5 Lecture 5. Proof of the maximum principle

5.1 The tent method

We continue with the static nonlinear optimization problem:

$$\begin{aligned} &\min \ g_0(x) \\ &\text{subject to } g_i(x) \le 0, \quad i = 1, \cdots, m \end{aligned} \tag{LM}$$

in which $\{g_i\}_{i=0}^m \in C^1(\mathbb{R}^n; \mathbb{R})$. Suppose that the problem is feasible, i.e., there exists an admissible x_* which minimizes $g_0(x)$.

Recall that we defined the following sets:

$$\Omega_i = \{x \in \mathbb{R}^n : g_i(x) \le 0\}, \quad i = 1, \cdots, m$$

and for $x_1 \in \mathbb{R}^n$, let

$$\Omega_0 = \{x : g_0(x) < g_0(x_1)\} \cup \{x_1\}.$$

We have shown that x_1 is a minimizer if and only if the set

$$\Sigma := \Omega_0 \cap \Omega_1 \cap \cdots \cap \Omega_m$$

is $\{x_1\}$.

As an example, let $m = 1$ and Figure 1 shows two sets Ω_0 and Ω_1 on a plane. In this figure, Ω_1 and Ω_0 intersect only at the point x_1. So x_1 is a minimizer.

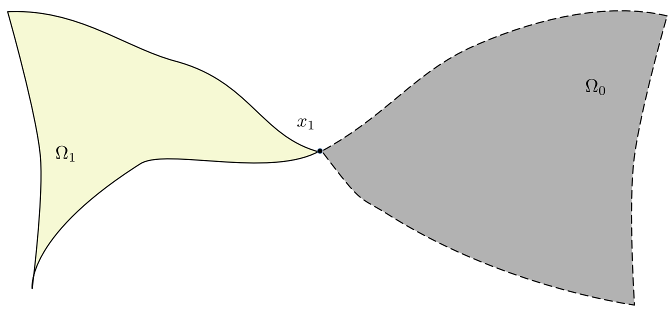

Figure 1: Separating 2-dim tents.

However, in Figure 2, the intersection of the two sets contains also some other points. Thus x_1 is not a minimizer.

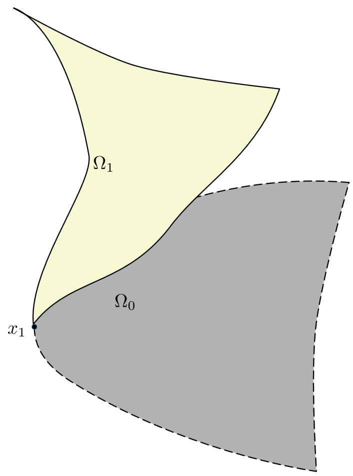

Figure 2: Separating 2-dim tents.

Heuristically, in order that Ω_0 and Ω_1 intersect at only one point, we shall be able to separate them at the intersecting point. That is, there exists a hyperplane passing through the point x_1 so that Ω_0 and Ω_1 lie on different sides of the hyperplane. But Ω_0 and Ω_1 are some nonlinear sets which are not easy to work with. However, we can linearize them! The linearization is achieved by taking their tangent cones. Given a set $\Omega \subseteq \mathbb{R}^n$ (may be non-convex), the *tangent cone* at $x \in \Omega$ is defined as

$$T_x\Omega := \left\{ v \in \mathbb{R}^n \;\middle|\; \begin{array}{l} \exists \{x_i\}_{i=1}^{\infty} \subseteq \Omega,\ \exists \{t_i\}_{i=1}^{\top} \subseteq \mathbb{R}_{>0}, \text{ s.t.} \\ t_i \downarrow 0,\ x_i \to x,\ \text{ and } (x_i - x)/t_i \to v \end{array} \right\}$$

see Figure 3. In practice, you choose a curve $\gamma(t)$ in Ω passing through x_0, then if γ is differentiable, $\gamma'(0)$ should lie in $T_x\Omega$.

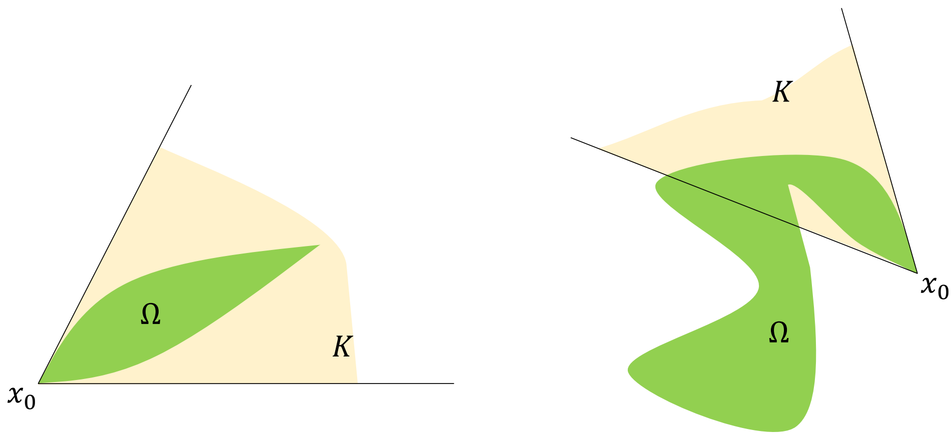

Figure 3: Tangent cones of convex and non-convex sets Ω.

