LECTURE 3

4. The Maximum Principle

In the previous lecture we saw how far calculus of variations can take us when it comes to solving optimal control problems. We now move on to consider Pontryagin's maximum principle which provides us with a more powerful tool for solving optimal control problems.

4.1. **Basic variable-time, fixed-endpoint problem.** For the basic variable-time fixed endpoint problem we assume that $f = f(x, u)$, $L = L(x, u)$, $K \equiv 0$ and $S = [t_0, \infty) \times \{x_1\}$.

Let $u^* : [t_0, t_f] \to U$ be an optimal control and let $x^* : [t_0, t_f] \to \mathbb{R}^n$ be the corresponding trajectory, with $x^*(t_0) = x_0$ and $x^*(t_f) = x_1$. Then there exists a $p^* : [t_0, t_f] \to \mathbb{R}^n$ and a $p_0^* \leq 0$ such that $(p^*(t), p_0^*) \neq (0, 0)$, $\forall t \in [t_0, t_f]$ (non-triviality), and

(1) $\dot{x}^* = H_p\big|_*$ and $\dot{p}^* = -H_x\big|_*$, where
$$H(x, u, p, p_0) = \langle p, f(x, u)\rangle + p_0 L(x, u).$$

(2) For fixed t, the function $u \mapsto H(x^*(t), u, p^*(t), p_0^*)$ takes a global maximum at $u = u^*(t)$. Hence,
$$H(x^*(t), u^*(t), p^*(t), p_0^*) \geq H(x^*(t), u, p^*(t), p_0^*),$$
for all $u \in U$ and all $t \in [t_0, t_f]$.

(3) $H(x^*(t), u^*(t), p^*(t), p_0^*) = 0$, for all $t \in [t_0, t_f]$.

4.2. **Basic variable-time, variable-endpoint problem.** In the variable-endpoint problem we change the final set to $S = [t_0, \infty) \times S_1$. Here, $S_1 = \{x \in \mathbb{R}^n : h_1(x) = h_2(x) = \ldots = h_{n-k}(x) = 0\}$ with h_i, $i = 1, \ldots, n-k$, $\mathcal{C}^1$ functions and assuming that every point $x \in S_1$ is a regular point of S_1.

To be valid for this case the maximum principle for the basic fixed-endpoint problem given above has to be changed to $x^*(t_f) \in S_1$ and we have to add the following (transversality condition):

(4) $\langle p^*(t_f), d\rangle = 0$ for all $d \in T_{x^*(t_f)} S_1$.

With these changes we have freed k dimensions of the final state, but constrained k dimensions of the adjoint state in the final time, thus preserving the number of boundary conditions to the canonical equations.

4.3. **Proof of the maximum principle.** To prove the maximum principle we assume that we have found an optimal control u^* giving the trajectory x^*. We then consider a set of perturbations that shows that the various criteria in the maximum principle have to hold for the given control and trajectory to be optimal. The proof of the Maximum principle is divided into 10 steps

Step 1: First we move from Lagrange form to Mayer form by introducing the additional state x^0, with $\dot{x}^0 = L(x, u)$ and $x^0(t_0) = 0$. We then get the new differential equation
$$\dot{y} = \begin{pmatrix} \dot{x}^0 \\ \dot{x} \end{pmatrix} = \begin{pmatrix} L(x, u) \\ f(x, u) \end{pmatrix} =: g(y, u),$$

where $y := \begin{pmatrix} x^0 \\ x \end{pmatrix}$. The cost functional can now be written

$$J(u) = \int_{t_0}^{t_f} \dot{x}^0(t)dt = x^0(t_f)$$

and the new target set is $[t_0, \infty) \times \mathbb{R} \times \{x_1\} =: [t_0, \infty) \times S'$.

