# Chapter 4

# The Maximum Principle

## 4.1 The statement of the Maximum Principle

For the Maximum Principle , A-F is a primary reference. Supplementary references: the original book by Pontryagin et al.; Sussmann's lecture notes (more rigorous and modern treatment).

$\longrightarrow$ All standing assumptions on the general control problem still hold, we'll just make them more specific.

We will state and prove the Maximum Principle for two special problems, and later discuss how all other cases can be deduced from this.

### 4.1.1 Special Problem 1

System: $\dot{x} = f(x, u)$, $x(t_0) = x_0$ (given), no $t$ dependence, $f$ is $\mathcal{C}^1$ in $x$ (no differentiability with respect to $u$!), $f$ and $f_x$ are $\mathcal{C}^0$ in $u$ (these are conditions for existence and uniqueness we had earlier). $x \in \mathbb{R}^n$, $u \in \mathcal{U}$—some subset of $\mathbb{R}^m$. Can be the whole $\mathbb{R}^m$. Usually closed (Sussmann assumes this, but A-F and Pontryagin don't).

$L = L(x, u)$, no $t$ dependence. Assume that $L$ satisfies the same regularity assumptions as $f$.

$K = 0$ —no terminal cost.

Target set: $S = \mathbb{R} \times \{x_1\}$—fixed-endpoint, free-time problem.

$J(u) = \int_{t_0}^{t_1} L(x, u)dt$, where $t_1$ is the *first time* at which $x(t) = x_1$.

**Theorem 4 (Maximum principle for Special Problem 1 [Theorem 5-5P in [A-F, p. 305])** *Let $u^*(t)$ be an optimal control and let $x^*(t)$ be the corresponding trajectory. Then there exists a function $p^*(t)$ and a constant $p_0^* \leq 0$, satisfying $(p_0^*, p^*(t)) \neq (0, 0)\ \forall\, t$, such that:*

*1) $x^*(t)$ and $p^*(t)$ satisfy the canonical equations*

$$\dot{x}^* = H_p(x^*, u^*, p^*, p_0^*)$$
$$\dot{p}^* = -H_x(x^*, u^*, p^*, p_0^*)$$

*with boundary conditions $x^*(t_0) = x_0$, $x^*(t_1) = x_1$, where the Hamiltonian is defined as*

$$H(x, u, p, p_0) := \langle p, f(x, u)\rangle + p_0 L(x, u) \qquad \forall\, u \in \mathcal{U},\ \forall\, t \in [t_0, t_1]$$

*2)* $H(x^*(t), u^*(t), p^*(t), p_0^*) \geq H(x^*(t), u, p^*(t), p_0^*)\ \forall\, u \in \mathcal{U},\ \forall\, t \in [t_0, t_1]$*, i.e.,* $u^*$ *is a global maximum of* $H(x^*(t), \cdot, p^*(t))\ \forall\, t$.

*3)* $H(x^*(t), u^*(t), p^*(t), p_0^*) \equiv 0,\ t \in [t_0, t_1]$.

Comments:

- $u^*$ is *optimal*—by this we mean that for any other control $u$ taking values in $\mathcal{U}$ which transfers $x$ from $x_0$ at $t = t_0$ to $x_1$ at some (unspecified) time $t = t_1$,

$$J(u^*) \leq J(u)$$

  (We could relax the assumption of *global* minimum to a *local* minimum, but need to know what topology to use, so we won't do that. The proof will make the topology more clear.)

- $p_0$ is the *abnormal multiplier*. We saw its analog in calculus of variations . It handles pathological cases. If $p_0 \neq 0$, then we can always normalize it to $p_0 = -1$ and forget it, which is what one does when applying the Maximum Principle (almost always). For this reason, we will also suppress the argument $p_0$ of $H$ from now on.

- We saw in the variational approach (just a little while ago) that $H = \mathit{const}$, but $H \equiv 0$ may seem surprising. This is a feature of the *free-time problem* (earlier we were looking at the fixed-time case). More on this later.

How restrictive is Special Problem 1?

- Time-independence—no problem: if time-dependent, can always set $x_{n+1} = t$ and eliminate $t$ (but we need $f$ to be $C^1$ with respect to $t$);

- No terminal cost—no problem: if we have terminal cost, can do

$$K(t_f, x_f) = K(t_0, x_0) + \int_{t_0}^{t_1} (K_t + K_x f(t, x, u))\, dt$$

  and define $\hat{L} := L + K_t + K_x f$.

- $S = \mathbb{R} \times \{x_1\}$—this is not very general, need more general target set.

### 4.1.2 Special Problem 2

Same as Special Problem 1, *except* $S = \mathbb{R} \times S_1$, where $S_1$ is a $k$-dimensional surface in $\mathbb{R}^n$, $1 \leq k \leq n$. This is defined as follows:

$$S_1 = \{x \in \mathbb{R}^n : h_1(x) = h_2(x) = \cdots = h_{n-k}(x) = 0\}$$

Here, $h_i$ are smooth functions from $\mathbb{R}^n$ to $\mathbb{R}$. We also assume that each $x \in S_1$ is a regular point, in the sense discussed before: the gradients $(h_i)_x$ are linearly independent at $x$.

The special case $k = n$ means $S_1 = \mathbb{R}^n$.

Special Problem 1 can be thought of as a special case with $k = 0$.

Does the case $S_1 = \mathbb{R}^n$ (i.e., both $x_f$ and $t_f$ are free) make sense? When do we stop? Why even start moving?

Answer: yes, because $L$ may be negative. (Also if we have $K$, which translates to the same thing: moving decreases the cost.)

$\longrightarrow$ When we say "cost", we may think implicitly that $L \geq 0$, but it is not necessarily true.

The difference between the next theorem and the previous one is only in the boundary conditions for the canonical system.

**Theorem 5 (Maximum principle for Special Problem 2 [Theorem 5-6P in [A-F, p. 306])** *Same as the Maximum Principle for Special Problem 1, except:*

*The final boundary condition in the canonical equations is*

$$x^*(t_1) \in S_1$$

*and there is one more necessary condition:*

*4) ("transversality condition")*
*The vector $p^*(t_1)$ is orthogonal to the tangent space to $S_1$ at $x^*(t_1)$:*

$$\langle p^*(t_1), d \rangle = 0 \qquad \forall\, d \in T_{x^*(t_1)} S_1$$

We know this means that $p^*(t_1)$ is a linear combination of $(h_i)_x(x^*(t_1))$, since

$$T_{x^*(t_1)} S_1 = \{d \in \mathbb{R}^n : \langle (h_i)_x(x^*(t_1)), d \rangle = 0,\ i = 1, 2, \dots, n-k\}$$

If $k = n$, then $S_1 = \mathbb{R}^n$. In this case, the transversality condition says (or can be extended to say)

$$p^*(t_1) = 0$$

($p^*$ must be orthogonal to all $d \in \mathbb{R}^n$, since there are no $h_i$).

Note that we still have $n$ boundary conditions at $t = t_1$: $k$ degrees of freedom for $x^*(t_1) \in S_1$ correspond to $n - k$ equations, and $n - k$ degrees of freedom for $p^*(t_1) \perp S_1$ correspond to $k$ equations/constraints.

$\longrightarrow$ The freer the state, the less free the costate. Each additional degree of freedom for $x^*(t_1)$ kills one degree of freedom for $p^*(t_1)$.

## 4.2 The proof of the Maximum Principle

We now come to the culmination of the course!