Example 1. Let $f : \mathbb{R}^n \to \mathbb{R}$ be a C^1 function. Consider the set

$$M = \{x : f(x) = 0\}.$$

Choose $\gamma(t)$ in M, s.t. $\gamma(0) = x_0$, then

$$\lim_{t \to 0+} \frac{\gamma(t) - \gamma(0)}{t} \cdot \nabla f(\gamma(t)) = \frac{df(\gamma(t))}{dt} = 0.$$

Thus if $\nabla f(x_0) \neq 0$, the tangent cone at x_0 of M is the orthogonal complement of $\nabla f(x_0)$, i.e., the tangent space of M. More generally, when Ω is a smooth sub-manifold – think of a smooth surface in $\mathbb{R}^n$ – then $T_x\Omega$ is the tangent space of Ω at x. Hence the notation coincides with tangent space in the smooth case. Now consider the set

$$N = \{x : f(x) \leq 0\},$$

If $x_0 \in \partial N$, i.e., $f(x_0) = 0$, choose $\gamma(t) \in N$ with $\gamma(0) = x_0$,

$$\lim_{t \to 0+} \frac{\gamma(t) - \gamma(0)}{t} \cdot \nabla f(\gamma(t)) = \lim_{t \to 0+} \frac{f(\gamma(t)) - f(\gamma(0))}{t} \leq 0.$$

Thus $T_{x_0}N = \{x : (x - x_0)^\top \nabla f(x_0) \leq 0\}$. That is, $T_{x_0}N$ is a half space.

Sometimes in practice, it is hard to compute the tangent cone of a given set, e.g., the tangent cone of the reachable set as we will soon see. As a compromise, we may compute some subcones of the tangent cone. The most useful ones are those called *tents.*

Definition 1 (Tent)**.** Given a set Ω and the closure of the tangent cone $T_x\Omega$ at x, a *tent* is a convex cone $K \subseteq T_x\Omega$ with apex x.

Notice that this definition of a tent is not the original one, see [1], which is simplified to avoid technical complications. But at the end of the day, this definition will be sufficient for our purpose.

There are two reasons why we need the notion of tent instead of a tangent cone. The first is already mentioned – the tangent cone is sometimes difficult to calculate. The second reason is that a tangent cone of a set may be non-convex and that non-convex objects are hard to work with. In Figure 4, K_0 represent the tangent cones while K_1 some tents.

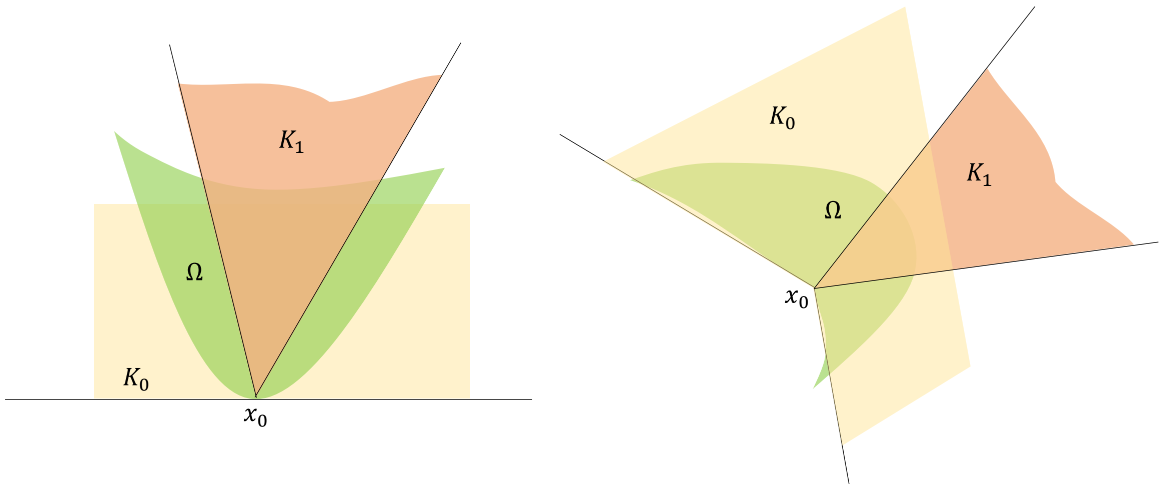

Figure 4: Tents.