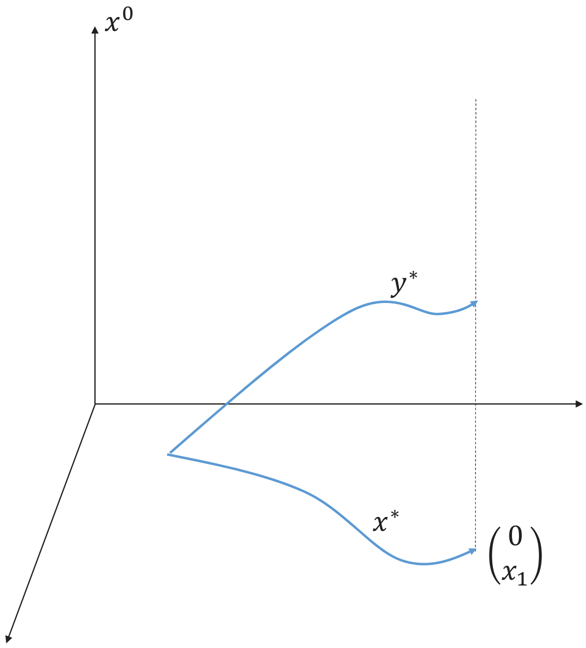

FIGURE 1. The trajectories x^* and y^*.

The Hamiltonian can now be written

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

Step 2: The next step is to consider what is referred to as *temporal control perturbations*. A temporal perturbation is achieved by adding a small perturbation to the final time t^*. Let $u_\tau(t) := u^*(\min\{t, t^*\})$, for $t \in [t_0, t^* + \varepsilon\tau]$.

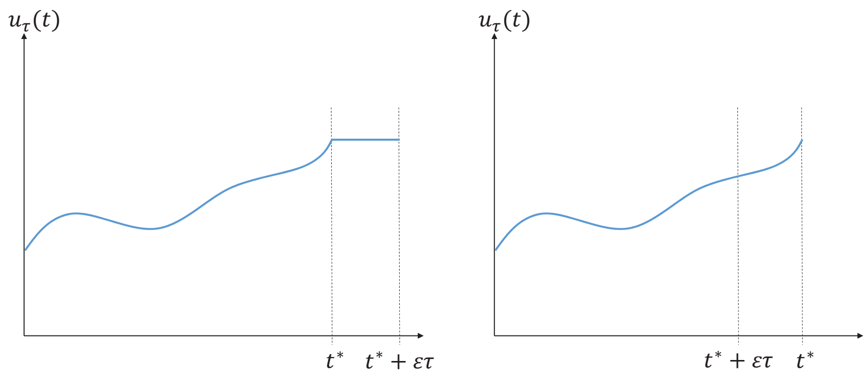

FIGURE 2. Temporal perturbations of u^*.

Let y be the perturbed trajectory. We first assume look at the case $\tau > 0$:

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

where $\delta(\tau)$ is a vector that is linear in τ. For $\tau < 0$ we have

$$\begin{aligned} y(t^* + \varepsilon\tau) &= y^*(t^* + \varepsilon\tau) \\ &= y^*(t^*) + g(y^*(t^*), u^*(t^*))\varepsilon\tau + o(\varepsilon) \\ &= y^*(t^*) + \varepsilon\delta(\tau) + o(\varepsilon). \end{aligned}$$

Hence, for a given small ε we get a vector in the enlarged state space as a result of the temporal control perturbation when varying τ over $\mathbb{R}$. We call this vector $\vec{\rho}$.

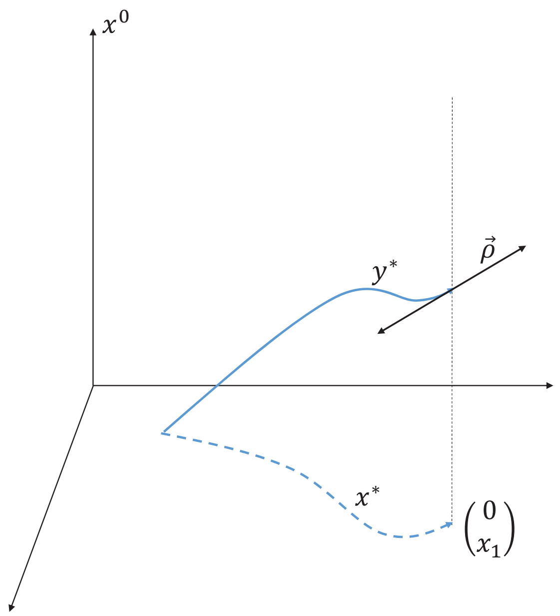

FIGURE 3. The result of temporal perturbations.

