Optimal Control Theory - Module 3 - Maximum Principle

Fall, 2015 - University of Notre Dame

7.1 - Statement of Maximum Principle

Consider the problem of minimizing

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

subject to $(t_f, x(t_f)) \in S = [t_0, \infty) \times S_1$ where S_1 is a k dimensional manifold in $\mathbb{R}^n$

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

where h_i are $\mathcal{C}^1$ functions from $\mathbb{R}^n$ to $\mathbb{R}$ subject to

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

for $u \in \mathcal{C}[t_0, T]$ and $u(t) \in U \subset \mathbb{R}^m$ with f and L being $\mathcal{C}^1$ functions.

Let $u^* : [t_0, t_f] \to \mathbb{R}$ be an optimal control with state trajectory $x^* : [t_0, t_f] \to \mathbb{R}^n$ and a constant. Then there exists a function $p^* : [t_0, t_f] \to \mathbb{R}^n$ and a constant $p_0^* \leq 0$ (not both zero) for all $t \in [t_0, t_f]$ such that

1. x^* and p^* satisfy Hamilton's canonical equations,

$$\begin{aligned} \dot{x}^* &= H_p(x^*, u^*, p^*, p_0^*) \quad (1) \\ \dot{p}^* &= -H_x(x^*, u^*, p^*, p_0^*) \quad (2) \end{aligned}$$

   with $x^*(t_0) = x_0$ and $x^*(t_f) \in \mathcal{S}_1$ where

$$H(x, u, p, p_0) = \langle p, f(x, u)\rangle + p_0 L(x, u) \quad (3)$$

2. For each $t \in [t_0, t_f]$ and $u \in \mathcal{U}$

$$H(x^*, u^*, p^*, p_0^*) \geq H(x^*, u, p^*, p_0^*) \quad (4)$$

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

4. The vector $p^*(t_f)$ is orthogonal to the tangent space of S_1 at $x^*(t_f)$. In other words,

$$\langle p^*(t_f), d\rangle = 0 \text{ for all } d \in T_{x^*(t_f)} S_1 \quad (5)$$

Equation (4) is called the *maximum principle*, *Pontryagin's Maximum Principle* or PMP for short. Equation (5) is a transversality condition. Equation (2) in the pair of Hamilton's equations is often called the co-state or adjoint equation.

The following comments are in order

- The above itemized conditions provide a necessary condition for optimality.
- If we let the abnormal multiplier $p_0^* = -1$, then we obtain our earlier Hamiltonian.
- $H = 0$ is a unique feature of the free-time problem.

- We assumed time-invariance in the plant. But this is not overly restrictive since time can be treated as another state, x_{n+1}, whose state equation is $\dot{x}_{n+1} = 1$ with initial condition $x_{n+1}(0) = 0$.
- The tangent space found in the transversality condition (5) may be written as

$$T_{x^*(t_f)} S_1 = \{d \in \mathbb{R}^n \, : \, \langle \nabla h_i(x^*(t_f)), d \rangle = 0 \text{ for } i = 1, 2, \ldots, n-k\}$$

- Note that if $S_1 = \{x_1\}$ (fixed endpoint problem) then the tangent space is the empty set so that $\langle p^*(t_f), d \rangle = 0$ for any $p^*(t_f)$.

7.2 Proof of Maximum Principle - Temporal and Needle Perturbations

We start by transforming the original problem into Mayer form. This is done by defining an additional variable $x^0 \in \mathbb{R}$ that satisfies the differential equation

$$\dot{x}^0 = L(x,u), \quad x^0(t_0) = 0 \tag{6}$$

with the augmented system being

$$\begin{array}{c} \dot{x}^0 = L(x,u) \\ \dot{x} = f(x,u) \end{array} , \quad \left[\begin{array}{c} x^0(t_0) \\ x(t_0) \end{array} \right] = \left[\begin{array}{c} 0 \\ x_0 \end{array} \right] \tag{7}$$

The cost function may now be rewritten as

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

which is clearly in Mayer form. For notational covenience we let $y = \left[\begin{array}{c} x^0 \\ x \end{array} \right]$ for which the state equation becomes,

$$\dot{y} = \left[\begin{array}{c} L(x,u) \\ f(x,u) \end{array} \right] = g(y,u)$$

The target set takes the form $S = [t_0, \infty) \times S_1'$.

It will be convenient to view the augmented state y as shown below in figure 1. In this figure, the x^0 state is on the axis point upwards and the other states are in the horizontal plane. Figure 1 shows the situation for the fixed endpoint problem in which the target set $S_1 = \{x_1\}$. The resulting S_1' state, then is a vertical line that hits the horizontal plane at $\left[\begin{array}{c} 0 \\ x_1 \end{array} \right]$. If we are given a trajectory y^* that is optimal, then since $(x^0)^*$ represents the optimal cost, any admissible perturbation of y^* must hit the S_1' manifold higher up. It is impossible for the perturbed trajectory y to hit S' below $y^*(t^*)$ (see right hand side of figure 1).

The basic plan of the proof is as follows. We first introduce a very special set of perturbations on the control u and the terminal time t_f. These perturbations are used to show that the cone of "admissible perturbations" (we earlier referred to this as the cone of feasible directions) cannot contain a "descent direction" for the problem. In other words, we adopt a geometric characterization of optimality very similar to our geometric characterizations for the extremals of NLP problems. We use this geometric characterization to verify that all of the necessary conditions listed above must be satisfied by a point at the vertex of this cone.

Let's first consider the perturbed due to varying the terminal time. These controls are written as

$$u_\tau(t) := u^*(\min\{t, t^*\}), \quad t \in [t_0, t^* + \epsilon\tau] \tag{8}$$