Example 2. A tent of the set $M = \{g(x) < g(x_1)\} \cup \{x_1\}$ at x_1 is $\{x : (x - x_1)^\top \nabla g(x_1) \leq 0\}$.

Now let's come back to our nonlinear optimization problem. Intuitively, to be able to "separate" Ω_0 and Ω_1, the tents of the two sets at the intersecting point should also be separable in the sense that they intersect only at the apex. Or equivalently, there is a hyperplane passing through x_1 which separates the tents of Ω_0 and Ω_1, see Figure 5.

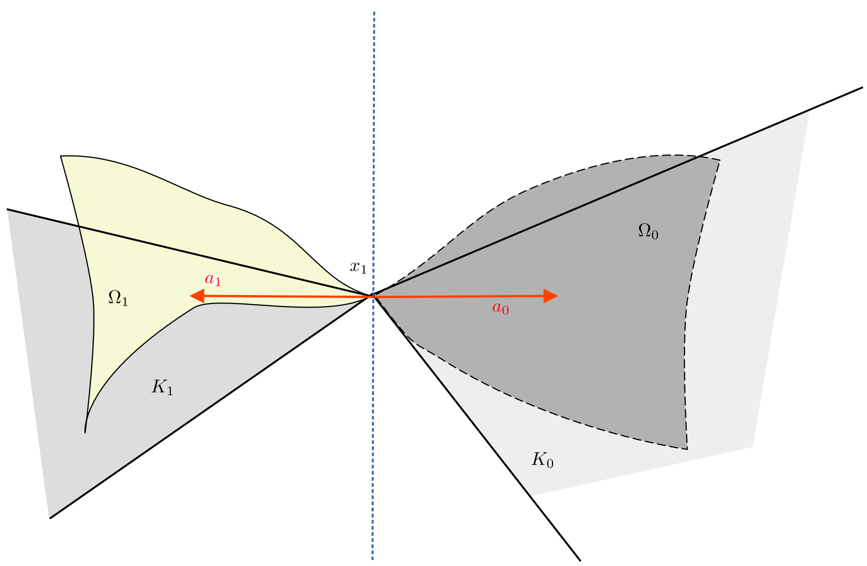

Figure 5: Separating 2-dim tents.

In Figure 5, let us choose two arbitrary nonzero vectors a_0 and a_1 perpendicular to the separating hyperplane such that

$$a_0 + a_1 = 0 \tag{1}$$

Then we notice that

$$a_i^\top (x - x_1) \geq 0, \quad \forall x \in K_i,\ i = 0, 1. \tag{2}$$

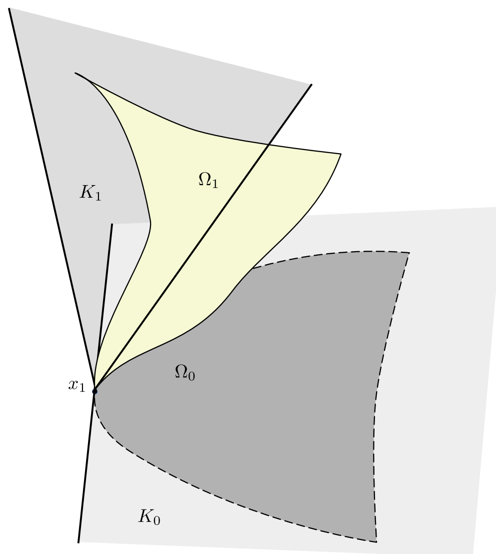

Figure 6: Tents not separable.

Thus if we can find K_0 and K_1, we can obtain a necessary condition based on the relation (2). For problem (LM), this is easy since g_0 and g_1 are smooth and from the above example with know that

$$K_i = \{x : \nabla g_i(x_1)(x - x_1) \leq 0\}, \quad i = 0, 1.$$

(It is easy to see that $K_0 = T_{x_1}\{x : g_0(x) \leq g_0(x_1)\}$) That is, K_i are half spaces, see Figure 7.

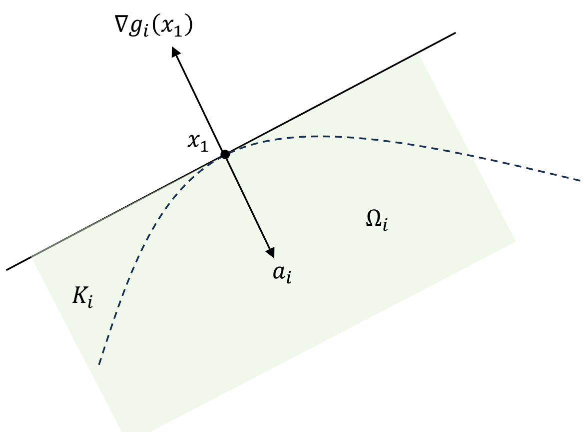

Figure 7: The tents are half spaces.