Proof outline:

1) Lagrange $\mapsto$ Mayer;
2) Principle of optimality;
3) Temporal variations of $u^*$;
4) Spatial variations of $u^*$;
5) Variational equation;
6) Terminal cone;
7) Key topological lemma;
8) Separating hyperplane;
9) Adjoint equation;
10) Proof for Special Problem 1;
11) Proof for Special Problem 2.

We'll explain all concepts and logical steps, but will not necessarily hammer out all $\varepsilon$ and $\delta$.

$$\dot{x} = f(x,u), \qquad x(t_0) = x_0$$

$$J = \int_{t_0}^{t_1} L(x,u)dt, \qquad (t_1, x(t_1)) \in S = \mathbb{R} \times \{x_1\} \text{ (first)}$$

### 4.2.1 Lagrange $\mapsto$ Mayer

Introduce an additional state:

$$\dot{x}^0 = L(x,u), \qquad x^0(t_0) = 0$$

Then $J = \int_{t_0}^{t_1} \dot{x}^0 dt = x^0(t_1)$ —terminal cost.

Note: A-F use $x_0$, and their initial state is $\mathbf{x_0}$ (which is hard to do on the board). Pontryagin et al. use $x^0$, but they use superscripts everywhere.

The composite system is

$$\begin{aligned}\dot{x}^0 &= L(x,u)\\ \dot{x} &= f(x,u)\end{aligned} \tag{4.1}$$

and the initial condition is $\binom{0}{x_0}$.

Target set: if we are dealing with Special Problem 1, then $x(t_1) = x_1$, and the target set for the augmented system (4.1) is

$$\mathbb{R} \times \left\{ \begin{pmatrix} * \\ x_1 \end{pmatrix}, * - \text{arbitrary real number} \right\} = \mathbb{R} \times \mathbb{R} \times \{x_1\}$$

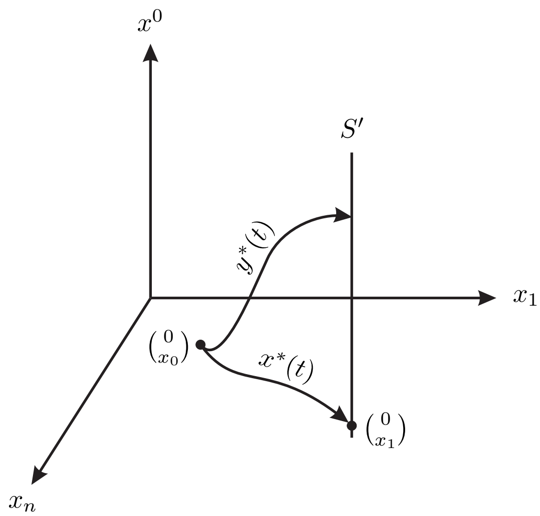


Figure 4.1: Optimal trajectory of the augmented system

Sometimes we write the augmented system (4.1) as

$$\dot{y} = g(y,u), \qquad y \in \mathbb{R}^{n+1} \tag{4.2}$$

where

$$g = \begin{pmatrix} L \\ f \end{pmatrix} \in \mathbb{R}^{n+1}$$

In the figure, $x^*(t)$ is the projection of $y^*(t)$ onto $\mathbb{R}^n$.

### 4.2.2 Principle of optimality

(this should be familiar from ECE 515)

Geometrically, optimality of $x^*(t)$ (or, what is the same, of $y^*(t)$) means that no other trajectory $y(t)$ of the system (4.2) hits $S'$ below $y^*(t)$, starting from $(0, x_0)$.

$\longrightarrow$ Related to the *reachable set* of (4.2): the final point must be on the boundary of the reachable set (see, e.g., [Agrachev-Sachkov, p. 134]). See figure 4.2. (We'll actually work with attainability cone—first-order approximation of the reachable set.)

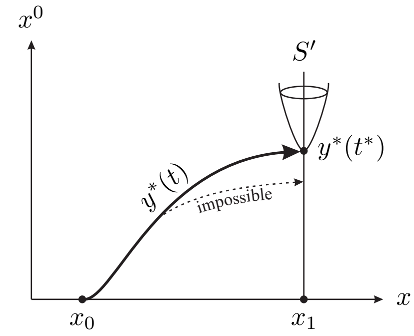


Figure 4.2: Principle of optimality

Consider again the optimal trajectory $y^*(t)$. For simplicity, draw $x$ in one dimension.

$\longrightarrow$ From now on we let $t^*$ be the terminal time of the optimal trajectory.

For all $t$, no other trajectory starting from $y^*(t) = (x^{0,*}(t), x^*(t))$ can hit $S'$ below $y^*(t^*)$.

**Exercise 12** [A] (due Mar 14) Prove the following stronger claim: if $y^*(t)$ is optimal, and if we take any $t_1$ and $t_2$ and let $S''$ be the vertical line passing through $x^*(t_2)$, then no trajectory starting at $y^*(t_1)$ can hit $S''$ below the value $x^{0,*}(t_2)$, *at any time* (even $\neq t_2$).

We will come back and use the above observation after we generate a (sufficiently rich) family of control variations and resulting trajectory perturbations. The fact that all these perturbed trajectories must hit $S'$ higher than $y^*$ (in an approximate sense) will lead to the desired necessary conditions.

### 4.2.3 Temporal variations of $u^*$

Let $\tau \in \mathbb{R}$ be arbitrary, and let $\varepsilon > 0$ be small.

$$u_\tau(t) := u^*(\min\{t, t^*\}), \qquad t \in [t_0, t^* + \varepsilon\tau]$$

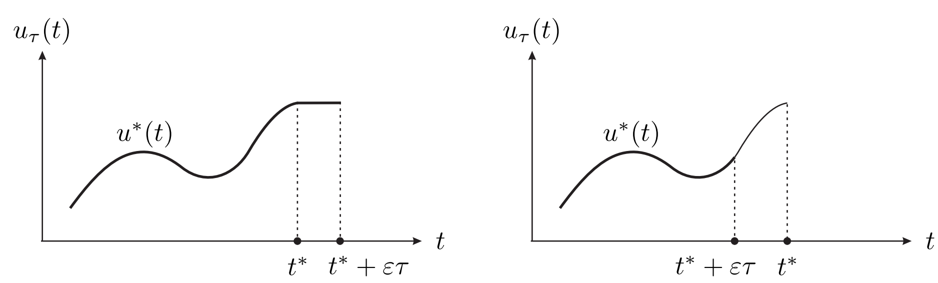


Figure 4.3: Temporal variation

The resulting solution will have, at the new final time $t^* + \varepsilon\tau$, the value

$$\begin{aligned} y(t^* + \varepsilon\tau) &= y^*(t^*) + \dot{y}(t^*)\varepsilon\tau + o(\varepsilon) \\ \text{(recall)} &= y^*(t^*) + g(y^*(t^*), u^*(t^*))\varepsilon\tau + o(\varepsilon) \\ &= y^*(t^*) + \begin{pmatrix} L(x^*(t^*), u^*(t^*)) \\ f(x^*(t^*), u^*(t^*)) \end{pmatrix} \varepsilon\tau + o(\varepsilon) \\ &=: y^*(t^*) + \delta(\tau) + o(\varepsilon) \end{aligned}$$

(note that $x^*(t^*) = x_1$). The vector $\delta(\tau)$ describes the infinitesimal (up to first order in $\varepsilon$) perturbation resulting from the temporal control variation.