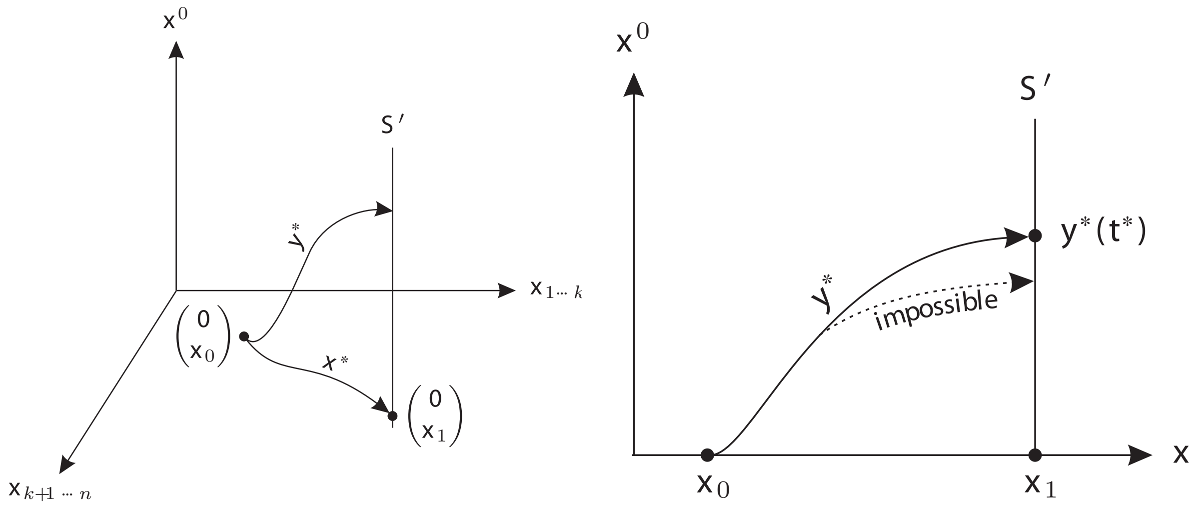

Figure 1: Trajectory y^* and its perturbation

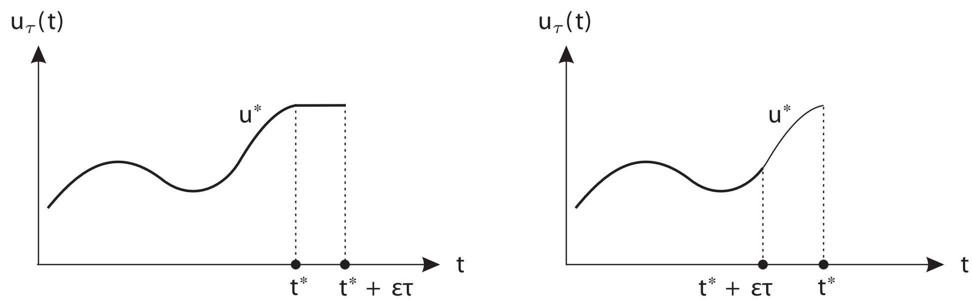

Figure 2: temporal perturbation of control

The perturbed control is graphically shown in figure 2 for $\tau > 0$ and $\tau < 0$. Note that for $\tau > 0$, we must extend u by simply keeping the same value it had at time t^*.

The impact that this temporal perturbation has on the state, y, trajectory may be written as

$$\begin{aligned} y(t^* + \epsilon\tau) &= y^*(t^*) + \dot{y}(t^*)\epsilon\tau + o(\epsilon) \\ &= y^*(t^*) + g(y^*(t^*), u^*(t^*))\epsilon\tau + o(\epsilon) \\ &= y^*(t^*) + \epsilon\delta(\tau) + o(\epsilon) \end{aligned} \tag{9}$$

For $\tau < 0$, we see $y(t^* + \epsilon\tau) = y^*(t^* + \epsilon\tau)$, so that the first-order series expansion of y^* about t^* is equivalent to y^*. The vector $\epsilon\delta(\tau)$ describes the infinitesimal first-order perturbation of the terminal point's impact on y^*. As we vary τ over $\mathbb{R}$, keeping ϵ fixed, the points $y^*(t^*) + \epsilon\delta(\tau)$ form a line through $y^*(t^*)$. This line will be denoted as $\vec{\rho} = \epsilon\delta(\tau)$.

The other variation we need is a *spatial variation*. This is often called a *needle perturbation*. Let $w \in \mathcal{U}$ be arbitrarily chosen and consider the interval $I = [b - \epsilon a, b] \in (t_0, t^*)$. We let $a > 0$ be arbitrary and $\epsilon > 0$ is infinitesimally small. We define the needle perturbation to the control as

$$u_{w,I}(t) := \begin{cases} u^*(t) & \text{if } t \notin I \\ w & \text{if } t \in I \end{cases} \tag{10}$$

Figure 3 shows the needle perturbation and its impact on the state trajectory.

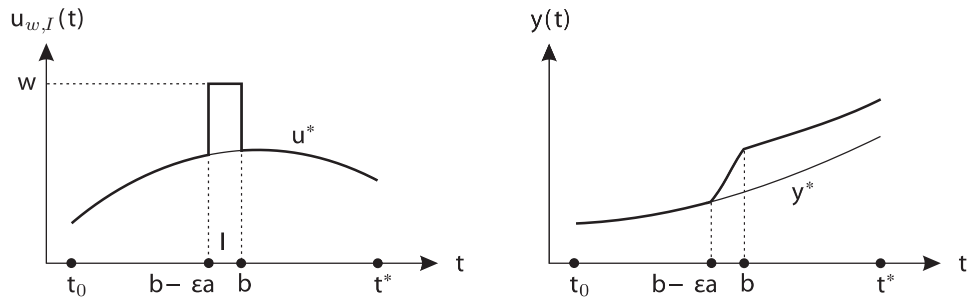

Figure 3: Needle Perturbation

Let $\approx$ denote equality up to $o(\epsilon)$. The first order Taylor series of the optimal trajectory y^* about $t = b$ is

$$y^*(b - \epsilon a) \approx y^*(b) - \dot{y}^*(b)\epsilon a$$

which implies that

$$y^*(b) \approx y^*(b - \epsilon a) + g(y^*(b), u^*(b))\epsilon a \tag{11}$$

A similar perturbation of the trajectory y at $t = b$ yields,

$$y(b) \approx y(b - \epsilon a) + g(y^*(b - \epsilon a), w)\epsilon a \tag{12}$$

Taylor series expansion of the last term in equation (12) yields,

$$g(y^*(b - \epsilon a), w)\epsilon a \approx g(y^*(b), w)\epsilon a + g_y(y^*(b), w)\left(y^*(b - \epsilon a) - y^*(b)\right)\epsilon a \tag{13}$$

Inserting this back into equation (**??**) produces,

$$\begin{aligned} y(b) &\approx y(b - \epsilon a) + g(y^*(b), w)\epsilon a + g_y(y^*(b), w)\left(y^*(b - \epsilon a) - y^*(b)\right)\epsilon a \\ &\approx y(b - \epsilon a) + g(y^*(b), w)\epsilon a + g_y(y^*(b), w)\left(-\dot{y}^*(b)\right)(\epsilon a)^2 \end{aligned} \tag{14}$$

The last term in equation 14 can be neglected since it is a second order term in ϵ^2 and we are only focusing on first order variations. We may, therefore conclude that

$$y(b) \approx y^*(b) + \nu_b(w)\epsilon a \tag{15}$$

where

$$\nu_b(w) := g(y^*(b), w) - g(y^*(b), u^*(b)) \tag{16}$$

This focuses on the perturbation in y^* immediately due to the needle perturbation. In other words, $y(b)$, in equation (15) is the variation immediately after the needle perturbation is over. This difference can be clearly seen in figure 3. We now need to consider what happens to y after $t = b$.