Step 3: Additional to the temporal perturbations we wish to apply *spatial control perturbations*. In the proof of the maximum principle we apply something called *needle perturbations* (Pontryagin-McShane perturbations).

Let $w \in U$ and $I := (b - \varepsilon a, b]$ where $a > 0$ and b is chosen such that $I \subset (t_0, t^*)$ for some $\varepsilon > 0$ and $u^*(t)$ is continuous at $t = b$. Define the perturbed control

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

Taylor expanding around $t = b$ gives

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

where $\approx$ denotes equality up to terms of order $o(\varepsilon)$. If we move to considering the perturbed trajectory, Taylor expansion around $t = b - \varepsilon a$ gives

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

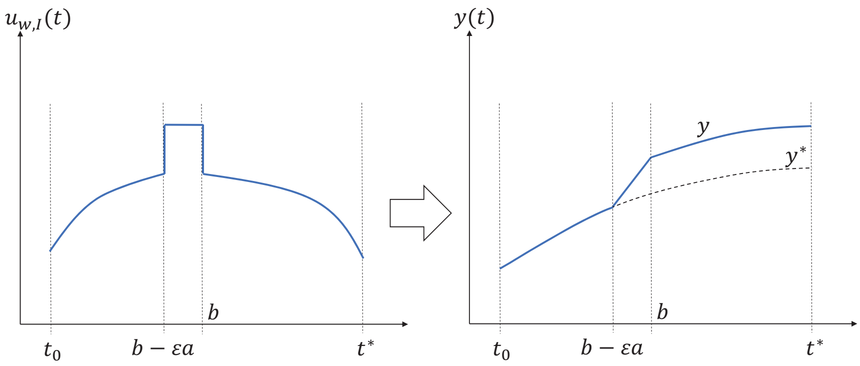

FIGURE 4. Spatial perturbations of u^*.

We want to find the difference between the optimal and the perturbed trajectory at time $t = b$ and Taylor expand the last part of this equation around $t = b$. We have

$$g(y^*(b - \varepsilon a), w)\varepsilon a \approx g(y^*(b), w)\varepsilon a + \underbrace{g_y(y^*(b), w)(y^*(b - \varepsilon a) - y^*(b))\varepsilon a}_{o(\varepsilon)}.$$

Hence,

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

where

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

In words this means that up to $o(\varepsilon)$ the difference between the optimal trajectory y^* and the perturbed trajectory y is the length of the interval times the difference in derivative at the end of the interval. This seems natural since u^* is continuous at b, and thus by the piecewise continuity assumption, continuous on the entire interval I for sufficiently small ε.

Step 4: We are now ready to write the variational equation for the needle perturbations. We want to find how, up to order $o(\varepsilon)$, the perturbation from Step 3 propagates up to the final time. For $t \geq b$ we let

$$y(t) = y^*(t) + \varepsilon\psi(t) + o(\varepsilon) =: y(t, \varepsilon), \tag{4.1}$$

where $\psi : [b, t^*] \to \mathbb{R}^{n+1}$ is the function we seek.

Differentiating $y(t, \varepsilon)$ w.r.t. ε we find $\psi(t) = y_\varepsilon(t, 0)$ and get the boundary condition $\psi(b) = y_\varepsilon(b, 0) = \nu_b(w)a$.

If we rewrite (4.1) as an integral equation we get

$$y(t, \varepsilon) = y(b, \varepsilon) + \int_b^t g(y(s, \varepsilon), u^*(s))ds.$$

Hence,

$$\psi(t) = y_\varepsilon(t, 0) = \nu_b(w)a + \int_b^t g_y(y^*(s), u^*(s))\psi(s)ds.$$

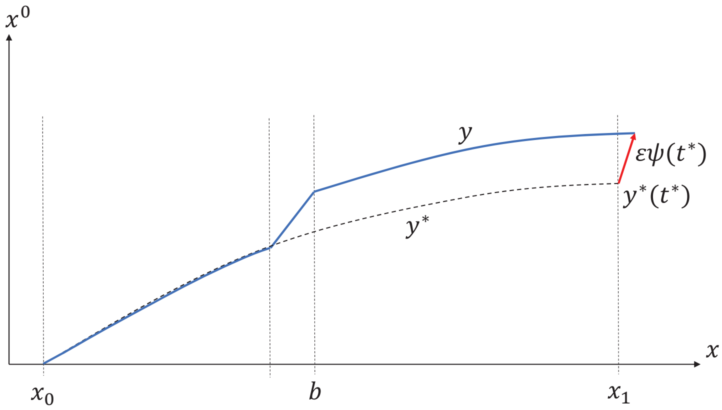

FIGURE 5. The result of spatial perturbations.