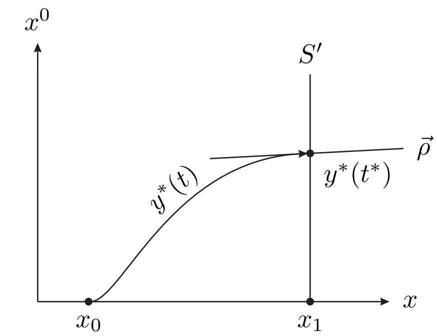


Figure 4.4: The effect of temporal variations

But $\delta(\tau)$ depends linearly on $\tau$, and $\tau$ can be positive or negative, so we get the entire line—call it $\vec{\rho}$. Any point on $\vec{\rho}$ corresponds to $u_\tau$ for some $\tau$. (Note that $\varepsilon\tau$ appear only as a product, so larger $\tau$ corresponds to larger perturbation for fixed $\varepsilon$ but it's only valid in the limit as $\varepsilon \to 0$ anyway.)

$\longrightarrow$ What really matters is *direction*, not magnitude.

### 4.2.4 Spatial variations of $u^*$

("needle" variations, or Pontryagin-McShane variations)

These are essentially the same as variations we saw in the proof of Weierstrass' E-function condition (that proof is due to McShane, 1939).

Let $w$ be an arbitrary element of the control set $\mathcal{U}$. Let $a > 0$ be arbitrary, $\varepsilon > 0$ be small. Consider the interval $I := (b - \varepsilon a, b] \subset [t_0, t^*)$ where $b \neq t^*$ is a point of continuity[1] of $u^*$.

$$u_{w,I}(t) := \begin{cases} u^*(t) & \text{if } t \notin I \\ w & \text{if } t \in I \end{cases}$$

The picture illustrates this control variation and the resulting trajectory perturbation:

We will deviate on $I$, after that will "run parallel" to $y^*$. Let's characterize this formally. First, we need to describe the perturbation obtained at $t = b$ (after that, we'll propagate it up to $t^*$).

[1]The reason for this is that the subsequent Taylor expansions assume that $y$ is $\mathcal{C}^1$ on $I$.

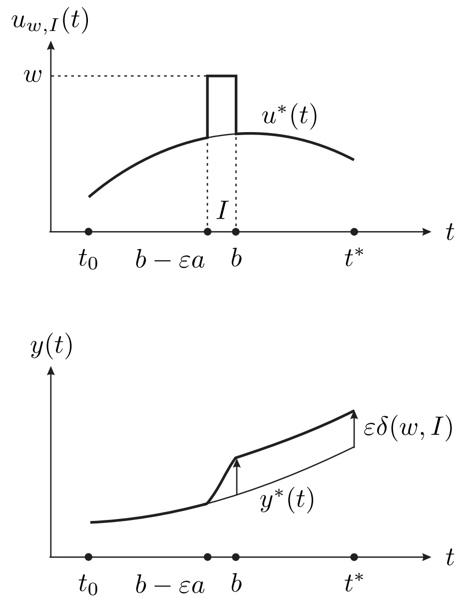


Figure 4.5: Spatial control variation and resulting trajectory perturbation

First, the first-order Taylor expansion of $y^*$ around $t = b$ gives

$$y^*(b - \varepsilon a) - y^*(b) \approx -\dot{y}^*(b)\varepsilon a \tag{4.3}$$

where here and below, $\approx$ stands for equality up to $o(\varepsilon)$. This implies

$$y^*(b) \approx y^*(b - \varepsilon a) + \dot{y}^*(b)\varepsilon a = y^*(b - \varepsilon a) + g(y^*(b), u^*(b))\varepsilon a \tag{4.4}$$

Second, the first-order Taylor expansion of $y$ around $t = b - \varepsilon a$ gives (using the fact that $y(b - \varepsilon a) = y^*(b - \varepsilon a)$ by construction)

$$\begin{aligned} y(b) &\approx y(b - \varepsilon a) + \dot{y}(b - \varepsilon a)\varepsilon a = y^*(b - \varepsilon a) + g(y^*(b - \varepsilon a), w)\varepsilon a \\ &= y^*(b - \varepsilon a) + g(y^*(b), w)\varepsilon a + (g(y^*(b - \varepsilon a), w) - g(y^*(b), w))\varepsilon a \end{aligned} \tag{4.5}$$

For the last term, we have using Taylor again

$$(g(y^*(b - \varepsilon a), w) - g(y^*(b), w))\varepsilon a \approx g_y(y^*(b), w)(y^*(b - \varepsilon a) - y^*(b))\varepsilon a$$

and in view of (4.3) we see that this is $o(\varepsilon^2)$. Thus, from (4.5) we obtain

$$y(b) \approx y^*(b - \varepsilon a) + g(y^*(b), w)\varepsilon a$$

Combining this with (4.4), we can write

$$y(b) \approx y^*(b) + \nu_b(w)\varepsilon a$$

where

$$\nu_b(w) := g(y^*(b), w) - g(y^*(b), u^*(b))$$

Or, writing this out in more detail,

$$x^0(b) = x^{0,*}(b) + (L(x^*(b), w) - L(x^*(b), u^*(b)))\,\varepsilon a + o(\varepsilon)$$

and

$$x(b) = x^*(b) + (f(x^*(b), w) - f(x^*(b), u^*(b)))\,\varepsilon a + o(\varepsilon)$$

### 4.2.5 The variational equation

What happens after $t = b$?

In the variational approach, we derived that a control perturbation $u = u^* + \alpha\eta$ and the corresponding state perturbation $x = x^* + \alpha\xi$ are related by the differential equation (3.3). There, we studied the effect of control perturbation on the trajectory with the *same initial condition*: $\xi(t_0) = 0$. However, we can also study in a similar way the effect of the *same control* on a *nearby trajectory*. Consider a perturbed trajectory

$$y = y^* + \varepsilon\varphi$$

where $\varphi$ is a perturbation. Suppose that both $y$ and $y^*$ have the same control $u^*$. Then we have

$$\dot{y} = \dot{y}^* + \varepsilon\dot{\varphi} \tag{4.6}$$

The left-hand side of (4.6) can be written as

$$\dot{y} = g(y, u^*) = g(y^*, u^*) + g_y|_* (y - y^*) + o(\varepsilon) = g(y^*, u^*) + g_y|_* \varepsilon\varphi + o(\varepsilon)$$

Since $\dot{y}^* = g(y^*, u^*)$, (4.6) gives

$$\dot{\varphi} = g_y|_* \varphi + \frac{o(\varepsilon)}{\varepsilon}$$

In fact, it is equivalent (and more convenient) to work with a perturbation $\psi$ satisfying just (no h.o.t. )

$$\dot{\psi} = g_y|_* \psi \tag{4.7}$$

with the same initial condition as $\varphi$.

*Claim:* In terms of $\psi$, the perturbed trajectory $y$ satisfies

$$y = y^* + \varepsilon\psi + o(\varepsilon) \tag{4.8}$$

**Exercise 12** [B] (due Mar 14) Prove this claim.

⟶ The equation (4.7) is sometimes called the *variational equation*. We can write it in the form

$$\dot{\psi} = A_*(t)\psi$$

where $A_*(t) = g_y|_*$. It's just linearization around $y^*(\cdot)$.

Recall that $g(y,u) = (L(x,u), f(x,u))$. So letting $\psi = (\xi^0, \xi)$ and writing out the variational equation in components, we can check that

$$\dot{\xi}^0 = (L_x)^T\Big|_* \xi$$
$$\dot{\xi} = f_x|_* \xi$$

(note that $L_{x^0}, f_{x^0} = 0$). In other words,

$$A_*(t) = \begin{pmatrix} 0 & (L_x)^T\Big|_* \\ 0 & f_x|_* \end{pmatrix}$$

(here $(L_x)^T\Big|_*$ is a row vector and $f_x|_*$ is an $n \times n$ matrix) and the perturbed trajectory of the augmented system (4.1) is described by (4.8).