For $t > b$, we write the perturbed trajectory as

$$y(t) = y^*(t) + \epsilon\psi(t) + o(\epsilon) = y(t, \epsilon) \tag{17}$$

for $b \leq t \leq t^*$. $\psi(t)$ is the perturbation in the y-trajectory due to the perturbed control. From our preceding analysis, we know that

$$\psi(b) = \nu_b(w)a \tag{18}$$

Also if we differentiate $y(t, \epsilon)$ with respect to ϵ we see that

$$\psi(t) = y_\epsilon(t, 0) \tag{19}$$

Let's rewrite $y(t, \epsilon)$ as an integral equation,

$$y(t, \epsilon) = y(b, \epsilon) + \int_b^t g(y(s, \epsilon), u^*(s))ds \tag{20}$$

and then differentiate $y(t, \epsilon)$ with respect to ϵ. This yields,

$$\begin{aligned} y_\epsilon(t, 0) &= \nu_b(w)a + \int_b^t g_y(y(s, 0), u^*(s))y_\epsilon(s, 0)ds \\ &= \nu_b(w)a + \int_b^t g_y(y^*(s), u^*(s))\psi(s)ds \end{aligned} \tag{21}$$

Taking the derivative of $\psi(t)$ with respect to time and using equation (21) we see that

$$\dot{\psi}(t) \;=\; g_y(y^*, u^*)\psi = g_y|_* \, \psi = A_*(t)\psi(t) \tag{22}$$

Recall that$y$ is the augmented state $\begin{bmatrix} x^0 \\ x \end{bmatrix}$. So let $\psi = \begin{bmatrix} \eta^0 \\ \eta \end{bmatrix}$ be the corresponding components of the spatial control perturbation ψ. From equation (22) we can now see that

$$\dot{\eta}^0 \;=\; (L_x)^T\Big|_* \eta \tag{23}$$

$$\dot{\eta} \;=\; f_x|_* \, \eta \tag{24}$$

which implies that the time-varying map, $A_*(t)$, from equation (**??**) may be written as

$$A_*(t) = \begin{bmatrix} 0 & (L_x)^T\Big|_* \\ 0 & f_x|_* \end{bmatrix} \tag{25}$$

which is a linear time-varying system matrix.

Let $\Phi_*(\cdot, \cdot)$ be the state transition matrix for A_*, then the perturbation at t^* is

$$\begin{aligned} \psi(t^*) &= \Phi_*(t^**, b)\psi(b) \\ &= \Phi_*(t^*, b)\nu_b(w)a \end{aligned}$$

which implies that the perturbed state trajectory at time t^* is

$$\begin{aligned} y(t^*) &= y^*(t^*) + \epsilon\psi(t^*) + o(\epsilon) \\ &\approx y^*(t^*) + \epsilon\Phi_*(t^*, b)\nu_b(w) \end{aligned} \tag{26}$$

and for notational convenience we let

$$\delta(w, I) = \Phi_*(t^*, b)\nu_b(w) \tag{27}$$

The term $\delta(w, I)$ therefore represents the first order variation to y^* after b due to the needle perturbation. To clearly define this, of course, δ, needs w (the strength of the needle perturbation) and I (the interval over which the perturbation acted).

Figure 11 illustrates the evolution of ψ from b up until the terminal time t^*.

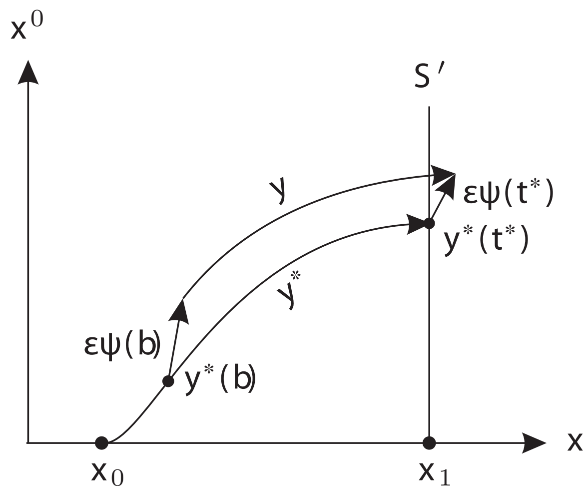

Figure 4: Impact of needle perturbation on y^* over the interval $[b, t^*]$.

In what follows, we want to consider the impact of multiple needle perturbations on the state trajectory y^* as shown in figure 5. In particular, let's concantenate two needle perturbations as shown in figure 5. At the end of the first needle perturbation we see

$$y(b_1) \approx y^*(b_1) + \nu_{b_1}(w_1)\epsilon a_1 \tag{28}$$

At the end of the second needle perturbation we have

$$y(b_2) \quad = \quad y^*(b_2) + \epsilon\left(\Phi_*(b_2, b_1)\nu_{b_1}(w_1)a_1 + \nu_{b_2}(w_2)a_2\right) + o(\epsilon) \tag{29}$$

Since the state transition matrix has the semi-group property, we know that $\Phi_*(t^*, b_2)\Phi_*(b_2, b_1) = \Phi_*(t^*, b_1)$ and we can rewrite $y(t^*)$ as

$$\begin{aligned} y(t^*) &\approx y^*(t^*) + \epsilon\Phi_*(t^*, b_2)\left(\Phi_*(b_2, b_1)\nu_{b_1}(w_1)a_1 + \nu_{b_2}(w_2)a_2\right) \\ &\approx y^*(t^*) + \epsilon\Phi_*(t^*, b_1)\nu_{b_1}(w_1)a_1 + \epsilon\Phi_*(t^*, b_2)\nu_{b_2}(w_2)a_2 \qquad (30) \\ &\approx y^*(t^*) + \epsilon\delta(w_1, I_1) + \epsilon(w_2, I_2) \qquad (31) \end{aligned}$$

In other words, the concatenated needle perturbations add up and so if we want to generate a perturbation on $y*$ in the direction $\epsilon\beta_1\delta(w_1, I_1) + \epsilon\beta_2\delta(w_2, I_2)$ for some $\beta_1, \beta_2 > 0$, we just adjust the interval length so that

$$\begin{aligned} u(t) = w_1 \text{ on } \overline{I}_1 = (b_1 - \epsilon a\beta_1 a_1, b_1] \\ u(t) = w_2 \text{ on } \overline{I}_2 = (b_2 - \epsilon\beta_2 a_2, b_2] \end{aligned}$$

This construction applies to any linear combination, so we can concatenate arbitrary needle perturbations to obtain

$$y(t_f) = y^*(t^*) + \epsilon\left(\beta_0\delta(\tau) + \sum_{i=1}^{m}\beta_i\delta(w_i, I_i)\right) \tag{32}$$

The first term represents the effect of the temporal perturbation on y and the second term in equation (32) represents the perturbation of multiple needle perturbations on the state trajectory. The proof of the Maximum Principle views these needle perturbations as forming a "cone" of feasible directions from which a geometric necessary condition for optimality will be obtained.

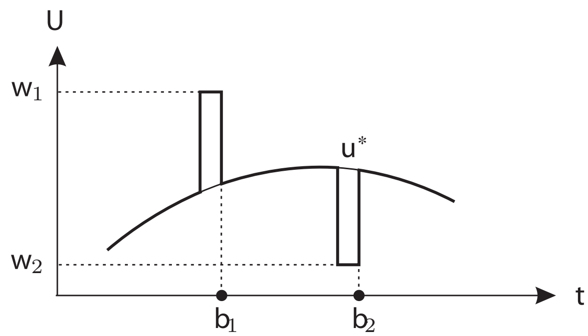

Figure 5: Concatenated Needle Perturbations

7.3 - Proof of the Maximum Principle - establishing the maximum principle

The previous lecture showed that the perturbed state trajectory y at time t_f can be written as the linear combination of a temporal perturbation and several needle perturbations.