Differentiating with respect to t we find that

$$\dot{\psi}(t) = g_y(y^*(t), u^*(t))\psi(t) = g_y\big|_*(t)\psi(t) =: A_*(t)\psi(t),$$
$$\psi(b) = \nu_b(w)a.$$

For later purposes we can divide ψ into two parts,

$$\dot{\psi} = \begin{pmatrix} \dot{\eta}^0 \\ \dot{\eta} \end{pmatrix} = \begin{pmatrix} 0 & (L_x)^\top\big|_* \\ 0 & f_x\big|_* \end{pmatrix} \begin{pmatrix} \eta^0 \\ \eta \end{pmatrix}. \tag{4.2}$$

Since ψ is the solution to a linear differential equation there is a transition matrix $\Psi_*(\cdot,\cdot)$, such that

$$\psi(t_2) = \Psi_*(t_2, t_1)\psi(t_1).$$

Hence,

$$\psi(t^*) = \Psi_*(t^*, b)\psi(b) = \Psi_*(t^*, b)\nu_b(w)a$$

which gives us

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

Letting $\delta(w, I) := \Psi_*(t^*, b)\nu_b(w)a$ we get

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

Observe here that, since $a > 0$, having a perturbation that gives $\delta(w, I)$ does not mean that there is a corresponding perturbation w', I' such that $\delta(w', I') = -\delta(w, I)$. Hence, as opposed to the temporal perturbations, spacial perturbations give rise (up to $o(\varepsilon)$) to a set of unidirectional vectors starting at $y^*(t^*)$.

Step 5: The next step is to define what is referred to as the *terminal cone*. When combining two spatial perturbations we get

$$y(t^*) = y^*(t^*) + \varepsilon\delta(w_1, I_1) + \varepsilon\delta(w_2, I_2) + o(\varepsilon).$$

Letting $\bar{a}_i = \beta_i a_i$, with $\beta_i \geq 0$ we can combine any number m of spatial perturbations to get

$$y(t^*) = y^*(t^*) + \varepsilon\sum_{i=1}^{m}\beta_i\delta(w_i, I_i) + o(\varepsilon).$$

We then add temporal perturbations with $\bar{\tau} = \beta_0\tau$, $\beta_0 \in \mathbb{R}$, and consider points of the form

$$y = y^*(t^*) + \varepsilon\left(\beta_0\delta(\tau) + \sum_{i=1}^{m}\beta_i\delta(w_i, I_i)\right) + o(\varepsilon).$$

These y will define a convex cone C_{t^*} with apex at $y^*(t^*)$.

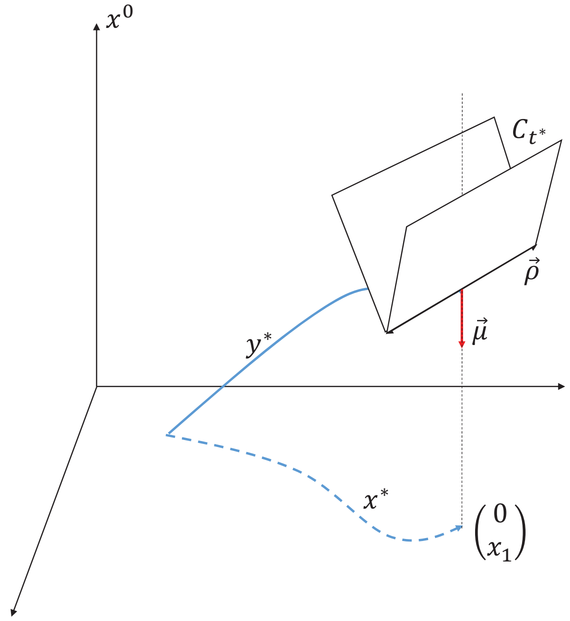

FIGURE 6. The terminal cone C_{t^*}.