$\psi$ describes how spatial perturbations propagate with time.

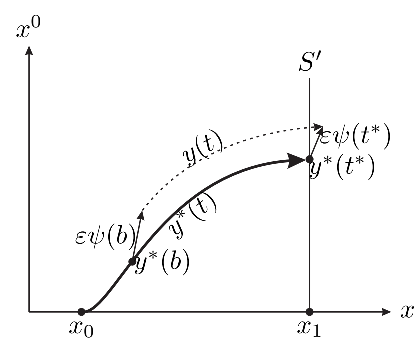


Figure 4.6: Propagation of a spatial variation

Specifically, we are interested in propagating the perturbation

$$\psi(b) = \nu_b(w)a$$

obtained in the previous subsection from time $t = b$ up to the final time $t = t^*$. The formula is $\psi(t^*) = \Phi_*(t^*, b)\psi(b)$, where $\Phi_*$ is the transition matrix for $A_*$.

### 4.2.6 Terminal cone

In the previous step we showed that the perturbation $\psi(b) = \nu_b(w)a$ will propagate according to $\dot{\psi} = A_*\psi$, so at time $t^*$ we will have

$$y(t^*) = y^*(t^*) + \varepsilon\Phi_*(t^*, b)\nu_b(w)a + o(\varepsilon)$$

where $\Phi_*$ is the transition matrix for $A_*$. Letting $\delta(w,I) := \Phi_*(t^*,b)\nu_b(w)a$, we write

$$y(t^*) = y^*(t^*) + \varepsilon\delta(w,I) + o(\varepsilon)$$

$\longrightarrow$ Note: $\delta(w,I)$ depends only on $w$, $a$, and $b$, but not on $\varepsilon$. Its contribution is scaled by multiplication by $\varepsilon$. And, tracing the formulas, we see that the *direction* is independent of $a$.

This vector $\delta(w,I)\varepsilon$ is what we had in Figure 4.6 as $\varepsilon\psi(t^*)$.

So, let $\vec{\rho}(w,b)$ be the ray in this direction originating at $y^*(t^*)$ (note that we only have a semi-line and not a line: $a, \varepsilon > 0$).

$\longrightarrow$ This is the *infinitesimal* perturbation direction which characterizes the effect of using $u = w$ on an infinitesimal interval prior to $t = b$.

For different choices of $w$ and $b$ we have different directions.

Is there a control variation that generates, e.g., $\delta(w_1, I_1) + \delta(w_2, I_2)$?

*Caution:* we cannot use $w_1 + w_2$ because nobody said that $w_1 + w_2 \in \mathcal{U}$! $\mathcal{U}$ can be non-convex, actually it can be discrete. (Also, $b_1 + b_2$ may be $> t^*$.)

The right answer is to perturb $u^*$ *both* at $b_1$ (by $w_1$) and $b_2$ (by $w_2$):

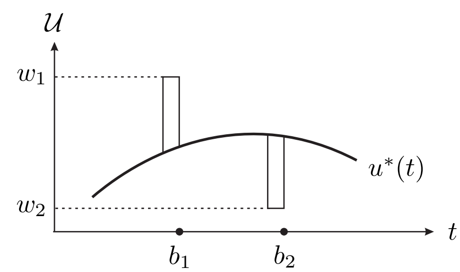


Figure 4.7: Adding spatial variations

$\longrightarrow$ This assumes $b_1 \neq b_2$ —OK, can always move them slightly to ensure this. $\varepsilon$ small $\Rightarrow I_1, I_2$ do not overlap.

$\longrightarrow$ Why will the overall perturbation be the sum of the two?

Because the variational equation $\dot{\psi} = A_*\psi$ is *linear*!

$$\Phi_*(t^*, b_2)\left(\psi_2(b_2) + \Phi_*(b_2, b_1)\psi_1(b_1)\right) = \Phi_*(t^*, b_2)\psi_2(b_2) + \Phi_*(t^*, b_1)\psi_1(b_1)$$

(recalling the semigroup property of the transition matrix).

More generally, if we want $\beta_1\delta(w_1, I_1) + \beta_2\delta(w_2, I_2)$ for some $\beta_1, \beta_2 > 0$, then we have to adjust the length of the intervals $I_1, I_2$: $I_1 = (b_1 - \varepsilon\beta_1 a_1, b_1]$, $I_2 = (b_2 - \varepsilon\beta_2 a_2, b_2]$. Same for more than two directions.

We see that these concatenated needle variations yield a *convex cone* of infinitesimal perturbations of the terminal point $y^*(t^*)$. We call this convex cone $\mathrm{co}(\vec{P})$. ($\vec{P}$ is the unconvexified cone of the individual directions.) *Convex cone* means $y_1, y_2 \in \vec{P} \Rightarrow \beta_1 y_1 + \beta_2 y_2 \in \vec{P}$ for all $\beta_1, \beta_2 > 0$ (this is for vertex at 0). Here, $y^*(t^*)$ is the *vertex* of our cone (so we should really write $y_1 - y^*(t^*)$, $y_2 - y^*(t^*)$).

Recall that we also have the line $\vec{\rho}$ of perturbations arising from the temporal variations of $u^*$. Adding it to the convex cone $\mathrm{co}(\vec{P})$ of perturbations arising from spatial (needle) variations, we obtain the *terminal cone*

$$C_{t^*} := \vec{\rho} + \mathrm{co}(\vec{P}) = \{y = y_1 + y_2 : y_1 \in \vec{\rho}, y_2 \in \mathrm{co}(\vec{P})\}$$

It is easy to check that $C_{t^*}$ is again a convex cone, with vertex at $y^*(t^*)$.

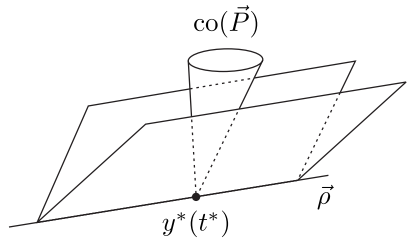


Figure 4.8: Terminal cone

Everything between the two hyperplanes is $C_{t^*}$.

Every vector (direction) in $C_{t^*}$ corresponds to an infinitesimal perturbation induced by a general (temporal plus spatial) control variation:

We will not formally prove this, but it is clear: spatial perturbations get propagated by the variational equations and add up as before, and at the end we also add the effect of the temporal variation.

Preview of the next steps:

- We didn't use optimality of $u^*$ yet. Next, we will use it to show that $C_{t^*}$ must face "upward";
- Then we will define the supporting hyperplane at $t^*$ which flows back along the adjoint equation.

### 4.2.7 Key topological lemma

Consider the terminal cone $C_{t^*}$.

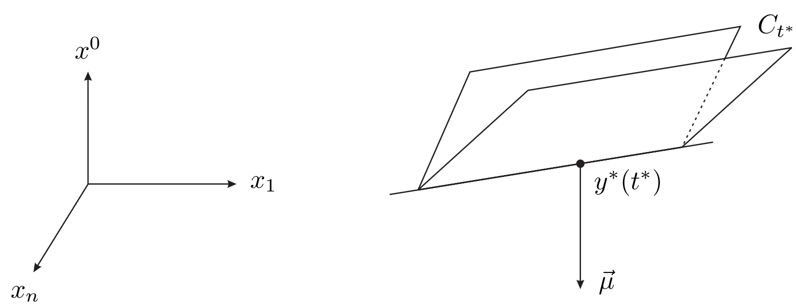


Figure 4.9: Illustrating the separation property

Let $\mu = (-1, 0, \ldots, 0)$, and let $\vec{\mu}$ be the ray generated by this vector (downward). Optimality suggests that $\vec{\mu} \notin C_{t^*}$. Let us make this precise.