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

Let $\vec{\rho}\ (w, b)$ denote the a direction of a perturbation due to the needle perturbation of strength w on interval I. The set of all such rays corresponding to all simple needle perturbations will be denoted as $\vec{P}$ and we let the convex cone of this set be denoted as co $\vec{P}$. Note that co $\vec{P}$ does not yet include the temporal variation $\delta(\tau)$. The temporal perturbations define a line, $\vec{\rho}$ going though $y^*(t^*)$. Adding these directions, we obtain the *terminal cone*, C_{t^*} whichis shown below in figure 6.

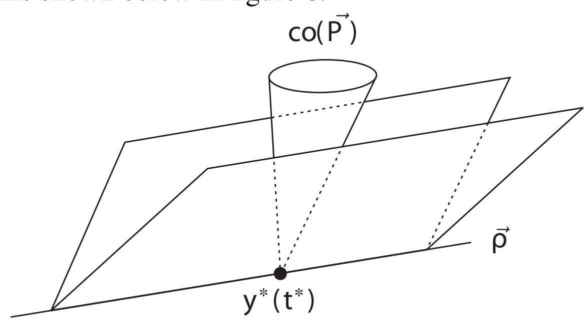

Figure 6: The Terminal Cone, C_{t^*}

Recall that the x^0 axis shown in figure 1 represents the *cost* of y. If $y^*(t^*)$ is indeed optimal, we would expect the vector

$$\vec{\mu} = \begin{bmatrix} -1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \in \mathbb{R}^{n+1} \tag{33}$$

to be directed outside of the terminal cone, C_{t^*}. Just as it would in our earlier geometric characterizations of optimality for finite-dimensional NLP problems. The following lemma establishes this geometric characterization of optimality for infinite-dimensional spaces.

Lemma: The vector $\vec{\mu}$ defined in equation (33) does not intersect the interior of the terminal cone, C_{t^*}.

We now outline the proof for this lemma. Let's assume the statement is false. This would mean we can pick a point $\hat{y}$ on the ray $\overrightarrow{\mu}$ below $y^*(t^*)$ such that $\hat{y} \in C_{t^*}$ along with an open ball $N_\epsilon(\hat{y})$ around it. Clearly $\hat{y}$ may be written as

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

for some $\beta > 0$ where μ is the vector defined by the ray $\overrightarrow{\mu}$ defined in equation (33). But if the lemma's statement is false, any point in $N_\epsilon(\hat{y} \in C_{t^*}$ which would mean that

$$\begin{aligned} y &= y^*(t^*) + \epsilon\left(\beta_0\delta(\tau) + \sum_{i=1}^{m}\beta_i\delta(w_i, I_i)\right) \\ &= y^*(t^*) + \epsilon\nu \end{aligned}$$

where we've introduced the vector ν, for notational convenience.

The actual terminal state, $y(t_f)$, due to the needle perturbations, however, is

$$y_{\text{actual}} = y^*(t^*) + \epsilon\nu + o(\epsilon)$$

We let $N_\epsilon(\hat{y})$ denote an ϵ-neighborhood about $\hat{y}$ and we let $B_\epsilon(\hat{y})$ be the corresponding set of points obtained by needle perturbations.

Even through N_ϵ always lies below $y^*(t^*)$ for sufficiently small ϵ. Is it possible that the $o(\epsilon)$ perturbations in y_{actual} cause some points of $B_\epsilon(\hat{y})$ to lie in the terminal cone as shown in figure 7.

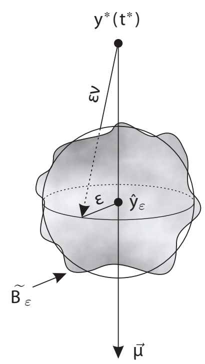

Figure 7: The difference between $N_\epsilon(\hat{y})$ and $B_\epsilon(\hat{y})$

Since $\lim_{\epsilon\to 0}\frac{o(\epsilon)}{\epsilon} = 0$, this implies that the $o(\epsilon)$ variation is less than ϵ if ϵ is small enough. So for sufficiently small ϵ, the ball B_ϵ must lie completely below $y^*(t^*)$, which contradicts the assertion that the lemma is false. $\square$

Since C_{t^*} is convex and doesn't intersect $\overrightarrow{\mu}$, we can use the separation theorem to find a hyperplane separating $y^*(t^*)$ from the $\overrightarrow{\mu}$. This means there exists a vector $\begin{bmatrix} p_0^* \\ p^*(t^*) \end{bmatrix} \in \mathbb{R}^{n+1}$ such that the hyperplane

$$H = \left\{ y \in \mathbb{R}^{n+1} \;:\; \left\langle \begin{bmatrix} p_0^* \\ p^*(t^*) \end{bmatrix}, y \right\rangle = \left\langle \begin{bmatrix} p_0^* \\ p^*(t^*) \end{bmatrix}, y^* \right\rangle \right\}$$

is seperating C_{t^*} from $\overrightarrow{\mu}$. In particular this separation property my be formally written as

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

for all δ such that $y^*(t^*) + \delta \in C_{t^*}$ and

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

Given the form of μ in equation (33) we can see that this last inequality ipmies that $p_0^* \leq 0$, which is one of the conditions found in our formal statement of the maximum principle.

We now derive the co-state equations. Before doing this, let's recall that two linear systems $\dot{x} = Ax$ and $\dot{z} = -A^T z$ are said to be *adjoint* to each other. In particular, the solutions x and z of these two equations must satisfy

$$\begin{aligned} \frac{d}{dt}\langle x, z\rangle &= \langle \dot{z}, x\rangle + \langle z, \dot{x}\rangle \\ &= \left(-A^T z\right)^T x + z^T A x = 0 \end{aligned}$$

We now consider a specific pair of adjoint systems on the time interval $[t_0, t^*]$.

The first system will be the variational equation for the needle perturbation

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

that we stated earlier in equation (25) where

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

The second system is the adjoint of the first. Namely,

$$\frac{d}{dt}\begin{bmatrix} p_0 \\ p \end{bmatrix} = \dot{z} = -A_*^T z = \begin{bmatrix} 0 & 0 \\ -\left.L_x\right|_* & -\left.(f_x)^T\right|_* \end{bmatrix} z$$

These equations imply that $\dot{p}_0 = 0$, so that p_0 is a constsant and

$$\begin{aligned} \dot{p} &= -\left.L_x\right|_* p_0 - \left.(f_x)^T\right|_* p \\ &= -H_x(x^*, u^*, p, p_0) \end{aligned}$$

for the terminal condition we let

$$z(t^*) = \begin{bmatrix} p_0^* \\ p^*(t^*) \end{bmatrix}$$

which is a vector determined by the separating hyperplane between the terminal cone C_{t^*} and $\vec{\mu}$. In other words, since p_0 is a constant equal to p_0^*, we can conclude that

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

which yields the second of Hamilton's canonical equations. So we've just verified the first major statement in the maximum principle.