Note here that since $I_i \subset (t_0, t^*)$ the temporal perturbations will not interfere with the spatial perturbations.

Step 6: Let $\mu := (-1, 0, \dots, 0)^\top \subset \mathbb{R}^{n+1}$ and let $\vec{\mu}$ be the ray generated by $a\mu$ with $a \geq 0$. A key topological lemma in the proof of the maximum principle is that, when (u^*, x^*) is optimal,

Lemma 4.1. *The intersection of $\vec{\mu}$ and int C_{t^*} is empty.*

A heuristic argument for this is given by assuming the opposite. Then there is a temporal and a spatial perturbation such that

$$y(t_f) = y^*(t^*) + \varepsilon\beta\mu + o(\varepsilon),$$

for some (arbitrary) $\beta > 0$. We have

$$\begin{aligned} J(u) &= J(u^*) - \varepsilon\beta + o(\varepsilon) \\ x(t_f) &= x_1 + o(\varepsilon). \end{aligned}$$

Now it might seem that we have a proof, but, since there is a $o(\varepsilon)$ difference between x_{t_f} and x_1, we need not hit the target set.

To obtain a formal proof assume that the lemma is false, we will then show that in this case there has to be a perturbation giving a $y(t_f)$ along the ray $\vec{\mu}$ under $y^*(t^*)$.

Proof. If the lemma is false we can choose a point $\hat{y}$ on the ray $\vec{\mu}$ with an ε-ball, B_ε, centered in $\hat{y}$, such that $B_\varepsilon \subset C_{t^*}$. For a suitable $\beta > 0$ we can write $\hat{y} = y^*(t^*) + \varepsilon\beta\mu$. Since $B_\varepsilon \subset C_{t^*}$ every point of B_ε can be written $y^*(t^*) + \varepsilon\nu$, where $\varepsilon\nu$ is a first order perturbation of the terminal point given by a combination of temporal and spatial perturbations. Since $\hat{y} \in C_{t^*}$, there is perturbation Δu_0, such that $\varepsilon\Delta u_0$ gives the first order perturbed terminal value $\hat{y}$.

To parameterize the perturbations leading to first-order terminal values in B_ε let e_i, $i = 1, \dots, n+1$, be $n+1$ perpendicular unit-vectors and chose $n+1$ perturbations Δu_i^+, such that $\varepsilon\Delta u_i^+$ leads to the first order perturbed terminal value $\hat{y} + \varepsilon e_i$ and $n+1$ perturbations Δu_i^-, such that $\varepsilon\Delta u_i^-$ leads to the first order perturbed terminal value $\hat{y} - \varepsilon e_i$.

Let

$$h^+(\rho) = \begin{cases} 0, & \text{for } \rho < 0, \\ \rho, & \text{for } \rho \geq 0, \end{cases}$$

and let

$$h^-(\rho) = \begin{cases} -\rho, & \text{for } \rho < 0, \\ 0, & \text{for } \rho \geq 0. \end{cases}$$

Then for $|\rho|^2 = \rho_1^2 + \ldots + \rho_{n+1}^2 \leq 1$ the perturbations

$$\Delta u(\rho_1, \ldots, \rho_{n+1}) := \varepsilon \left(1 - \sum_{i=1}^{n+1} |\rho_i|\right) \Delta u_0 + \varepsilon \sum_{i=1}^{n+1} h^+(\rho_i)\Delta u_i^+ + \varepsilon \sum_{i=1}^{n+1} h^-(\rho_i)\Delta u_i^- \tag{4.3}$$

give a parametrization of the perturbations leading to first-order terminal values in B_ε, with

$$\varepsilon\nu = \varepsilon \left(1 - \sum_{i=1}^{n+1} |\rho_i|\right) \nu_0 + \varepsilon \sum_{i=1}^{n+1} h^+(\rho_i)\nu_i^+ + \varepsilon \sum_{i=1}^{n+1} h^-(\rho_i)\nu_i^-,$$

where $\varepsilon\nu_0$, $\varepsilon\nu_i^+$ and $\varepsilon\nu_i^-$ are the first order-results of applying the perturbations $\varepsilon\Delta u_0$, $\varepsilon\Delta u_i^+$ and $\varepsilon\Delta u_i^-$, respectively.