**Lemma 6** *$\vec{\mu}$ does not intersect the interior of the cone $C_{t^*}$.*

Let's see if this is obvious. If the statement of the lemma is not true, $\vec{\mu} \in C_{t^*}$ which means that there exists a perturbation $u$ of $u^*$ and the corresponding trajectory $y(t)$ such that at its terminal time, $\hat{t}$, we have

$$y(\hat{t}) = y^*(t^*) + \varepsilon\beta\mu + o(\varepsilon)$$

for some $\beta > 0$. Recall that $y = \binom{J}{x}$, $y^* = \binom{J^*}{x^*}$, so

$$J(u) = J^* - \varepsilon\beta + o(\varepsilon)$$

$$x(\hat{t}) = x_1 + o(\varepsilon)$$

We are claiming that this cannot happen. Note that there isn't really any contradiction with optimality of $x^*$, since the perturbed trajectory $x$ need not hit the target set. What the lemma basically says is that in a small tube around $x^*$, the control $u^*$ should still be *approximately* (up to $o(\varepsilon)$) optimal. This is how we can get away with considering perturbations which destroy the terminal constraint. But we see now that the lemma is not obvious.

In [A-F] this approximate optimality claim is not made very precise. [A-F] also don't show why we need interior of $C_{t^*}$. Let's look into this more carefully. This is actually [Pontryagin et al., Lemma 3, pp. 109–112]. (Note: page numbers refer to the Russian version, that's the only one I have.)

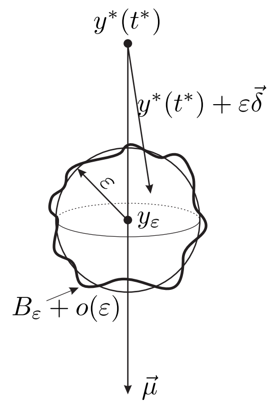


Figure 4.10: Proving the lemma

For some small $\varepsilon > 0$, consider a point $y_\varepsilon := y^*(t^*) + \varepsilon\beta\mu$, $\beta > 1$ on the ray $\vec{\mu}$. The ray $\vec{\mu}$ is in the *interior* of $C_{t^*}$ $\Rightarrow$ nearby directions are also in $C_{t^*}$. Consider the ball $B_\varepsilon$ of radius $\varepsilon$ centered at $y_\varepsilon$. Points in this ball are of the form $y^*(t^*) + \varepsilon\delta$ where $\delta$ are directions (vectors) from $C_{t^*}$, obtained using the earlier control variations.

But the actual points we can hit are

$$y^*(t^*) + \varepsilon\delta + o(\varepsilon)$$

So the set of points actually reachable by perturbed trajectories is of this form, which means it's a "warped" version of $B_\varepsilon$—it's $o(\varepsilon)$ away from $B_\varepsilon$.

Now, view $\varepsilon > 0$ as a parameter and let $\varepsilon \to 0$. The radius of $B_\varepsilon$ is $\varepsilon$, the "warping" is of the order $o(\varepsilon)$ $\Rightarrow$ for $\varepsilon$ small enough, the perturbed (warped) set will still intersect the ray $\vec{\mu}$. This means we can actually hit $x_1$ (i.e., find a perturbed trajectory as above but with $x(\hat{t}) = x_1$) with a lower value of the cost—*contradiction*.

This is the argument in Pontryagin et al., still not completely rigorous, they have a footnote on p. 112. Warped ball can have a hole (or dent) in it! However, this can be settled using continuity of the map $B_\varepsilon \to \tilde{B}_\varepsilon$ (warping). In fact, can show that the warped ball *contains*, for $\varepsilon$ small enough, a ball centered at $y_\varepsilon$ whose radius is of the order $\varepsilon - o(\varepsilon)$. One way to do it is to apply Brouwer's fixed point theorem. See Sussmann, Lemma 5.3.1 ($r \leftrightarrow \varepsilon$, $\rho \leftrightarrow o(\varepsilon)$).

### 4.2.8 Separating hyperplane

The interior of $C_{t^*}$ and the ray $\vec{\mu}$ are convex sets. By the *separating hyperplane theorem* there exists a hyperplane which

- passes through the point $y^*(t^*)$;
- contains $C_{t^*}$ in one closed half-space and $\vec{\mu}$ in the other.

The separating hyperplane theorem is a standard result in convex analysis. See, e.g., [A-F], Chap. 3, or Luenberger, or Rockafellar's book "Convex analysis" for more details. Note: the above discussion tacitly assumes that the interior of $C_{t^*}$ is non-empty, but the other, degenerate case can also be handled.

**Exercise 12** [C] (due Mar 14) Find a suitable version of the separating hyperplane theorem which applies in the above situation, without any extra assumptions (such as nonempty interior of $C_{t^*}$). Write down a precise statement and a reference where you took it from.

Let the normal to this hyperplane be

$$\begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix} \neq 0$$

(this is the definition of $p_0^*$ and $p^*(t^*)$). The equation of the hyperplane is

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, y \right\rangle = \left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, y^*(t^*) \right\rangle$$

The second property above says:

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \mu \right\rangle \geq 0$$

and

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \delta \right\rangle \leq 0 \qquad \forall\, \delta \in C_{t^*}$$

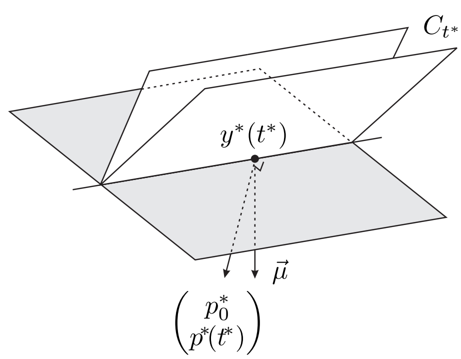


Figure 4.11: Separating hyperplane for Special Problem 1

Note that in view of the definition of $\mu$, the first inequality above says that $p_0^* \leq 0$ as needed.

Now, we will use $(p_0^*, p^*(t^*))$ as a boundary condition to specify the adjoint vector.

### 4.2.9 The adjoint equation

Recall that two systems of the form $\dot{x} = Ax$ and $\dot{z} = -A^T z$ are called *adjoint*. The meaning of this is

$$\langle x(t), z(t)\rangle = \text{const}$$

(or $x^T(t)z(t) = \text{const}$). Proof: $\frac{d}{dt}\langle x, z\rangle = \langle \dot{x}, z\rangle + \langle x, \dot{z}\rangle = x^T A^T z - x^T A^T z = 0$.

In our case, as the first system take the variational equation (4.7). We have

$$-A_*^T(t) = \begin{pmatrix} 0 & 0 \\ -L_x|_* & -(f_x)^T\Big|_* \end{pmatrix}$$

Denote the adjoint vector by $\binom{p_0}{p}$. Then the adjoint system is

$$\begin{aligned} \dot{p}_0 &= 0 \\ \dot{p} &= -L_x|_* p_0 - (f_x)^T\Big|_* p \end{aligned}$$

(the first term on the right-hand side is a column vector, the second is a matrix). The first ODE says that $p_0^* \equiv \text{const}$. The second is

$$\dot{p} = -H_x(x^*, u^*, p)$$

where recall that $H := \langle p, f(x, u)\rangle + p_0 L$.

