^{*0} = L(x^*, u^*), \quad \dot{x}^{*i} = f^i(x^*, u^*),$$

$$\dot{\lambda}_i^* = -\lambda_j^* \frac{\partial f^j}{\partial x^i}(x^*, u^*) + \frac{\partial L}{\partial x^i}(x^*, u^*),$$

which is the extended system (2) along with the adjoint equation (4) for $u = u^*$. A similar result follows for an abnormal extremal as well. To summarize, we have the following:

Proposition 2. *Let $u^*: [t_0, t_1^*] \rightarrow \mathcal{U}$ be an optimal control. Then the corresponding optimal trajectory and costate $(\hat{x}^*(t), [\hat{v}^*](t))$ for $t \in [t_0, t_1^*]$ satisfy the contact Hamiltonian system (11) corresponding to the optimal contact Hamiltonian $h(\cdot, u^*(\cdot)): \mathbb{P}(T^*\mathbb{R}^{n+1}) \rightarrow \mathbb{R}$.*

Remark 3. One may formulate the necessary condition on the cotangent bundle $T^*\mathbb{R}^{n+1}$ as a Hamiltonian system in the symplectic sense. However, this geometric setting is less suited for interpreting the duality between the separating hyperplanes and the costate (discussed in Section 2.5). The symplectic formulation involves the coordinate ν_0 for the costate vector $\hat{v} \in T^*\mathbb{R}^{n+1}$, but this is redundant because the corresponding Hamilton equations give $\dot{\nu}_0 = 0$; see (4). The contact-geometric view gets rid of this redundancy at the outset by projectivization; as a result, it naturally gives rise to the identification of hyperplanes and projective cotangent spaces. This is exactly the duality between the separating hyperplanes and the costate exploited in the maximum principle, as explained in Sections 2.2 and 2.5.

3. Geometry of optimal control on manifolds

We now replace the state space \mathbb{R}^n by an n -dimensional manifold M to consider optimal control of systems on the manifold M . We do not delve into the details on how to extend the proof of the maximum principle to manifolds; see, e.g., Agrachev and

Sachkov (2004, Chapter 12), Barbero-Liñán and Muñoz-Lecanda (2009), Chang (2011) and Jurdjevic (1997, Chapter 11). Skipping technical details, we briefly sketch how the geometric ideas in \mathbb{R}^n from the previous section can be generalized to manifolds. Our main references in contact geometry are Arnold (1989, Appendix 4), Cannas da Silva (2008, Chapters 10 & 11), Geiges (2008, Chapters 1 & 2) and Kushner et al. (2007, Chapter 10).

3.1. Optimal control on manifolds

Let $\tau: TM \rightarrow M$ be the tangent bundle of M , \mathcal{U} a compact subset of \mathbb{R}^m as before, and $pr: M \times \mathcal{U} \rightarrow M$ the projection to the first slot. We now have a map $f: M \times \mathcal{U} \rightarrow TM$ such that $\tau \circ f = pr$ and a cost function $L: M \times \mathcal{U} \rightarrow \mathbb{R}$, and then we can formulate the optimal control problem on M just as in (1).

Then we define the extended configuration space $\hat{M} := \mathbb{R} \times M = \{\hat{x} := (x^0, x)\}$, and also $\hat{f}: \hat{M} \times \mathcal{U} \rightarrow \mathbb{R} \times TM$ by $\hat{f}(\hat{x}, u) := (L(x, u), f(x, u))$; hence we may define the extended system (2).

The cone $C_{\hat{x}^*}$ is now more naturally a subset of the tangent space $T_{\hat{x}^*}\hat{M}$ and the costate vector $\hat{v}^*(t_1^*)$ is in the cotangent space $T_{\hat{x}^*}^*\hat{M}$. Likewise, the hyperplane $\mathcal{H}(t)$ and costate vector $\hat{v}^*(t)$ at $\hat{x}^*(t)$ are in the tangent and cotangent spaces, respectively, i.e., $\mathcal{H}(t) \subset T_{\hat{x}^*(t)}\hat{M}$ and $\hat{v}^*(t) \in T_{\hat{x}^*(t)}^*\hat{M}$. As briefly mentioned in Remark 1, one needs to propagate back the costate vector $\hat{v}^*(t_1^*)$ along $\hat{x}^*(t)$ by the cotangent lift of the flow $\phi_{t_1^*-t}: \hat{M} \rightarrow \hat{M}$ defined by the optimal solution $\hat{x}^* = \hat{f}(\hat{x}^*, u^*)$, i.e., $\hat{v}^*(t) := T^*\phi_{t_1^*-t}(\hat{v}^*(t_1^*))$, and hence

$$[\hat{v}^*(t)] = [T^*\phi_{t_1^*-t}(\hat{v}^*(t_1^*))] \in \mathbb{P}(T_{\hat{x}^*(t)}^*\hat{M}),$$

where $\mathbb{P}(T_{\hat{x}^*(t)}^*\hat{M})$ is the projectivization of the cotangent space $T_{\hat{x}^*(t)}^*\hat{M}$.