Therefore, a_i must be of the form

$$a_i = \lambda_i \nabla g_i(x_1)$$

for $\lambda_i \leq 0$. Since λ_i cannot be zero at the same time, $\lambda_i < 0$ for $i = 0, 1$. Thus the relation (1) becomes

$$\lambda_0 \nabla g_0(x_1) + \lambda_1 \nabla g_1(x_1) = 0$$

for $\lambda_0, \lambda_1 < 0$ or

$$\nabla g_0(x_1) + \lambda \nabla g_1(x_1) = 0$$

for some $\lambda > 0$. This is a special case of the famous KKT (Karush-Kuhn-Tucker) condition which we will be able to prove once we have generalize the above reasoning.

Exercise 1. Derive the KKT condition for the problem

$$\min g_0(x)$$

subject to

$$g_i(x) \leq 0$$
$$h_j(x) = 0.$$

To solve this exercise, we need to generalize our previous discussions to arbitrary finite many tents. The first thing to generalize is separability. Suppose we have three sets $\Omega_0, \Omega_1, \Omega_2$ or three cones K_0, K_1, K_2, how should we define the separability of them? The correct answer is the following one.

Definition 2 (Separability)**.** Let $K_0, \cdots, K_p$ be some closed, convex cones with a common apex x in $\mathbb{R}^n$. They are said to be *separable* if there exists a hyperplane Γ through x that separates one of the cones from the intersection of the others.

Lemma 1. *Let $K_0, \cdots, K_p$ be some closed, convex cones with a common apex x in $\mathbb{R}^n$. Then they are separable if and only if there exist vectors $a_i,\ i = 0, 1, \cdots, p$ fulfilling*[1]

$$a_i^\top (y - x) \leq 0, \quad \forall y \in K_i$$

and at least one of which is not zero and such that

$$a_0 + \cdots + a_p = 0.$$

Lemma 2. *Let $\Omega_0, \cdots, \Omega_p$ be sets in $\mathbb{R}^n$ satisfying*

$$\Omega_0 \cap \cdots \cap \Omega_p = \{x\},$$

and $K_0, \cdots, K_p$ be tents of these sets at x. If at least one of the tents is distinct from its affine hull. Then $K_0, \cdots, K_p$ is separable.

The proofs of the above two results are quite non-trivial. Interested readers are referred to [1]. Now we have all the ingredients to prove the maximum principle.

5.2 Optimal Control

5.2.1 Problem reduction

We start by recalling the problem formulation of optimal control under fixed terminal time. We focus on time-invariant control systems (as we mentioned before, time varying case can be turned into time invariant one):

$$\dot{x} = f(x, u), \tag{3}$$

where $x(t) \in \mathbb{R}^n,\ u(t) \in U \subset \mathbb{R}^m$ for all $t \in [0, t_f]$, the initial condition $x(0) = x_0$ is assumed to be fixed. The cost function is

$$J(u(\cdot)) = \varphi(x(t_f)) + \int_0^{t_f} L(x(s), u(s))ds,$$

where $\varphi(\cdot),\ f(\cdot, u),\ L(\cdot, u)$ are continuously differentiable for all u. The optimal control problem amounts to finding a process $u_*(t),\ x_*(t),\ 0 \leq t \leq t_f$, with a (measurable) controller $u_*(t)$ such that $x_*(t_f) \in M$ for some manifold M, and $J(u_*(\cdot))$ attains a minimum. We say that the problem is in 1) *Mayer form* if $L = 0$; 2) *Lagrange form* if $\varphi = 0$; 3) *Bolza form* if neither L nor φ is zero.

We claim that the preceding three types of optimal control problems can all be reduced to Mayer form. In fact, let

$$x_{n+1}(t) = \int_0^t L(x(s), u(s))ds$$

[1]Note that we can also use $a_i^\top (y - x) \geq 0$ by reversing the sign of a_i, see (2).

the system becomes

$$\begin{cases} \dot{x} = f(x,u) \\ \dot{x}_{n+1} = L(x,u) \end{cases} \tag{4}$$

with initial condition $(x_0, 0)$, and the cost function becomes

$$J = \varphi(x(t_f)) + x_{n+1}(t_f). \tag{5}$$

This is an optimal control problem of the Mayer form of a time-invariant system. Due to this reason, it suffices to study the optimal control problem with cost function:

$$J = \varphi(x(t_f)).$$

Introduce the following notations which will be used in the proof:

$$\begin{aligned} & x_1 := x_*(t_f) \\ & \Omega_0 = \{x_1\} \cup \{x : \varphi(x) < \varphi(x_1)\} \\ & \Omega_1 : \text{ reachability set from } x_0 \\ & \Omega_2 = M\text{: the terminal manifold} \end{aligned}$$