We now proceed to establish the second assertion in our earlier statement of the maximum principle. This second assertion is why we call this the *maximum* principle. It asserts that the optimal control, u^*, maximizes the Hamiltonian when it is evaluated along an optimal trajectory. In other words, we want to verify that

$$H(x^*, u^*, p^*, p_0^*) \geq H(x^*, u, p^*, p_0^*)$$

Recall that a needle perturbation results in the following perturbation at t^*,

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

For notational convenience, the second term above was denoted as $\epsilon\delta(w, I)$. Since $y^*(t^*) \in C_{t^*}$, we know that

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

where Φ_* is the state transition matrix of the variation matrix A_* in equation (25). Recall, however, that for adjoint systems, $\frac{d}{dt}\langle z, x\rangle = 0$. So the inner product given above is constant over time and we can conclude at $t = b$ that

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

where

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

Inserting this expression for ν_b back into our earlier inner product, we see that

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

But since $H = p_0 L + pf$, we can easily see that the above inner product is identical to the Hamiltonian maximization principle at the point $t = b$.

$$H(x^*(b), w, p^*(b), p_0^*) \leq H(x^*(b), u^*(b), p^*(b), p_0^*)$$

The choice of b, however, is arbitrary. So this is precisely the maximum principle found in our original statement of the theorem.

7.4 - Proof of the Maximum Principle - Transversality Conditions

The previous lecture proved those statements in our formulation of the maximum principle characterizing why we call it the maximum principle. Namely that the Hamiltonian, along the optimal state trajectory, is optimized when we use u^*. We now proceed to establish the transversality conditions found in our original statement of the maximum principle. We'll divide our discussion into two parts. First we verify our statements regarding the fixed endpoint problem (free time) and then generalize to the transversality conditions required by the variable endpoint problem.

The separation property holds when we only have a temporal perturbation (i.e. $\delta(\tau) \in C_{t^*}$). We've already showed that

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

where τ can be positive or negative. The separation property requires that

$$\begin{aligned} 0 &= \left\langle \begin{bmatrix} p_0^* \\ p^*(t^*) \end{bmatrix}, \begin{bmatrix} L(x^*(t^*), u^*(t^*)) \\ f(x^*(t^*), u^*(t^*)) \end{bmatrix} \right\rangle \\ &= H(x^*(t^*), u^*(t^*), p^*(t^*), p_0^*) \end{aligned}$$

So we can conclude that under the optimal control the Hamiltonian is zero at the terminal time t^*.

We want to show, however, that $H \equiv 0$ for all t along the optimal trajectory. To do this we first establish that H is a continuous function of time and then show that $H|_*$ is constant.

To establish continuity, let t be any point of discontinuity in u^*. The optimal state and co-states, x^* and p^*, are continuous everywhere. Applying the Hamiltonian maximization property established in the previous lecture for $b < t$ and $w = u^*(t^*)$ and making b approach t from the left, we obtain

$$H(x^*(t), u^*(t^+), p^*(t), p_0^*) \leq H(x^*(t), u^*(t^-), p^*(t), p_0^*)$$

Repeating this argument with $b > t$ and $w = u^*(t^-)$ we also see that

$$H(x^*(t), u^*(t^*), p^*(t), p_0^*) \geq H(x^*(t), u^*(t^-), p^*(t), p_0^*)$$

which means that both limits are equal and we can conclude H is continuous.

To establish that H is constant, we cannot simply differentiate since H is only continuous (not necessarily differentiable). But in view of the maximum principle we can see that

$$\begin{aligned} H(x^*(t), u^*(t), p^*(t), p_0^*) &= \max_{u \in \mathcal{U}} H(x^*(t), u, p^*(t), p_0^*) \\ &\equiv m(x^*(t), p^*(t)) \end{aligned}$$

We adopt the notation $m(x^*(t), p^*(t))$ for notational simplicity and we just showed in the preceding paragraph that m is a continuous function of time t. So let t and t' be two arbitrary times and we have

$$H(x^*(t'), u^*(t'), p^*(t'), p_0^*) - H(x^*(t), u^*(t), p^*(t), p_0^*) \leq m(x^*(t'), p^*(t')) - m(x^*(t), p^*(t))$$

Note that H is $\mathcal{C}^1$ for each fixed $\bar{t} \in [t_0, t^*]$ which means it is also locally Lipschitz. Since it is locally Lipschitz, the function is absolutely continuous and hence differentiable.

Let's now consider the derivative of m. In particular, consider

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

You can use the maximum principle to show that this must be zero. (Homework problem). Since we know H at t^* and we know H is continuous, then because m' is zero for all t we can conclude $H = 0$ for all time. This therefore completes our characterization of the conditions found in the fixed endpoint version of the maximum principle.

Turning to the variable endpoint problem, we must establish a *transversality condition*. In the variable endpoint problem we again need to revisit the topological lemma discussed earlier. Recall that we proved this lemma by showing that a contradiction to optimality results if our terminal point has a cost lower than $x^{0,*}(t^*)$ whose x component is in S_1. In the variable endpoint problem, the manifold that our terminal point reaches is no longer a simple hyperplane. Let D be the set of all points in S_1 that have a lower cost than $x^(0, *)(t^*)$. A linear approximation to D is given by the tangent space of S_1. We denote this approximation as T and define it as

$$T = \left\{ y \in \mathbb{R}^{n+1} \,:\, y = y^*(t^*) + \begin{bmatrix} 0 \\ d \end{bmatrix} + \beta u,\, d \in T_{x^*(t^*)} S_1,\, \beta \geq 0 \right\} \tag{34}$$

The geometry is shown below on the left side of figure 8.

We use the same construction as before, where we select a point $\hat{y}$ on T and consider an open ball $N_\epsilon(\hat{y})$ which is completely below the top line shown on the right side of figure 8. The actual states generated by needle perturbations are $B_\epsilon(\tilde{y}) \approx N_\epsilon(\hat{y})$. Since the difference between both sets is $o(\epsilon)$, we can expect to

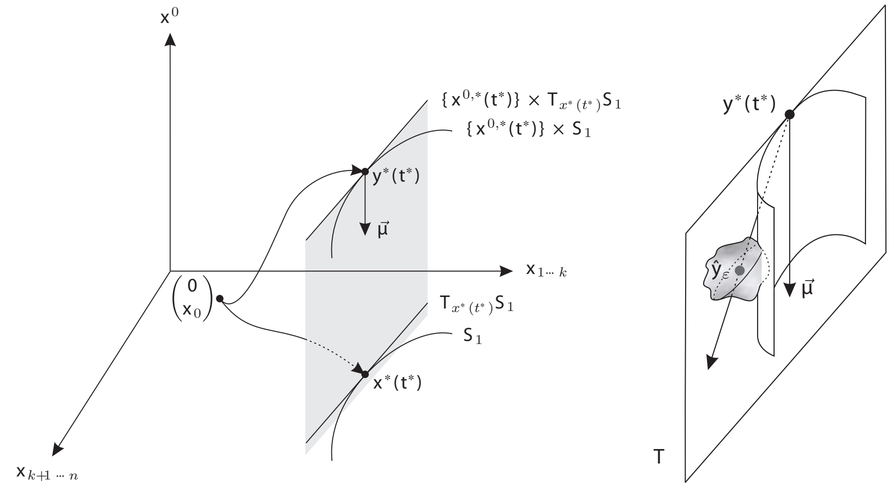

Figure 8: Geometry beneath the topological lemma for the variable endpoint problem (left), constructing a ball that has lower cost (right)

$B_\epsilon(\tilde{y})$ to also be below the optimal cost for a small enough ϵ. This is precisely the same type of reason that lead to our topological lemma for the fixed endpoint problem. We can therefore conclude that T *does not intersect the interior of the terminal cone* C_{t^*}. As before, we invoke the separation lemma to assert the existence of a hyperplane separating T and C_{t^*}. We let this hyper plane be characterized by the vector $p(t^*) = \begin{bmatrix} p_0^* \\ p^*(t^*) \end{bmatrix}$. The separation assertion requires

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

For any d in this tangent space, $T_{x^*(t^*)}S_1$, we also know that $-d \in T_{x^*(t^*)}S_1$ and so we can conclude that

$$\langle p(t^*), -d \rangle \leq 0$$

Since the choice of d was arbitrary, we can conclude for $d \in T_{x^*(t^*)}S_1$ that

$$\langle\langle p(t^*), d \rangle = 0$$

which is the transversality condition for the variable endpoint problem stated in our theorem.