The function defined in (4.3) mapping ρ to a point in B_ε is continuous and bijective and thus has a continuous inverse. Since the actual terminal point depends continuously on the perturbation, we can define a continuous "warping" map F by

$$F(y^*(t^*) + \varepsilon\nu) = y_{\rho_1,\ldots,\rho_{n+1}}(t_f),$$

Where $y_{\rho_1,\ldots,\rho_{n+1}}$ is the actual trajectory obtained when adding the perturbation $\Delta u(\rho_1, \ldots, \rho_{n+1})$ and t_f is the corresponding terminal time. The map F should be seen as taking first-order perturbations to actual perturbed terminal points, but can also be seen as a map from B_ε onto a warped version of B_ε which we denote $\tilde{B}_\varepsilon$.

Let $\hat{y}_\varepsilon = y^*(t^*) + \varepsilon\beta\mu$. Then $\hat{y}_\varepsilon \to y^*(t^*)$ along $\vec{\mu}$ as $\varepsilon \to 0$. B_ε that is now centered in $\hat{y}_\varepsilon$ is still in C_{t^*} and still consists of points $y^*(t^*) + \varepsilon\nu$ for all ε.

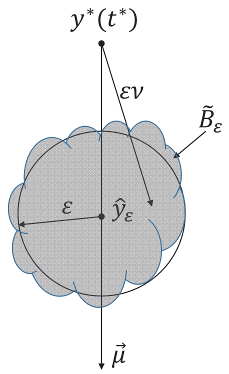

FIGURE 7. The first order perturbed terminal values and the actual perturbed terminal values.

To finish the proof we need to show that, for sufficiently small ε, $\tilde{B}_\varepsilon$ will contain points along $\vec{\mu}$ under $y^*(t^*)$.

For all $\alpha \in (0, 1)$ we have that $|o(\varepsilon)| < \alpha\varepsilon$, for sufficiently small ε. For an arbitrary z in the $(1-\alpha)\varepsilon$-ball around $\hat{y}_\varepsilon$ we want to find a y in B_ε such that $F(y) = z$, or equivalently if we define

$$G(y) := y - F(y) + z,$$

then $G(y) = y$.

Let $y \in B_\varepsilon$, then $|y - F(y)| = o(\varepsilon)$, so that $|y - F(y) + z - \hat{y}_\varepsilon| \le o(\varepsilon) + (1-\alpha)\varepsilon \le \varepsilon$. Hence, G is a continuous map from the compact convex set B_ε into itself. Hence, by Brouwer's fixed point theorem there is a y such that $G(y) = y$. Since z was arbitrary this finishes the proof. $\square$

Step 7: Since $\operatorname{int} C_{t^*}$ and $\vec{\mu}$ are convex disjoint sets there exists a hyperplane that separates them. This hyperplane must pass through $y^*(t^*)$. Lets denote by

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

a vector normal to this hyperplane. The equation for the hyperplane is then

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

and separation means that

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

for all $\delta \in \mathbb{R}^{n+1}$ such that $y^*(t^*) + \delta \in C_{t^*}$, and

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

Step 8: The two linear systems

$$\dot{x} = Ax \quad \text{and} \quad \dot{z} = -A^\top z$$

are called adjoint. We have that

$$\frac{d}{dt}\langle z, x\rangle = \langle \dot{z}, x\rangle + \langle z, \dot{x}\rangle = -x^\top A^\top z + x^\top A^\top z = 0.$$

Hence, the angle that the state vectors for a pair of adjoint systems make is constant over time. Remembering the variational equation

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

The adjoint to this system is

$$\dot{z} = -A_*^\top(t) z = \begin{pmatrix} 0 & 0 \\ -L_x\big|_* & -(f_x)^\top\big|_* \end{pmatrix} z.$$

Hence with $z = \begin{pmatrix} p_0^* \\ p^*(t) \end{pmatrix}$ we get $\dot{p}_0^* = 0$ and

$$\dot{p}^* = -L_x\big|_* p_0^* - \langle f_x\big|_*, p^*\rangle = -H_x(x^*, u^*, p^*, p_0^*).$$

The way the costate was defined in the maximum principle we see that, for all $t \in [t_0, t^*]$,

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

Gemoetrically this means that we can see $\begin{pmatrix} p_0^* \\ p^*(t) \end{pmatrix} \neq 0$ as the normal to a hyperplane passing through $y^*(t^*)$ such that any perturbation always pushes the curve to the same side of this hyperplane.