Let $u_*(t)$, $x_*(t)$, $0 \le t \le t_f$ be an optimal process. Then it is easily seen that

$$\Omega_0 \cap \Omega_1 \cap \Omega_2 = \{x_1\}. \tag{6}$$

The reader should immediately realize that such type of condition implies separability of tents of the three sets, this is the content of Lemma 2. Denote K_i the tent of Ω_i at x_1. It thus remains to find the tents K_i. The tents K_0 and K_2 can be easily computed:

$$\begin{aligned} & K_0 = \{x \in \mathbb{R}^n : \nabla\varphi(x_1)(x - x_1) \le 0\} \\ & K_2 = T_{x_1}\Omega_2 \end{aligned}$$

(note that Ω_2 is a fixed manifold).

Therefore, our problem boils down to calculating the tent of Ω_1 at x_1: K_1. By definition, a tent is only a convex subcone of the tangent cone of Ω_1 at x_0, we should however, try to find a tent as big as possible, since the bigger the tent, the more necessary information it conveys. This is the main non-trivial step in proving the maximum principle (if we already know Lemma 1, 2). This tent was found by Boltyanskii, via the so called *needle variation*.

5.2.2 Needle variation

Suppose at the moment that the optimal control $u_* : [0, t_f] \to U$ is continuous. Fix $\tau \in (0, t_f]$ and consider the following *needle shaped variation* of u_* for small $\varepsilon > 0$:

$$u_\varepsilon(t) = \begin{cases} w, & t \in (\tau - \varepsilon, \tau] \\ u_*(t), & \text{otherwise} \end{cases}$$

where $w \in U$ is some constant, see Figure 8.

Denote $t \mapsto x_\varepsilon(t)$ the solution to $\dot{x} = f(x, u_\varepsilon)$, i.e.,

$$\dot{x}_\epsilon(t) = \begin{cases} f(x_*(t), u_*(t), & t \in [0, \tau - \epsilon] \\ f(x_\epsilon(t), w), & t \in (\tau - \epsilon, \tau] \\ f(x_\epsilon(t), u_*(t)), & t \in (\tau, t_f]. \end{cases}$$

Obviously, $u_\epsilon(\cdot)$ is admissible, thus $x_\epsilon(t_f)$ belongs to the reachable set at t_f, i.e., $x_\epsilon(t_f) \in \Omega_1$ for all ϵ chosen above. Thus by definition, $\left.\frac{\partial x_\varepsilon(t_f)}{\partial \varepsilon}\right|_{\varepsilon=0+}$ must belong to the tangent cone of Ω_1. Denote

$$v(t) = \left.\frac{\partial x_\varepsilon(t)}{\partial \varepsilon}\right|_{\varepsilon=0+}, \quad t \in [\tau, t_f]$$

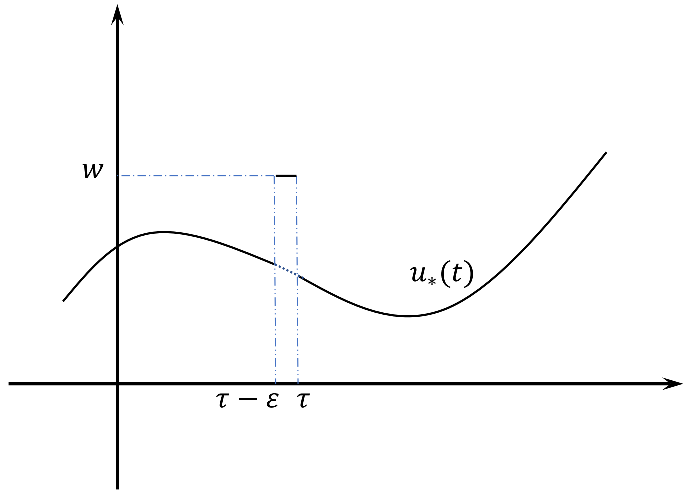

Figure 8: Needle variation.