3.2. Contact geometry and necessary condition

The space $\mathcal{C} := \mathbb{R}^{n+1} \times \mathbb{H}$ introduced in Section 2.6 is our prototype of what is called a *manifold of contact elements*, which we define now: let \hat{M} be a manifold and $T\hat{M}$ its tangent bundle. A *contact element* on \hat{M} is a point $\hat{x} \in \hat{M}$ along with a hyperplane (passing through the origin) $\mathcal{H}_{\hat{x}} \subset T_{\hat{x}}\hat{M}$. The collection \mathcal{C} of contact elements $(\hat{x}, \mathcal{H}_{\hat{x}})$ is called a *manifold of contact elements* of \hat{M} . Note that an element in the manifold \mathcal{C} is a point $\hat{x} \in \hat{M}$ along with a hyperplane $\mathcal{H}_{\hat{x}} \subset T_{\hat{x}}\hat{M}$ attached to the point \hat{x} .

Likewise, the space $\mathbb{P}(T^*\mathbb{R}^{n+1}) := \mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1})$ from Section 2.6 is an example of the *projectivized cotangent bundle* $\mathbb{P}(T^*\hat{M})$ defined by

$$\mathbb{P}(T^*\hat{M}) := \bigcup_{\hat{x} \in \hat{M}} \mathbb{P}(T_{\hat{x}}^*\hat{M}).$$

Then the hyperplane $\mathcal{H}_{\hat{x}} \subset T_{\hat{x}}\hat{M}$ is identified with $[\hat{v}]_{\hat{x}} \in \mathbb{P}(T_{\hat{x}}^*\hat{M})$, and thus \mathcal{C} is identified with the projectivized cotangent bundle $\mathbb{P}(T^*\hat{M})$. Note that then $[\hat{v}^*(t)]$ and $\mathcal{H}(t) := \ker \hat{v}^*(t)$ are curves in $\mathbb{P}(T^*\hat{M})$ and \mathcal{C} , respectively. An element $[\hat{v}]_{\hat{x}}$ in each fiber $\mathbb{P}(T_{\hat{x}}^*\hat{M})$ is parametrized by the homogeneous coordinates $\{\lambda_i := -\nu_i/\nu_0\}_{i=1}^n$ defined in (9) for normal extremals. Therefore, $(x^0, \dots, x^n, \lambda_1, \dots, \lambda_n)$ gives local coordinates for $\mathbb{P}(T^*\hat{M})$ just as with the case with $\mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1})$; in fact, for $\hat{M} = \mathbb{R}^{n+1}$, $\mathbb{P}(T^*\hat{M}) \cong \mathbb{R}^{n+1} \times \mathbb{P}(\mathbb{R}^{n+1})$.

Let $\pi: T^*\hat{M} \rightarrow \hat{M}$ be the cotangent bundle projection and $[\pi]: \mathbb{P}(T^*\hat{M}) \rightarrow \hat{M}$ be its projectivization, i.e., $[\pi]([\hat{v}]_{\hat{x}}) = \hat{x}$ for any $[\hat{v}]_{\hat{x}} := (\hat{x}, [\hat{v}]) \in \mathbb{P}(T^*\hat{M})$. Now, we define a one-form θ on $\mathbb{P}(T^*\hat{M})$ as follows:

$$\theta|_{[\hat{v}]_{\hat{x}}} \cdot w|_{[\hat{v}]_{\hat{x}}} = \hat{v}_{\hat{x}} \cdot T|_{[\hat{v}]_{\hat{x}}}[\pi](w|_{[\hat{v}]_{\hat{x}}})$$

for any $w_{[\hat{v}]_x} \in T_{[\hat{v}]_x} \mathbb{P}(T^* \hat{M})$. We also define a two-form ω on $\mathbb{P}(T^* \hat{M})$ by $\omega := -\mathbf{d}\theta$. Locally, we have

$$\theta_{[\hat{v}]_x} = -\mathbf{d}x^0 + \lambda_1 \mathbf{d}x^1 + \cdots + \lambda_n \mathbf{d}x^n$$

and

$$\omega_{[\hat{v}]_x} = \mathbf{d}x^1 \wedge \mathbf{d}\lambda_1 + \cdots + \mathbf{d}x^n \wedge \mathbf{d}\lambda_n.$$