Step 9: We are now ready to tie up the sack for the fixed-endpoint problem by showing that the Hamiltonian has the properties stated in the maximum principle.

Hamiltonian maximization: From Step 7 we know that

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

which by the adjoint property of previous step gives that

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

for any spatial perturbation. But $\nu_b(w)$ was defined as

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

hence

$$\underbrace{\left\langle \begin{pmatrix} p_0^* \\ p^*(b) \end{pmatrix}, \begin{pmatrix} L(x^*(b), w) \\ f(x^*(b), w) \end{pmatrix} \right\rangle}_{H(x^*(b), w, p^*(b), p_0^*)} \leq \underbrace{\left\langle \begin{pmatrix} p_0^* \\ p^*(b) \end{pmatrix}, \begin{pmatrix} L(x^*(b), u^*(b)) \\ f(x^*(b), u^*(b)) \end{pmatrix} \right\rangle}_{H(x^*(b), u^*(b), p^*(b), p_0^*)}.$$

This shows that $u^*(b)$ maximizes $H(x^*(b), \cdot, p^*(b), p_0^*)$ for all b such that u^* is continuous at b. However, H is continuous in both x^* and p^* which are continuous in t. Hence, $H(x^*(t), \cdot, p^*(t), p_0^*)$ is continuous in time. Now since u^* is piecewise continuous with right or left limits the Hamiltonian will be maximized by $u^*(t)$ for all $t \in [t_0, t^*]$.

$H\big|_* \equiv 0$: For temporal perturbations we have

$$\delta(\tau) = \begin{pmatrix} L\big|_*(t^*) \\ f\big|_*(t^*) \end{pmatrix} \tau, \quad \text{for } \tau \in \mathbb{R}.$$

By the separation property we must then have that

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \begin{pmatrix} L\big|_*(t^*) \\ f\big|_*(t^*) \end{pmatrix} \tau \right\rangle \leq 0, \quad \text{for } \tau \in \mathbb{R}.$$

Hence,

$$H\big|_*(t^*) = \left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \begin{pmatrix} L\big|_*(t^*) \\ f\big|_*(t^*) \end{pmatrix} \right\rangle = 0.$$

We now show that $H\big|_*$ is continuous in t and has time-derivative 0. Assume that u^* is discontinuous at t_c and let $t \nearrow t_c$. By the Hamiltonian maximization property, and since x^* and p^* are both continuous in t, we have

$$H(x^*(t_c), u^*(t_c^-), p^*(t_c), p_0^*) \geq H(x^*(t_c), \underbrace{u^*(t_c^+)}_{w}, p^*(t_c), p_0^*).$$

Letting $t \searrow t_c$ gives us the opposite relation which finishes the proof of continuity.

To compute the time derivative of $H\big|_*$ we must be careful since H_u might not exist. We write $H\big|_*(t) = m(x^*(t), p^*(t))$, where

$$m(x, p) := \max_{u \in U} H(x, u, p, p_0^*).$$

First note that $m^*(x^*(t), p^*(t))$ is continuous in t. Let t and t' be two times

$$\begin{aligned} H(x^*(t'), u^*(t), p^*(t'), p_0^*) - H\big|_*(t) &\leq m(x^*(t'), p^*(t')) - m(x^*(t), p^*(t)) \\ &\leq H\big|_*(t') - H(x^*(t), u^*(t'), p^*(t), p_0^*) \end{aligned}$$

Dividing by $t' - t$ and letting $t' \searrow t$ we get

$$\begin{aligned} \lim_{t' \searrow t} \frac{m(x^*(t'), p^*(t')) - m(x^*(t), p^*(t))}{t' - t} &\geq \lim_{t' \searrow t} \frac{H(x^*(t'), u^*(t), p^*(t'), p_0^*) - H\big|_*(t)}{t' - t} \\ &= \langle H_x\big|_*, \dot{x}^* \rangle + \langle H_p\big|_*, \dot{p}^* \rangle \\ &= \langle H_x\big|_*, H_p\big|_* \rangle + \langle H_p\big|_*, -H_x\big|_* \rangle = 0. \end{aligned}$$

With $t' \nearrow t$ we get the opposite relation and thus have that $\frac{dm}{dt}\Big|_* = 0$. Hence, $\dot{H}\Big|_* = 0$ a.e. and $H\big|_*(t^*) = 0$ which shows that $H\big|_* \equiv 0$.