then it remains to find $v(t_f)$. We call $v(t_f)$ a *deviation vector*. To find the deviation vector, first we need to characterize $x_\epsilon(t)$. Denote $v_\epsilon(t) = \frac{\partial x_\varepsilon(t)}{\partial \varepsilon}$, since $u_\epsilon(t) = u_*(t)$ for $t \geq \tau$, it follows that

$$\begin{aligned}\frac{dv_\epsilon(t)}{dt} &= \frac{\partial}{\partial \epsilon} f(x_\epsilon(t), u_*(t)) = \frac{\partial f}{\partial x}(x_\epsilon(t), u_*(t)) \frac{\partial x_\epsilon(t)}{\partial t} \\ &= \frac{\partial f}{\partial x}(x_\epsilon(t), u_*(t)) v_\epsilon(t), \quad \forall t \in (\tau, t_f]\end{aligned}$$

Evaluating at $\epsilon = 0+$, we get $\dot{v}(t) = \frac{\partial f}{\partial x}(x_*(t), u_*(t))v(t)$. That is, $v(t)$ satisfies a linear ODE. It still remains to find the initial condition $v(\tau)$. Note that

$$\begin{aligned}x_\varepsilon(\tau) &= x_*(\tau - \epsilon) + \int_{\tau-\epsilon}^{\tau} f(x_\epsilon(s), w)\mathrm{d}s, \\ &= x_*(\tau - \epsilon) + \int_{\tau-\epsilon}^{\tau} f(x_*(s), u_*(s))\mathrm{d}s + \int_{\tau-\epsilon}^{\tau} [f(x_\epsilon(s), w) - f(x_*(s), u_*(s))]\mathrm{d}s \\ &= x_*(\tau) + \int_{\tau-\epsilon}^{\tau} [f(x_\epsilon(s), w) - f(x_*(s), u_*(s))]\mathrm{d}s\end{aligned}$$

thus

$$\begin{aligned}v(\tau) &= \lim_{\varepsilon \to 0+} \frac{x_\varepsilon(\tau) - x_*(\tau)}{\varepsilon} \\ &= \lim_{\varepsilon \to 0+} \frac{1}{\varepsilon} \left[\int_{\tau-\varepsilon}^{\tau} f(x_\varepsilon(t), w)dt - f(x_*(t), u_*(t))\mathrm{d}t \right] \\ &= f(x_*(\tau), w) - f(x_*(\tau), u_*(\tau)).\end{aligned} \tag{7}$$

To summarize, $v(\cdot)$ is the solution to the following Cauchy problem

$$\begin{cases} \dot{v} = \dfrac{\partial f}{\partial x}(x_*(t), u_*(t))v, \quad \forall t \in [\tau, t_f] \\ v(\tau) = f(x_*(\tau), w) - f(x_*(\tau), u_*(\tau)). \end{cases}$$

To construct more deviation vectors, let $v_1(t_f), \cdots, v_r(t_f)$ be some different deviation vectors obtained as above corresponding to some distinct time instants $\tau_1 < \cdots < \tau_r$ and constant inputs $w_1, \cdots, w_r$. Consider the combined needle variation

$$u_{\varepsilon,k}(t) = \begin{cases} w_i, & t \in (\tau_i - k_i\varepsilon, \tau_i] \text{ for some } i \in \{1, \cdots, r\} \\ u_*(t), & \text{otherwise} \end{cases}$$

where k_i are non-negative constants. One can show that

$$\sum_{i=1}^{r} k_i v_i(t_f) = \left.\frac{\partial x(t_f, u_{\varepsilon,k})}{\partial \varepsilon}\right|_{\varepsilon=0+}$$

which implies that $\sum_{i=1}^{r} k_i v_i(t_f)$ are again in $T_{x_1}\Omega_1$. Still call these vectors deviation vectors and define K_1 to be the set of all deviation vectors, i.e.,

$$K_1 = \left\{ \sum_{i=1}^{r} k_i v_i(t_f) \;\middle|\; \begin{array}{l} \exists r \in \mathbb{Z}_+,\ \tau_i \in [0, t_f),\ w_i \in U,\ k_i \geq 0, \\ v_i(t_f) \text{ the deviation vector obtained from needle} \\ \text{variation at} \tau_i \text{ with spike } w_i \end{array} \right\}$$

Then K_1 is a tent of Ω_1 at x_1.

So we have obtained a tent of the reachable set at x_1, although it is somehow abstract. We will see next how to use the expression of K_1.