7.5 - Other versions of the Maximum Principle

The previous lectures stated and proved the maximum principle for the fixed and variable endpoint problems. There are other useful problems, however, whose statements differ slightly from these two canonical problems. In particular we consider the case when the terminal time is fixed, when the system is time-varying, and when we have an additional terminal cost. For each of these cases, we can find a slight variation of the maximum principle that can be related back to our earlier proof.

Fixed Terminal Time: With a fixed terminal time, we see that the temporal perturbation $\delta(\tau)$ is no longer needed. As can be expected, nothing really changes in the analysis we gave earlier with the exception that we can no longer ensure $H = 0$ along the optimal trajectory. In fact, all we can guarantee is that the Hamiltonian is constant (though not necessarily zero) along the optimal trajectory.

This observation may be confirmed by transforming the fixed time problem into a free-time problem. Let's introduce a new variable $x_{n+1} = t$ with augmented differential equation constraints

$$\begin{aligned} \dot{x} &= f(x,u), \quad x(t_0) = x_0 \\ \dot{x}_{n+1} &= 1, \quad x_{n+1}(t_0) = t_0 \end{aligned}$$

The fixed time problem's target set is $\{t_1\} \times S_1$ whereas the new target set for our transformed problem is $[t_0, \infty) \times S_1 \times \{t_1\}$ with the terminal time no longer being explicitly controlled. The Hamiltonian for this reformulated problem is

$$\overline{H} = \langle p, f \rangle + p_{n+1} + p_0 L = H + p_{n+1}$$

This Hamiltonian must be zero along the optimal trajectory so that

$$H|_* + p^*_{n+1} = \overline{H}\big|_* = 0$$

which implies that the Hamiltonian for the original fixed-time problem must satisfy

$$H|_* = -p^*_{n+1}$$

which is constant, though not necessary zero

Time-varying Systems: Our original derivation of the maximum principle assume a time-invariant differential constraint $\dot{x} = f(x, u)$ and Lagrangian $L(x, u)$. In many applications, these objects may be functions of time, which would mean that the Hamiltonian is also a function of time,

$$H(t, x, u, p, p_0) := \langle p, f(t, x, u) \rangle + p_0 L(t, x, u)$$

We may add time as another state variable,

$$x_{n+1} := t$$

and augment the system equations to

$$\begin{array}{rcl} \dot{x} & = & f(x_{n+1}, x, u), \quad x(t_0) = x_0 \\ \dot{x}_{n_1} & = & 1, \quad x_{n+1}(t_0) = t_0 \end{array}$$

This is identical (in approach) to what we used for the fixed terminal time problem above. Based on our earlier analysis we see that

$$\dot{p}^*_{n+1} = -\left. H_t \right|_*$$

The difference is that this is not zero and p^*_{n+1} and hence $H|_* = -p^*_{n+1}$ are no longer constant. Instead we have a differential equation

$$\frac{d}{dt} \left. H \right|_* = \left. H_t \right|_*$$

with boundary condition $\left. H \right|_* (t_f) = -p^*_{n+1}(t_f)$. If t_f is free then x_{n+1} is free and the transversality condition becomes $p^*_{n+1}(t_f) = 0$.

Terminal Cost: Let's now assume that the cost functional has a terminal cost $K(x_f)$. For simplicity, we confine our attention to the Mayer problem where $L \equiv 0$. We use a transformation to reduce this to the basic variable endpoint problem

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

So this is an equivalent problem in Lagrange form with $L = \langle K_x, f \rangle$. For this problem the Hamiltonian is

$$\begin{array}{rcl} \overline{H} & = & \langle \overline{p}, f \rangle + p_0 \langle K_x, f \rangle \\ & = & \langle \overline{p} + p_0 K_x, f \rangle \end{array}$$

Applying the maximum principle, we obtain the co-state's differential equation,

$$\begin{aligned}\dot{\overline{p}}^* &= -\left.\overline{H}\right|_* \\ &= -\left.(f_x)^T\right|_* \overline{p}^* - p_0^* \left.K_{xx}\right|_* \left.f\right|_* - p_0^* \left.(f_x)^T\right|_* \left.K_x\right|_*\end{aligned}$$

with boundary condition $\overline{p}^*(t_f) = 0$.

8.1 - Simple Examples of the Maximum Principle

This lecture discusses some particular examples of the maximum principle.

Example 1: To start, let's consider a fixed time free-endpoint problem. The problem is to minimize

$$J(u) = m(x(1) - 1)^2 + \int_0^1 \sqrt{1 + u^2(t)}dt$$

where$m \geq 0$ is a positive constant, $u \in \mathcal{C}^1[0, T]$ and where $T = 1$ such that the state satisfies the state equation

$$\dot{x}(t) = u, \quad x(0) = 0$$

with $x(1)$ free. This is, therefore, a fixed-time optimal control problem in which the terminal point is free. We'll use the maximum principle to identify a candidate optimal control.

To use the Maximum Principle, we first need to define the problem's Hamiltonian. Based on our earlier discussion this function is

$$H(x, u, p) = -\sqrt{1 + u^2} + pu$$

(We've taken the abnormal multiplier p_0 to be -1).

We apply the maximum principle to determine u^*. In particular, the maximum principle allows us to immediately determine u^* as part of a point-wise maximization

$$\begin{aligned}u^*(t) &= \arg\max_{u \in \mathbb{R}} H(x^*(t), u(t), p^*(t)) \\ &= \arg\max_{u \in \mathbb{R}} \left(-\sqrt{1 + u^2} + p^*(t)u\right)\end{aligned}$$

Since u^* is chosen from the entire real line, $\mathbb{R}$ and since H is differentiable, we can use the first partial derivative of H with respect to u to identify a candidate u^*. In particular, we consider those u^* such that

$$H_u = 0 = -\frac{u^*}{\sqrt{1 + (u^*)^2}} + p^* \tag{35}$$

which implies that

$$p^*(t) = \frac{u^*(t)}{\sqrt{1 + (u^*(t))^2}} \tag{36}$$

for all $t \in [0, 1]$. This equation can be inverted to identify a control of the form,

$$u^*(t) = \pm \frac{p^*(t)}{\sqrt{1 - (p^*(t)(^2}}$$

So the control is determined once we know the co-state $p^*(t)$ for all $t \in [0, 1]$.