And we have

$$\left\langle \begin{pmatrix} p_0 \\ p \end{pmatrix}, \psi \right\rangle = \text{const}, \qquad t \in [t_0, t^*]$$

Geometrically, we can associate $(p_0, p(t))$ with a hyperplane passing through $y^*(t)$ with normal $(p_0, p(t))$:

$$\left\{ y : \left\langle \begin{pmatrix} p_0 \\ p(t) \end{pmatrix}, y \right\rangle = \left\langle \begin{pmatrix} p_0 \\ p(t) \end{pmatrix}, y^*(t) \right\rangle \right\}$$

Thus each solution of the adjoint system corresponds to a family of hyperplanes, "moving" along the optimal trajectory.

$\longrightarrow$ And in particular we see that the perturbed trajectory always remains on one side of the hyperplane.

We let $p^*(t)$ denote the solution of the adjoint equation with the above terminal condition.

$\longrightarrow$ The hyperplane with the above separation property is not unique, but once it's fixed, $p^*(t)$ is unique up to a constant multiple.

We think of the adjoint equation as *flowing back*.

It is clear that the vector

$$\begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix} \neq 0$$

because it is a normal to the supporting hyperplane. Since

$$\frac{d}{dt}\begin{pmatrix} p_0^* \\ p^*(t) \end{pmatrix} = -A_*^T(t)\begin{pmatrix} p_0^* \\ p^*(t) \end{pmatrix}$$

is a homogeneous (unforced) LTV system, we have

$$\begin{pmatrix} p_0^* \\ p^*(t) \end{pmatrix} \neq 0 \qquad \forall\, t$$

as claimed.

### 4.2.10 Proof of the Maximum Principle for Special Problem 1

We are now home free, can prove everything.

- *Hamiltonian maximization:*

  Recall that the perturbation caused by a needle variation with parameters $w, b, a$ was

  $$y(t) = y^*(t) + \varepsilon\Phi_*(t, b)\nu_b(w)a + o(\varepsilon)$$

  where

  $$\nu_b(w) = g(y^*(b), w) - g(y^*(b), u^*(b)) = \begin{pmatrix} L(x^*(b), w) - L(x^*(b), u^*(b)) \\ f(x^*(b), w) - f(x^*(b), u^*(b)) \end{pmatrix}$$

  ($\nu_b(w)$ is propagated by $\Phi_*(\cdot, b)$).

  $\Phi_*(t^*, b)\nu_b(w)a = \psi(t^*)$ is a direction in $C_{t^*}$, so using the second inequality for the separating hyperplane with this as $\delta$, we get

  $$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \Phi_*(t^*, b)\nu_b(w) \right\rangle \leq 0$$

  Flowing back to time $t = b$ (or any other $t$), we have

  $$\left\langle \begin{pmatrix} p_0^* \\ p^*(b) \end{pmatrix}, \nu_b(w) \right\rangle \leq 0$$

  Writing out the components of $\nu_b(w)$:

  $$p_0^* L(x^*(b), w) + \langle p^*(b), f(x^*(b), w)\rangle \leq p_0^* L(x^*(b), u^*(b)) + \langle p^*(b), f(x^*(b), u^*(b))\rangle$$

  which is

  $$H(x^*(b), w, p^*(b)) \leq H(x^*(b), u^*(b), p^*(b))$$

  Since $b$ and $w$ can be arbitrary, the claim is established.

  In fact, we could have written the Hamiltonian from the beginning as

  $$H(x, u, p, p_0) = \left\langle \begin{pmatrix} p_0^* \\ p^* \end{pmatrix}, \begin{pmatrix} L(x, u) \\ f(x, u) \end{pmatrix} \right\rangle$$

  So, the geometric interpretation of the Hamiltonian is the inner product of the adjoint vector with the velocity vector of the state trajectory ($y$), and all perturbations away from $y^*$ can only decrease it because they act in the directions which make non-positive inner product with the adjoint vector.

  $\longrightarrow$ $b$ cannot be a discontinuity of $u^*$. So, the maximization condition actually holds *almost everywhere.* This is OK because modifying $u^*$ on a set of measure 0 is inconsequential.

- $H|_* \equiv 0$:

  First, recall that temporal variations give perturbations

$$\delta(\tau) = \begin{pmatrix} L(x^*(t^*), u^*(t^*)) \\ f(x^*(t^*), u^*(t^*)) \end{pmatrix} \tau \in C_{t^*}$$

  We must have

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \delta(\tau) \right\rangle \leq 0$$

  but since $\tau$ can be of any sign, this is $= 0$. We conclude that

$$p_0^* L(x^*(t^*), u^*(t^*)) + \langle p^*(t^*), f(x^*(t^*), u^*(t^*)) = 0$$

  which means $H(x^*(t^*), p^*(t^*), u^*(t^*)) = 0$. So $H|_*$ is 0 at the terminal time; we need to show that it is 0 everywhere.

  First, let's see that $H(x^*(t^*), p^*(t^*), u^*(t^*))$ as a function of $t$ is *continuous*, even though $u^*$ is not! Let $t$ be a point of jump in $u^*$. This is basically the same argument as the one used to show the second Weierstrass-Erdmann condition (see Problem 9). Of course, $x^*$ and $p^*$ are continuous. We have

$$H(x^*(t), p^*(t), u^*(t^-)) \geq H(x^*(t), p^*(t), u^*(t^+))$$

  from applying the maximization condition at $t^-$, and similarly

$$H(x^*(t), p^*(t), u^*(t^-)) \leq H(x^*(t), p^*(t), u^*(t^+))$$

  from applying the maximization condition at $t^+$.

  Finally, let's show that on intervals where the control is continuous, the above function is actually constant. In calculus of variations we showed this by simply differentiating, but here we need to be somewhat more careful because the existence of $H_u$ is not assumed. Can still prove using the maximization condition.

  **Exercise 12** [D] Show that on any interval where $u^*$ is continuous, $H(x^*(\cdot), p^*(\cdot), u^*(\cdot))$ as a function of $t$ is constant (has 0 derivative). Hint: for a pair of nearby times $t$ and $t'$, consider the expression

$$\lim_{t' \to t} \frac{H(x^*(t'), p^*(t'), u^*(t')) - H(x^*(t), p^*(t), u^*(t))}{t' - t}$$

  and use the Hamiltonian maximization condition.

  The function $H(x^*(\cdot), p^*(\cdot), u^*(\cdot))$ of $t$ is:

  - 0 at $t = t^*$;
  - continuous;
  - has 0 time derivative a.e.

  Therefore, it is identically 0.

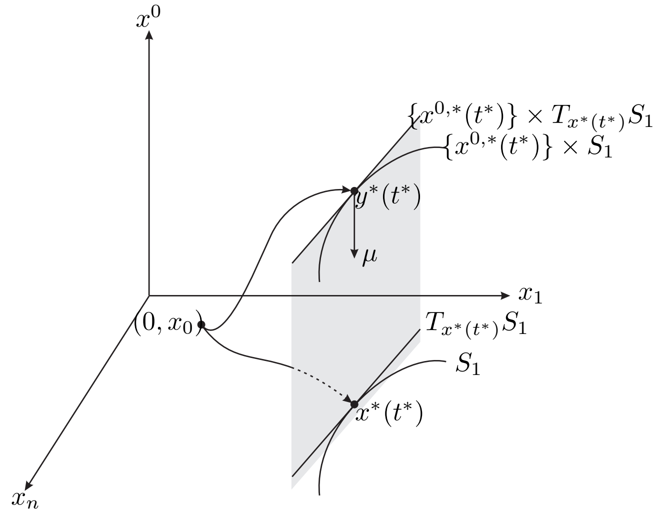


Figure 4.12: Deriving transversality conditions

### 4.2.11 Proof of the Maximum Principle for Special Problem 2: transversality conditions

In the previous case, $C_{t^*}$ was separated from $\vec{\mu}$, because hitting a point on $\vec{\mu}$ below $y^*(t^*)$ contradicted optimality. In the present case, hitting a point with lower cost whose projection onto the $x$-space is any point in $S_1$ (not necessarily $x^*(t^*)$) will contradict optimality. Since we are dealing with infinitesimal directions, can show that $C_{t^*}$ is separated from the set

$$T := \left\{ y : y = y^*(t^*) + \begin{pmatrix} 0 \\ d \end{pmatrix} + \beta\mu,\ d \in T_{x^*(t^*)}S_1,\ \beta \geq 0 \right\}$$

$d \in \mathbb{R}^n$ is a tangent direction. $T_{x^*(t^*)}S_1$ is viewed as a *subspace* which, when translated to $y^*(t^*)$, gives the upper line shown in the picture.