5.2.3 Final step: the costate equation and the maximum principle

Condition (6) implies that K_0, K_1, K_2 are separable. Invoking Lemma 1 and Lemma 2, we deduce that there exist three vectors a_i, at least one of which is nonzero, satisfying

$$a_i^\top v \leq 0, \quad v \in K_i,\ i = 0, 1, 2 \tag{8}$$

and

$$a_0 + a_1 + a_2 = 0. \tag{9}$$

In particular, $a_1^\top v(t_f) \leq 0$ for any deviation vector $v(t_f)$. Now we introduce a small trick: if we are able to construct some function $p : [0, t_f] \to \mathbb{R}^n$ such that $p(t)^\top v(t) \equiv$ constant with $p(t_f) = a_1$, then we obtain immediately $p(t)^\top v(t) = a_1^\top v(t_f) \leq 0$ for all $t \in [0, t_f]$. In other words, we propagate the inequality at the end point to the previous time instants. In particular, if v is the deviation vector obtained by needle variation at time τ with spike w, then $v(\tau) = f(x_*(\tau), w) - f(x_*(\tau), u_*(\tau)$. Thus at $t = \tau$, $p(\tau)^\top [f(x_*(\tau), w) - f(x_*(\tau), u_*(\tau))] \leq 0$ or

$$p(\tau)^\top f(x_*(\tau), u_*(\tau)) \geq p(\tau)^\top f(x_*(\tau), w)) \tag{10}$$

For convenience, define

$$H(x, u, p) := p^\top f(x, u)$$

which is the Hamiltonian associated with the system . Now that the spike can be any $w \in U$ and $t \in [0, t_f)$, it follows from (10) that

$$H(x_*(t), u_*(t), p(t)) = \max_{u \in U} H(x_*(t), u, p(t)) = \text{constant}, \quad \forall t \in [0, t_f). \tag{11}$$

This is the maximum principle that we have been looking for! Except two things: the interval $[0, t_f)$ doesn't include the endpoint t_f and the function p hasn't been determined yet. The first issue can be fixed if everything is continuous in the above formula, which is indeed true as long as we have shown p is, since f, x_* and u_* are continuous as assumed. For the second issue, let us recall the following simple fact:

Lemma 3. *Consider two linear ODE*

$$\begin{aligned} \dot{x} &= A(t)x \\ \dot{p} &= -A(t)^\top p \end{aligned}$$

where $x, p \in \mathbb{R}^n$. Then $p(t)^\top x(t) = p(t')^\top x(t')$ for any $t, t' \in \mathbb{R}$.

With this lemma, we can now construct p to be the solution of the following ODE

$$\begin{aligned} \dot{p} &= -\left[\frac{\partial f}{\partial x}(x_*(t), u_*(t))\right]^\top p \\ &= -H_x^\top(x_*, u_*, p) \end{aligned} \tag{12}$$

with terminal state $p(t_f) = a_1$ (note that this is exactly the costate equation).

Recall that

$$K_0 = \{x \in \mathbb{R}^n : \nabla\varphi(x_1)(x - x_1) \leq 0\}$$
$$K_2 = T_{x_1}\Omega_2$$

For a_0, since K_0 is a half space, $a_0^\top v \leq 0$ for $v \in K_0$ implies $a_0 = \lambda\nabla\varphi(x_1)^\top$ for some constant $\lambda \geq 0$. For a_2, since K_2 is a sub-manifold, $a_2 \perp K_2$. It follows from (9) that (recall $a_1 = p(t_f)$):

$$\lambda\nabla\varphi(x_*(t_f))^\top + p(t_f) \perp T_{x_*(t_f)}M \tag{13}$$

for some constant $\lambda \geq 0$. If $\lambda > 0$, then it is equivalent to $p(t_f) + \nabla\varphi(x_*(t_f)) \perp T_{x_*(t_f)}M$ by changing a_1 to λa_1. As in man textbooks, we ignore the pathological case $\lambda = 0$.

Up to now, we have prove the maximum principle for the Mayer problem under the assumption that u_* is continuous. For u not continuous, only the condition (11) needs to be modified by noticing that the limits in (7) exist for almost all $t \in [0, t_f]$. Summarizing, we have proved the following.

Theorem 1. *Suppose that the Mayer form optimal control problem admits a piecese-continuous optimal law $u_*(\cdot)$ with corresponding trajectory $x_*(\cdot)$. Then there is a solution to the costate equation (12), such that the triple $(x_*(t), u_*(t), p(t))$ satisfies the maximum principle (11) for almost all t (all t on the interval on which $u_*(\cdot)$ is continuous) and the transversality condition (13).*

We have so far considered the optimal control problem under the condition that t_f is fixed. It can be easily extended to the case of free terminal time: it is obvious that all the necessary conditions of Theorem 1 still need to hold. The mere difference is that now one can also make the variation of the terminal time. For example, consider a needle variation at τ, let $v(t_f)$ be the corresponding deviation vector. Fix some $\mu > 0$, since $x_\epsilon(t_f + \epsilon\mu) \in \Omega_1$, $\left.\frac{\partial x_\varepsilon(t_f+\varepsilon\mu)}{\partial\varepsilon}\right|_{0+}$ must also lie in the tangent cone of Ω_1, but

$$\left.\frac{\partial x_\varepsilon(t_f + \varepsilon\mu)}{\partial\varepsilon}\right|_{\varepsilon=0+} = \left.\frac{\partial x_\varepsilon(t_f)}{\partial\varepsilon}\right|_{\varepsilon=0+} + \left.\frac{\partial x_*(t_f + \epsilon\mu)}{\partial\epsilon}\right|_{\varepsilon=0+} = v(t_f) + \mu f(x_*(t_f), u_*(t_f))$$