The co-state can be determined by solving Hamilton's second canonical equation

$$\dot{p}^* = -H_x = 0$$

which implies that

$$p^*(t) = \text{constant}$$

for all $t \in [0,1]$. In particular, since this is a fixed time problem, we can use the transversality condition to determine p^* at the terminal time. In other words, the transversality condition requires

$$p^*(1) = -2m(x^*(1) - 1) \Rightarrow x^*(1) = 1 - \frac{p^*(1)}{2m} \tag{37}$$

and so we know that $p^*(t) = p^*(1)$ for all $t \in [0,1]$. We can therefore conclude that the optimal control u^* is a *constant* that satisfies

$$u^* \quad = \quad \pm \frac{p^*(1)}{\sqrt{1 - (p^*(1))^2}}$$

Using the fact that u^* is constant and inserting equation (36) into (37), we see that the terminal state must satisfy

$$x^*(1) = 1 - \frac{u^*}{2m\sqrt{(u^*)^2 + 1}} \tag{38}$$

Finally, we know from the first Hamilton equation that

$$x^*(1) = \int_0^1 u^* dt = u^*$$

So inserting this into equation 38 we obtain

$$u^* = 1 - \frac{u^*}{2m\sqrt{(u^*)^2 + 1}} \Rightarrow 1 = u^* \left(1 + \frac{1}{2m\sqrt{(u^*)^2 + 1}} \right)$$

which we can then solve to determine u^*.

The following Matlab script was used to compute u^* for various values of the penalty weight, m.

```
xdat = [];
dm = .1;
mstop = 10;
for m=.01:dm:mstop;
ustar = fzero(@(u) exam1_func(u,m),0.5);
J = m*(ustar-1)^2 + sqrt(1+ustar^2);
xdat= [ xdat; [m ustar J]];
end;
figure(1);
plot(xdat(:,1),xdat(:,2),xdat(:,1),xdat(:,3));
axis([0 mstop 0 2]);
legend('ustar','Jstar');
xlabel('terminal penalty, m');
ylabel('control/cost');
```

The results from this script are shown in figure 9. From the figure, one can see that as m increases, the control effort increases as well. This is a reasonable expectation since larger m means that missing the desired terminal condition $x(1) = 0$ is larger, and hence a larger control, u^*, will be needed to bring it down. In the limit as m goes to zero, we see the optimal control go to zero as well. Which indicates that there is no penalty for a large $x(1)$. Since all of the cost if found in the running cost (which is only a function of u), the logical choice is to use no control to achieve the objective.

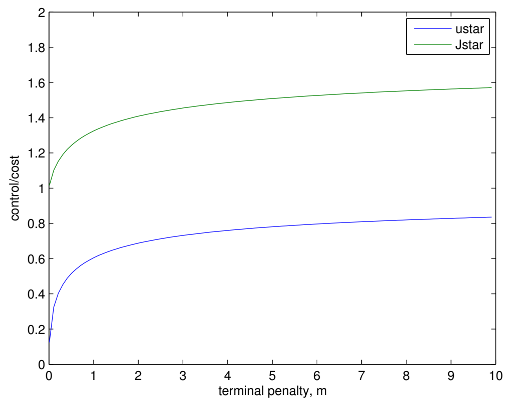

Figure 9: Example 1

Example 2: We now consider free time fixed endpoint problem. In this case, however, we assume that the controls are constrained to lie in the inteval $[0, 1]$. In particular, we want to minimize

$$J(u) = \int_0^T L(x,u,t)dt = \int_0^T \frac{u}{t}dt$$

where $u(t) \in [0, 1]$ for all $t \in [0, T]$ and the system states satisfy

$$\begin{aligned}\dot{x}_1 &= 1 - 2u \\ \dot{x}_2 &= -1\end{aligned}$$

with fixed initial and terminal conditions

$$\begin{bmatrix} x_1(0) \\ x_2(0) \end{bmatrix} = \begin{bmatrix} 0 \\ h \end{bmatrix}, \qquad \begin{bmatrix} x_1(T) \\ x_2(T) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

In this case the terminal time T is free.

We start by forming the Hamiltonian for this problem.

$$H(x,u,p) = -\frac{u}{t} - p_2 + p_1(1-2u)$$

Due to the constrained nature of u, we can no longer rely on the derivative of H with respect to u to see if we have an optimal point. So we directly look for a maximum. Fortunately, because the maximum principle does this maximization in a pointwise manner, it is relatively easy to see what the optimal u will be. In particular, the maximum principle requires that

$$\begin{aligned}u^*(t) &= \arg\max_{u\in[0,1]} H(x^*(t), u, p^*(t)) \\ &= \arg\max_{u\in[0,1]} \left\{ p_1^*(t) - p_2^*(t) - u\left(\frac{1}{t} + 2p_1^*(t)\right)\right\} \\ &= \arg\max_{u\in[0,1]} \left\{ -u\left(\frac{1}{t} + 2p_1^*(t)\right)\right\}\end{aligned}$$

Note that if $1 + 2tp_1^*(t)$ is positive then the optimal $u^*(t) = 1$. Otherwise if this term is negative we find that $u^* = 0$. In other words,

$$u^* = \begin{cases} 1 & 1 + 2tp_1^*(t) < 0 \\ 0 & 1 + 2tp_1^*(t) > 0 \end{cases}$$

We can rewrite this as

$$u^*(t) = -\frac{\mathrm{sgn}(1 + 2tp_1^*(t)) - 1}{2}$$

To determine u^* we need to find $p_1^*(t)$. We can do this from the co-state's differential equation.

$$-\begin{bmatrix} \dot{p}_1 \\ \dot{p}_2 \end{bmatrix} = \left.\frac{\partial H}{\partial x}\right|_* = 0$$

which implies that $p_1(t) = P_1 =$ constant and $p_2(t) = P_2 =$ constant for all $0 < t < T$. We therefore see that the optimal control must satisfy,

$$u^* = \begin{cases} 1 & \text{if } 1 + 2tP_1 < 0 \\ 0 & \text{if } 1 + 2tP_1 > 0 \end{cases}$$

We can think of the function

$$s(t) = 1 + 2tP_1$$

as a *switching* function that switches the control between 0 and 1. Note that $s(t)$ is a linear function of time. At time $t = 0$, we see that $s(0) = 1$. Depending on the sign of P_1, the switching curve either has a positive or negative slope as shown of the left side of figure 10. What we can see from this figure is that if $P_1 > 0$, then the switching curve never changes sign and if $P_1 < 0$ then there is at most one sign change. In other words, the optimal control for this system is a piecewise constant signal that has at most one switch. This switching time t_s can be determined from the equation $0 = s(t) = 1 + 2tP_1$.