**Lemma 7** *$T$ does not intersect the interior of the cone $C_{t^*}$.*

What follows is the argument from [Pontryagin et al., Lemma 10] (generalization of their Lemma 3 we used earlier). The argument given there is quite detailed.

We are looking at the "semi-plane" $T$ and the surface (manifold with boundary), call it $D$, which lies "directly under" $\{x^{0,*}(t^*)\} \times S_1$. These are the plane and the curved surface in Figure 4.13, respectively. Both are bounded from above but extend infinitely far down. They intersect along $\vec{\mu}$, and $T$ is tangent to $D$ along $\vec{\mu}$.

Suppose the lemma is not true. Then some direction $\delta$ in the "semi-plane" $T$ lies in $C_{t^*}$. As before, consider a small ball $B_\varepsilon$, of radius $\varepsilon$ and centered at some point on $\delta$. The points we can actually hit are as before given by its "warped" version, $B_\varepsilon + o(\varepsilon)$. On the other hand, using tangency property, for small $\varepsilon$ we have that the plane $T$ and the surface $D$ are $o(\varepsilon)$ apart. Hence, for $\varepsilon$ small enough, the "warped ball" $B_\varepsilon + o(\varepsilon)$ actually intersects the surface $D$, since $B_\varepsilon$ has radius $\varepsilon$ and is centered on $\delta \in T$, while the perturbations that give us $B_\varepsilon + o(\varepsilon)$ from $B_\varepsilon$ and $D$ from $T$ are of order $o(\varepsilon)$.

As before, there exists a separating hyperplane with normal that we denote by $(p_0*, p^*(t^*))$, and we have

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \begin{pmatrix} 0 \\ d \end{pmatrix} + \beta\mu \right\rangle \geq 0 \qquad \forall\, d \in T_{x^*(t^*)}S_1,\ \forall\, \beta \geq 0$$

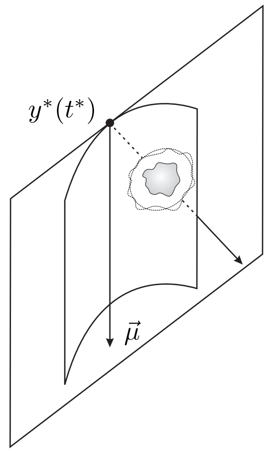


Figure 4.13: Finishing the proof

and

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \delta \right\rangle \leq 0 \qquad \forall\, \delta \in C_{t^*}$$

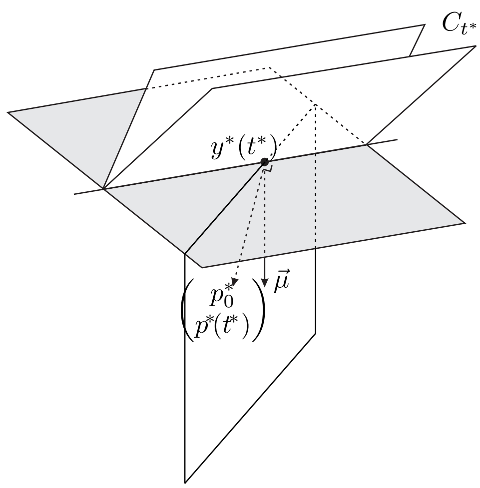


Figure 4.14: Separating hyperplane for Special Problem 2

In particular, from the first inequality we still have

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \mu \right\rangle \geq 0$$

(since $d = 0 \in T_{x^*(t^*)}S_1$). So we see that all previous conclusions still hold. We also have

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \begin{pmatrix} 0 \\ d \end{pmatrix} \right\rangle \geq 0$$

for each tangent direction $d$. $d \in T_{x^*(t^*)}S_1 \Rightarrow -d \in T_{x^*(t^*)}S_1 \Rightarrow \langle p^*(t^*), d\rangle = 0$ for each tangent direction. This is the desired transversality condition.

$\longrightarrow$ In fact it's clear that the separating hyperplane must pass through $T_{x^*(t^*)}S_1$, as well as through $\vec{\rho}$.

If $S_1 = \mathbb{R}^n$ then $C_{t^*}$ is separated from the entire $(n+1)$-dimensional half-space below $y^*(t^*)$. The separating hyperplane must be horizontal $\Rightarrow$ the normal is vertical $\Rightarrow p^*(t^*) = 0$.

## 4.3 A discussion of the Maximum Principle

- It is a *necessary* condition—it helps single out candidates for the optimal $u^*$, then need to analyze them separately to determine the optimal one. Optimal control may not even exist—more on this issue later.

- (continuing from the previous point) It is a necessary condition for *local* optimality (and in particular for global optimality). Conditions of the Maximum Principle hold if $u^*$ provides a lower cost compared to other $u$ which produce close trajectories.

  $\longrightarrow$ $u$ itself need not be close to $u^*$ in the supremum norm—the claim on p. 348 in [A-F] is *wrong*.

- The maximization of $H$ condition is a *global* one, over all $u \in \mathcal{U}$.

  Begin optional:________________________________________

- the Maximum Principle (in its classical version) relies on first-order analysis, in the sense that we get perturbations of the form

  $$y^*(t^*) + \varepsilon\delta + o(\varepsilon)$$

  and then use optimality to conclude that $\delta$ are suitably constrained (don't point in the direction of decreasing cost). However, we could study those needle variations that give $\delta = 0$, and then second-order terms become important:

  $$y^*(t^*) + \varepsilon^2\nu + o(\varepsilon^2)$$

  and $\nu$ can be used to obtain *higher-order* versions of the Maximum Principle [Krener, Bressan, Agrachev].

  Non-constant needle variations are also used (to prove Goh condition [Agrachev-Sachkov] for singular trajectories).

  End optional ________________________________________

As we discussed, we can get the Maximum Principle for other problems by reducing them, via changes of variables, to Special Problem 1 or 2. Alternatively (and more instructively), knowing the proof, in some cases we can see how it can be modified to get the result. Let's consider some specific scenarios ([A-F] have all details).

- *Fixed terminal time:*

  If $t^*$ is fixed, there are no temporal variations to consider, so the terminal cone $C_{t^*}$ will not have the line $\vec{\rho}$ coming from temporal perturbations $\delta(\tau)$.

  Where are $\delta(\tau)$ used in the proof?

  In one place only: to show that $H(t^*) = 0$. So we conclude that $H$ will be a constant but not necessarily 0 along the optimal trajectory (cf. calculus of variations with fixed time, earlier).