This finishes the proof of the maximum principle for the fixed-endpoint case!

Step 10: In the variable-endpoint setting the constraint $x(t_f) = x_1$ is changed to $x(t_f) \in S_1$. We thus have a contradiction of optimality if $y(t_f) = \begin{pmatrix} x^0(t_f) \\ x(t_f) \end{pmatrix}$, with $x(t_f) \in S_1$ and $x^0(t_f) < x^{0,*}(t_f)$.

We let D denote the set of all $y = \begin{pmatrix} x^0 \\ x \end{pmatrix}$, with $x \in S_1$ and $x^0 \leq x^{0,*}(t_f)$. We also introduce a linear approximation of this set

$$T := \{y \in \mathbb{R}^{n+1} : y = y^*(t^*) + \begin{pmatrix} 0 \\ d \end{pmatrix} + \beta\mu, \text{ with } d \in T_{x^*(t^*)}S_1, \text{ and } \beta \geq 0\}.$$

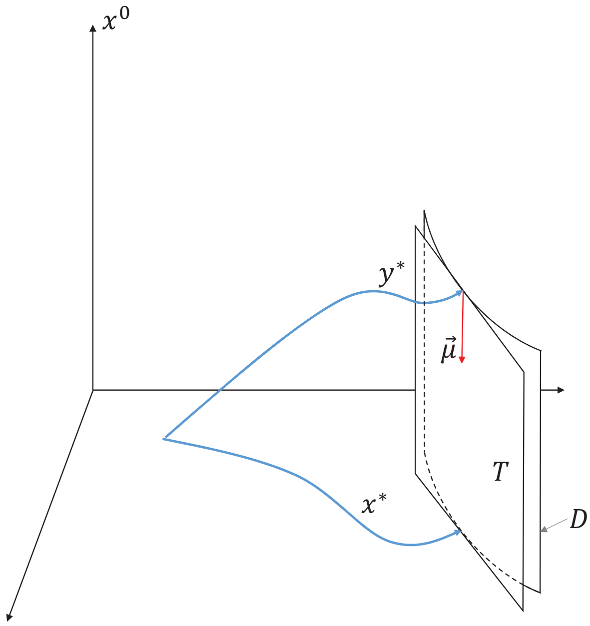

FIGURE 8. The sets D and T in the transversality condition.

Lemma 4.2. *The intersection of T and int C_{t^*} is empty.*

To prove this lemma we choose an ε-ball around $\hat{y}_\varepsilon = y^*(t^*) + \varepsilon \begin{pmatrix} 0 \\ d \end{pmatrix} + \varepsilon\beta\mu$, for some $d \in T_{x^*(t^*)}S_1$ and $\beta > 0$. Since T and D are tangent in $y^*(t^*)$ and hence the difference between T and D along $\varepsilon \begin{pmatrix} 0 \\ d \end{pmatrix}$ is $o(\varepsilon)$, $\tilde{B}_\varepsilon$ will meet D for sufficiently small ε. □

Following the same reasoning as in Step 7 we find that

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, y - y^*(t^*) \right\rangle \geq 0, \quad \forall y \in T,$$

which specifically means that

$$\left\langle \begin{pmatrix} p_0^* \\ p^*(t^*) \end{pmatrix}, \begin{pmatrix} 0 \\ d \end{pmatrix} \right\rangle = \langle p^*(t^*), d\rangle \geq 0, \quad \forall d \in T_{x^*(t^*)}S_1.$$

Now, if $d \in T_{x^*(t^*)}S_1$ then also $-d \in T_{x^*(t^*)}S_1$ and we arrive at the transversality condition

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

Note that in the special case when $S_1 = \mathbb{R}^n$ we have $T_{x^*(t^*)}S_1 = \mathbb{R}^n$ and thus $p^*(t^*) = 0$.