The one-form θ then defines the hyperplane $\ker \theta_{[\hat{v}]_x}$ at every point $[\hat{v}]_x$ of $\mathbb{P}(T^* \hat{M})$ so that ω is non-degenerate on the hyperplane, i.e., the hyperplane field $\ker \theta$ defines a *contact structure* on $\mathbb{P}(T^* \hat{M})$, which makes itself a *contact manifold*.

We may then define contact Hamiltonian systems just as in Section 2.7 and can generalize Proposition 2 to control systems on the manifold M , where the optimal contact Hamiltonian h is defined on $\mathbb{P}(T^* \hat{M})$.

4. Application: terminal cost and transversality condition

4.1. Optimal control to a general target with terminal cost

Let $K: \mathbb{R}^n \rightarrow \mathbb{R}$ be a smooth function defined on the configuration space $M = \mathbb{R}^n$ and consider the following variant of the optimal control problem (1) with the terminal cost $K(x(t_1))$:

$$\min_{u(\cdot) \in \mathcal{U}} \left[K(x(t_1)) + \int_{t_0}^{t_1} L(x(t), u(t)) dt \right]$$

$$\text{subject to } \dot{x} = f(x, u), \quad x(t_0) = x_0, \quad x(t_1) \in S_1,$$

where the conditions for the end points are the same as before. One may now define

$$x^0(t) := K(x(t)) + \int_{t_0}^t L(x(s), u(s)) ds$$

with the initial condition $x^0(t_0) = K(x(t_0))$. Then we have

$$\frac{d}{dt} [x^0(t) - K(x(t))] = L(x(t), u(t)),$$

or defining $\hat{y} = (y^0, y) := (x^0 - K(x), x)$, we have

$$\dot{y}^0 = L(y, u), \quad \dot{y} = f(y, u) \tag{12}$$

with the initial condition $y^0(t_0) = 0$; thus we have the same extended system as in (2). Therefore, we may apply Proposition 2 to the extended system (12), and so the optimal flow is given by the contact Hamiltonian flow X_h with the contact Hamiltonian $h: \mathbb{P}(T^* \mathbb{R}^{n+1}) \rightarrow \mathbb{R}$ defined by

$$h(\hat{y}, \mu) := \mu_j f^j(y, u^*) - L(y, u^*),$$

where $[(-1, \mu)] \in \mathbb{P}(T^* \mathbb{R}^{n+1})$ is the costate corresponding to y .

4.2. Transversality condition via contact transformation

One may be then tempted to write $\mu^*(t_1^*) \in (T_{x_1^*} S_1)^\perp$ as the transversality condition as in (3), but this is incorrect because $[(-1, \mu)] \in \mathbb{P}(T^* \mathbb{R}^{n+1})$ is the costate corresponding to the state variables $(y^0, y) \in \mathbb{R}^{n+1}$ and

$$y^0(t_1) = \int_{t_0}^{t_1} L(x(t), u(t)) dt$$

is *not* the quantity to be minimized. Instead, it is $x^0(t_1)$ that is to be minimized, and therefore we may write the *correct* transversality condition

$$\lambda^*(t_1^*) \in (T_{x_1^*} S_1)^\perp \tag{13}$$

for the costate $[(-1, \lambda)] \in \mathbb{P}(T^* \mathbb{R}^{n+1})$ corresponding to the original variable x .