Let's confirm this by reducing to the free-time problem via change of variables. Time becomes state: $\dot{x}_{n+1} = 1$; system becomes

$$\begin{aligned} \dot{x}^0 &= L \\ \dot{x} &= f \\ \dot{x}_{n+1} &= 1 \end{aligned}$$

Target set is

$$\left\{ \begin{pmatrix} x \\ x_{n+1} \end{pmatrix} : x_{n+1} = t_1, \text{ plus whatever constraints we had on } x(t_1) \right\}$$

Time is no longer explicitly constrained, and our version of the Maximum Principle for Special Problem 2 applies. The Hamiltonian for this new problem will be 0:

$$\bar{H} = \underbrace{p_0^* L + p^* \cdot f}_{\text{original } H} + p_{n+1}^* \cdot 1 \equiv 0$$

$p_{n+1}^*$ satisfies $\dot{p}_{n+1}^* = -H_{x_{n+1}}\big|_* = 0 \Rightarrow p_{n+1}^* \equiv \text{const}$, and so the original Hamiltonian $H$ will be equal to the constant $-p_{n+1}^*$.

- What if in the above discussion, $L$ or $f$ depend explicitly on time? Then in the new auxiliary problem, $\bar{H}$ will depend on $x_{n+1}$, so $p_{n+1}^*$ will not be constant and hence the original $H$ will not be constant.

- *Terminal cost:*

  $K(x(t_f))$. For simplicity, suppose that there is no running cost ($L \equiv 0$). Then, no need to consider the $x^0$ coordinate. Suppose also the *free-endpoint, free-time case* (just to be specific; free time is not so important but free endpoint is important).

  Consider again perturbations $x^*(t^*) + \varepsilon\vec{\delta} + o(\varepsilon)$. We expect that by optimality, these perturbations cannot decrease $K$:

  $$\langle K_x(x^*(t^*)), \delta\rangle \geq 0 \qquad \forall\, \delta \in C_{t^*}$$

  Geometrically, $C_{t^*}$ is separated from the half-space defined by vectors emanating from $x^*(t^*)$ and making non-positive inner product with $K_x(x^*(t^*))$. Recall that the endpoint is free, and so any deviation from $x^*(t^*)$ is in principle admissible; this situation is similar to Special Problem 2 where $S_1$ is the entire $\mathbb{R}^n$ and $-K_x(x^*(t^*))$ plays the role of $\mu$. Also, since $x^*(t_1)$ is free, $p^*(t^*)$ should be completely constrained (so that no degrees of freedom are left and the necessary condition hopefully gives a unique candidate). So, how should we define $p^*(t^*)$ to have $\langle p^*(t^*), \delta\rangle \leq 0$ $\forall\, \delta \in C_{t^*}$? Obvious: set $p^*(t^*) = -K_x(x^*(t^*))$, or its positive multiple. This corresponds to a hyperplane orthogonal to the gradient of $K$.

  *Note:* $\langle p^*(t^*), \delta\rangle \leq 0$ is the inequality that is really crucial in the proof of the Maximum Principle . The other inequality, $\langle p^*(t^*), \mu\rangle \geq 0$, was only needed to show that $p_0^* \leq 0$. Here we don't have this inequality, and we don't need it since we don't have $p_0^*$.

  Let's again confirm by reducing to Special Problem 2 via a change of variable.

  $$K(x(t_f)) = K(x_0) + \int_{t_0}^{t_f} \langle K_x, f(x,u)\rangle\, dt$$

So, ignoring $K(x_0)$ which is a known constant, we can work with

$$L = \langle K_x, f \rangle$$

For this modified problem, the Hamiltonian is

$$\bar{H} = \bar{p}_0 \langle K_x, f \rangle + \langle \bar{p}, f \rangle = \langle \bar{p}_0 K_x + \bar{p}, f \rangle$$

*Idea:* define $\bar{p}_0^* K_x(x^*(t)) + \bar{p}^*(t)$ to be the new adjoint vector, $p^*(t)$.

What about boundary condition? The vector $\bar{p}^*$ satisfies $\bar{p}^*(t^*) = 0$ (because the endpoint is free). So, we have

$$\bar{p}_0^* K_x(x^*(t^*)) + \bar{p}^*(t^*) = \bar{p}_0^* K_x(x^*(t^*))$$

Thus the terminal condition we get for $p^*$ is $p^*(t^*) = \bar{p}_0^* K_x(x^*(t^*))$ and this is consistent with the previous discussion because $\bar{p}_0^* \leq 0$.

But we need to verify that this $p^*$ satisfies the correct canonical equation $\dot{p}^* = -H_x|_*$ with respect to the Hamiltonian $H = \langle p, f \rangle$!

$$p^* = \bar{p}_0^* \left. K_x \right|_* + \bar{p}^*$$

where $\dot{\bar{p}}^* = -\left.\bar{H}_x\right|_*$, $\bar{H} = \bar{p}_0^* \langle K_x, f \rangle + \langle \bar{p}^*, f \rangle$. We have

$$\begin{aligned} \dot{p}^* &= \bar{p}_0^* \left\langle \left. K_{xx} \right|_*, f \right\rangle - \bar{p}_0^* \left\langle \left. K_{xx} \right|_*, f \right\rangle - \bar{p}_0^* \left.(f_x)^T\right|_* \left. K_x \right|_* - \left.(f_x)^T\right|_* \bar{p}^* \\ &= -\left.(f_x)^T\right|_* p^* = -\left. H_x \right|_* \end{aligned}$$

where $H = \langle p, f \rangle$.

**Exercise 13** (due Mar 28)

Consider the general problem

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

$$J = \int_{t_0}^{t_f} L(t, x, u) dt + K(x(t_f))$$

($t_f$ is free), $S = \mathbb{R} \times S_1$ where $S_1$ is a $k$-dimensional surface. Reduce this by a change of variables to Special Problem 2 and arrive at a precise statement of the Maximum Principle for this general problem.

The statement of the Maximum Principle contains the condition

$$\begin{pmatrix} p_0^* \\ p^*(t) \end{pmatrix} \neq 0 \qquad \forall\, t$$

(this is proved in Step 9). When can we say that $p^*(t) \neq 0\ \forall\, t$?

**Example 6** $L \neq 0$, free terminal time (e.g., time-optimal: $L \equiv 1$). $H = p_0^* L + \langle p^*, f \rangle \equiv 0$. If $p^*(t) = 0$ then $p_0^* L = 0$ at $t \Rightarrow p_0^* = 0$ (since $L \neq 0$) and this is a contradiction, hence $p^*(t) \neq 0\ \forall\, t$.

□

**Example 7** Suppose $L_x = 0$ (again, true for time-optimal: $L \equiv 1$, or when there is no running cost).

$$\dot{p}^* = -p_0^* L_x - (f_x)^T\Big|_* p^* = -(f_x)^T\Big|_* p^*$$

which is (again) a homogeneous LTV system. We have a terminal condition like, e.g., $p^*(t^*) = K_x(x^*(t^*))$. If this is not 0, then $p^*(t) \neq 0 \; \forall \, t$.

□

By similar reasoning, we see that if we have $(p_0^*, p^*(t)) = (0, 0)$ for *some* $t$, then it is $(0, 0)$ for *all* $t$ (because $p_0^*$ is constant and $p^*$ then satisfies an LTV equation). So, "$\neq (0,0) \; \forall \, t$" in the statement of the Maximum Principle could actually be replaced by "$\neq 0$ for some $t$".

What does it mean that $p_0^* = 0$?

Supporting hyperplane is vertical, $C_{t^*}$ lies on one side:

Projecting onto the $x$-space:

All perturbations bring us to one side—very unlikely. Cf. abnormal cases in calculus of variations .

This is more common:

Also, we cannot take $p_0^* \neq 0 \Rightarrow$ cannot tilt the hyperplane $\Rightarrow C_{t^*}$ fills the whole half-space—really unlikely.

$\longrightarrow$ In all the cases that we'll encounter, $p_0^* \neq 0$. *Can normalize it to* $-1$ *and forget it.*