Thus we can construct another tent of Ω_1 at x_1 as

$$K_1' = \{v(t_f) + \mu f(x_*(t_f), u_*(t_f)) : v(t_f) \in K_1,\ \mu \in \mathbb{R}\}.$$

It follows that one can obtain a finer condition than (11):

$$H(x_*(t), u_*(t), p(t)) = \max_{u\in U} H(x_*(t), u, p(t)) = 0, \quad \forall t \in [0, t_f).$$

Indeed, take $v(t_f) = 0$ (no needle variation), then $a_1^\top(\mu f(x_*(t_f), u_*(t_f))) \leq 0$ for any $\mu \in \mathbb{R}$ implies that $a_1^\top f(x_*(t_f), u_*(t_f)) = 0$.

Let us use Theorem 1 to derive the maximum principle for Bolza form. Recall that the system model and cost function of the Bolza problem can be equally written as (4) and (5). Suppose that the terminal manifold $\Omega_2 = M$, then for the augmented system (4), the terminal manifold is $\tilde{\Omega}_2 = \Omega_2 \times \mathbb{R}$. The Hamiltonian becomes

$$H(x, u, p, p_0) = p^\top f(x, u) + p_0 L(x, u)$$

and the costate equation still reads $\dot{p} = -H_x^\top$, and $\dot{p}_0 = 0$ since H doesn't depend on x_{n+1}. Thus p_0 is a constant. The transversal condition reads

$$\lambda\begin{bmatrix}\nabla\varphi(x_*(t_f))^\top \\ 1\end{bmatrix} + \begin{bmatrix}p(t_f) \\ p_0\end{bmatrix} \perp T_{x_1}\Omega_2 \times \mathbb{R}$$

for some $\lambda \geq 0$, from which it follows that $p_0 = -\lambda \leq 0$ and

$$p(t_f) - p_0\nabla\varphi(x_*(t_f))^\top \perp T_{x_*(t_f)}\Omega_2$$

When p_0 is nonzero, one can take $p_0 = -1$, when p_0 is zero, then $p(t_f) \perp T_{x_1}\Omega_2$. Thus we are done with the general Bolza form problem.

For u not continuous, only the condition (11) needs to be modified by noticing that the limits in (7) exist for almost all $t \in [0, t_f]$. Summarizing, we have proved the following.

Theorem 2. *Consider the system* $\dot{x} = f(x, u)$ *with cost function*

$$J(u) = \varphi(x(t_f)) + \int_0^{t_f} L(x, u)\mathrm{d}t$$

and boundary constraint

$$x(t_f) \in M \subseteq \mathbb{R}^n$$

Assume f, φ *and* L *are* C^1 *in* x. *Let* $(x^*(\cdot), u^*(\cdot))$ *correspond to the optimal solution to the minimization problem*

$$\min_{u \in \mathcal{U}_{\mathrm{ad}}} J(u)$$

in which $\mathcal{U}_{\mathrm{ad}} = \{u : [0, t_f] \to U \subseteq \mathbb{R}^m\}$. *Define the Hamiltonian function*

$$H(x, u, p, p_0) = p^\top f(x, u) + p_0 L(x, u)$$

Then there exists a function $p^* : [0, t_f] \to \mathbb{R}^n$ *and a constant* $p_0^* \in \{0, -1\}$, *satisfying* $(p_0^*, p^*(t)) \not\equiv (0, 0)$ *such that*

1) $(x^*(\cdot), p^*(\cdot))$ *satisfies the canonical equation*

$$\begin{aligned} \dot{x} &= H_p^\top \\ \dot{p} &= -H_x^\top \end{aligned}$$

with initial condition $x^*(0) = x_0$. *The second equation is called the costate equation, and* p *is the costate.*

2) The transversality condition holds:

$$p^*(t_f) - p_0^* \varphi_x^\top(x^*(t_f)) \perp T_{x^*(t_f)} M.$$

3) The maximum principle holds:

$$H(x^*(t), u^*(t), p^*(t), p_0^*) = \max_{u \in U \subseteq \mathbb{R}^m} H(x^*(t), u, p^*(t), p_0^*) = \text{constant} \tag{14}$$

for all $t \in [0, t_f]$. *In particular, this constant is zero if* t_f *is free.*

References

[1] Vladimir Grigorevich Boltyanskii. The method of tents in the theory of extremal problems. *Russian Mathematical Surveys*, 30(3):1, 1975.