The question is then: how does one rewrite (13) in terms of μ ? Contact geometry provides a simple and elegant answer to this question and leads us to a simple derivation of the correct transversality condition for μ : the discussion in the previous subsection motivates us to define the diffeomorphism $\Phi_K: \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n+1}$ defined by

$$\Phi_K(x^0, x) := (x^0 - K(x), x) = (y^0, y).$$

Clearly its inverse is given by $\Phi_K^{-1} = \Phi_{-K}$. Now let $\mathcal{H}_{\hat{x}_1^*}$ be the separating hyperplane at \hat{x}_1^* . Then $T_{\hat{x}_1^*} \Phi_K(\mathcal{H}_{\hat{x}_1^*})$ gives the separating hyperplane at \hat{y}_1^* , and we have

$$0 = \langle \hat{v}_{\hat{x}_1^*}, \mathcal{H}_{\hat{x}_1^*} \rangle = \langle T_{\hat{x}_1^*} \Phi_K^{-1}(\hat{v}_{\hat{x}_1^*}), T_{\hat{x}_1^*} \Phi_K(\mathcal{H}_{\hat{x}_1^*}) \rangle.$$

So $T_{\hat{x}_1^*} \Phi_K^{-1}(\hat{v}_{\hat{x}_1^*})$ is a costate vector in $T_{\hat{y}_1^*} \mathbb{R}^{n+1}$ and hence

$$[(-1, \mu^*(t_1^*))] = [T_{\hat{x}_1^*} \Phi_K^{-1}(\hat{v}_{\hat{x}_1^*})],$$

where $[\hat{v}_{\hat{x}_1^*}] = [(-1, \lambda^*(t_1^*))]$. As a result, the costates $[(-1, \lambda)]$ and $[(-1, \mu)]$ are related by the projectivization

$$\Psi_K := [T^* \Phi_K]: \mathbb{P}(T^* \mathbb{R}^{n+1}) \rightarrow \mathbb{P}(T^* \mathbb{R}^{n+1})$$

of $T^* \Phi_K: T^* \mathbb{R}^{n+1} \rightarrow T^* \mathbb{R}^{n+1}$; specifically,

$$(\hat{x}, \lambda) = \Psi_K(\hat{y}, \mu),$$

i.e., the diagram below commutes.

$$\begin{array}{ccc} \mathbb{P}(T^* \mathbb{R}^{n+1}) & \xleftarrow{\Psi_K} & \mathbb{P}(T^* \mathbb{R}^{n+1}) & & (\hat{x}, \lambda) & \longleftarrow & (\hat{y}, \mu) \\ \downarrow [\pi] & & \downarrow [\pi] & & \downarrow & & \downarrow \\ \mathbb{R}^{n+1} & \xrightarrow{\Phi_K} & \mathbb{R}^{n+1} & & \hat{x} & \longleftarrow & \hat{y} \end{array}$$

Therefore, we have

$$(x^0, x, \lambda) = \Psi_K(y^0, y, \mu) = (y^0 + K(y), y, \mu + \mathbf{d}K(y)),$$

and then applying (13) to the above equation gives the *correct* transversality condition for μ :

$$\mu^*(t_1^*) + \mathbf{d}K(x_1^*) \in (T_{x_1^*} S_1)^\perp.$$

Remark 4. One may, in principle, work with costate vectors in $T^* \mathbb{R}^{n+1}$ without projectivization, but as mentioned in Remark 3, the result comes with a redundancy in the costate vector, i.e., $\mu_0 = \lambda_0$ but λ_0 is left arbitrary; whereas the projectivization gets rid of the redundancy at the outset.

References

Agrachev, A. A., & Sachkov, Y. L. (2004). *Control theory from the geometric viewpoint*. Springer.

Arnold, V. I. (1989). *Mathematical methods of classical mechanics*. Springer.

Barbero-Liñán, M., & Muñoz-Lecanda, M. C. (2009). Geometric approach to Pontryagin's maximum principle. *Acta Applicandae Mathematicae*, 108(2), 429–485.

Cannas da Silva, A. (2008). *Lectures on symplectic geometry*. New York: Springer.

Chang, D. E. (2011). A simple proof of the Pontryagin maximum principle on manifolds. *Automatica*, 47(3), 630–633.

Geiges, H. (2008). *An introduction to contact topology*. Cambridge University Press.

Jurdjevic, V. (1997). *Geometric control theory*. Cambridge: Cambridge University Press.

Kushner, A., Lychagin, V. V., & Rubtsov, V. (2007). *Contact geometry and non-linear differential equations*. Cambridge, UK: Cambridge University Press.

Lewis, A.D. The maximum principle of Pontryagin in control and in optimal control. (unpublished).

Liberzon, D. (2012). *Calculus of variations and optimal control theory*. Princeton, NJ: Princeton University Press.

Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., & Mishchenko, E. F. (1962). *The mathematical theory of optimal processes*. John Wiley & Sons, Inc.