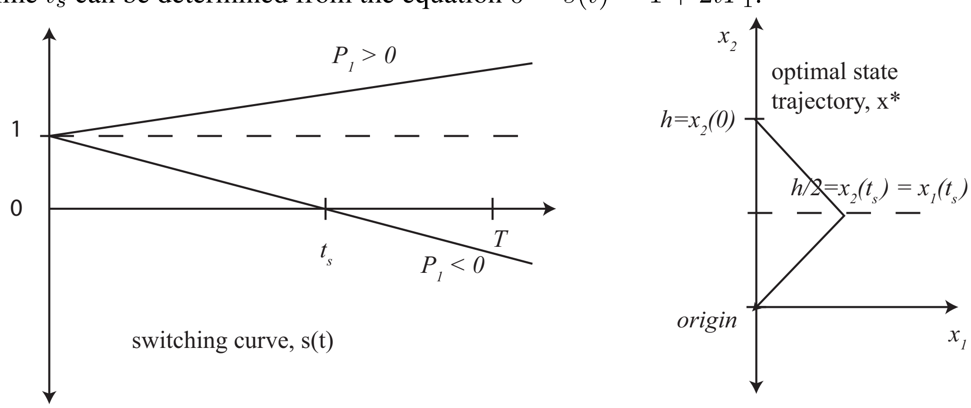

Figure 10: Example 2: switching function

If $P_1 \geq 0$, then we have an *inadmissible* control since $u^*(t) = 0$ for all t and we can never reach the terminal state at the origin. If $P_1 < 0$, then there is a single switching time occuring at $t_s = -\frac{1}{2P_1}$. At this time, the control switches from 0 to 1. To be admissible, the controlled system's state must reach the origin at time T. In other words,

$$x^*(t) = \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} = \begin{cases} \begin{bmatrix} t \\ h - t \end{bmatrix} & t \in [0, -\frac{1}{2P_1}] \\ \begin{bmatrix} t_s - (t - t_s) \\ h - t \end{bmatrix} & t \in [t_s, T] \end{cases}$$

If we let $t_s = h/2$, then $T = h$ and the desired terminal condition is reached under the optimal control. The system first coasts from $x_2(0) = h$ to $x_2(t_s) = h/2$ and then uses the control $u = 1$ to derive x_2 to zero by the terminal time T. The resulting optimal state trajectory is shown on the right side of figure 10.

The actual values of the co-state vector are obtained from the transversality condition. Since this is a free-time fixed-endpoint problem we know that the Hamiltonian H is zero for the entire trajectory. In particular at the terminal time, we know that

$$\begin{aligned} H(x^*(T), u^*(T), p^*(T)) &= -\frac{u^*(T)}{T} - P_2 + P_1(1 - 2u^*(T)) \\ &= -\frac{1}{h} - P_2^* - P_1^* = 0 \end{aligned}$$

which implies that

$$h = -\frac{1}{P_1^* + P_2^*}$$

This gives one equation constraining the costate. To find the other equation, recall that at the switching time $t_s = h/2$, that $s(t_s) = 0$, which means

$$1 + 2P_1^* \frac{h}{2} = 0 \rightarrow P_1^* = -\frac{1}{h}$$

So we can conclude that $P_1^* = -\frac{1}{h}$ and $P_2^* = 0$.

8.2 - Maximum Principle: double integrator and harmonic oscillator

The prior lecture considered examples where the state equation's right hand side $f(x, u)$ was only a function of the control input u. This greatly simplified the problem of determining the co-state, since $\dot{p}$ would equal zero and we could deduce that $p(t)$ was constant for all time. We now consider two simple problems in which $f(x, u)$ is a linear function of the state and control. The first system examines optimal control of a double integrator and the second examines the optimal control of a harmonic oscillator.

Minimum time Double Integrator: Consider the one dimensional problem of moving an object along a line. The coordinate x_1 measures the displacement of the object from its desired final position. The state of the system is determined by the object's position and velocity, so we need another state variable, x_2, for the velocity. The force acting on the object is controllable and is represented by the control variable u. We assume the control is bounded so that $|u| \leq 1$. We suppose further that the object is initially at a distance X_1 from home and is moving with velocity X_2. The objective is to find the control that moves the object to its home state (the origin) by the final time t_f. In other words, we require that $x_1(t_f) = x_2(t_f) = 0$. This problem, therefore has a fixed terminal point and a free terminal time.

The state equations for this system take the form

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= u \end{aligned} \tag{39}$$

The control is integrable and bounded so that

$$-1 \leq u(t) \leq 1$$

for all $t \in [0, t_f]$. The initial and final states are

$$\begin{aligned} x_1(0) = X_1, &\qquad x_2(0) = X_2 \\ x_1(t_f) = 0, &\qquad x_2(t_f) = 0 \end{aligned}$$

We first consider the cost functional

$$J(u, t_f) = \int_0^{t_f} dt = t_f$$

So this problem we first consider is to minimize the terminal time.

The Hamiltonian for this problem takes the form

$$H(x, u, p) = -1 + p_1 x_2 + p_2 u$$

The maximum principle implies

$$\begin{aligned} u^*(t) &= \arg \max_{u \in [-1,1]} H(x^*, u, p^*) \\ &= \arg \max_{u \in [-1,1]} \{-1 + p_1^*(t) x_2^*(t) + p_2^*(t) u\} \\ &= \arg \max_{u \in [-1,1]} \{p_2^*(t) u\} \end{aligned}$$

Using the same reasoning employed in the preceding example, we see that

$$u^*(t) = \begin{cases} -1 & \text{if } p_2^*(t) < 0 \\ 1 & \text{if } p_2^*(t) > 0 \end{cases}$$

So we can treat the costate $p_2^*(t)$ as the switching function. We determine p_2^* from its costate equations.

$$\begin{bmatrix} \dot{p}_1^* \\ \dot{p}_2^* \end{bmatrix} = - \left. H_x \right|_* = \begin{bmatrix} 0 \\ -p_1^* \end{bmatrix} \tag{40}$$

which implies that $p_1(t) = A$ (constant) and $p_2(t) = B - At$ where B is another constant of integration.

Note that $p_2^*(t)$ is a linear function of time. So we can again use the same reasoning employed in the preceding example to deduce that the optimal control exhibits at most one switch. We will use this fact to

To determine the switching time and the actual controls, however, we first need to identify the admissible controls (i.e. those that allow us to reach home from the initial condition).

To determine the admissible control, we trace solutions back from the terminal (home) state. Let t_s denote the switching time. If there is no switch we let $t_s = 0$. If the home state is approached with $u = +1$, then we know that

$$\begin{aligned} x_2(t) &= t - t_f \\ x_1(t) &= \frac{1}{2}(t - t_f)^2 \end{aligned}$$

for $t \in [t_s, t_f]$. The preceding equations represent a curve parameterized by t. We can remove this parameter by relating the two equations through $t - t_f$ and obtain

$$x_2^2(t) = 2x_1(t) \tag{41}$$

for $t \in [t_s, t_f]$. Note that x_2 is negative along this trajectory. If X_2 is negative and $X_2^2 = 2X_1$, then the control $u = 1$ can be applied without switching to reach the target in optimal time $t_f = -X_2$. Otherwise, we must switch to $u = -1$ at some time $t_s \in [0, t_f]$. As no further switching can take place, the solution that passes through the initial state is given by

$$\begin{aligned} x_2(t) &= X_2 - t \\ x_1(t) &= X_1 + X_2 t - \frac{1}{2}t^2 \end{aligned}$$

for $t \in [0, t_s]$. The parameter t can be removed to obtain the following equation

$$x_2^2(t) + 2x_1(t) = X_2^2 + 2X_1 \tag{42}$$

The trajectories in equations (41) and (42) intersect at the switching point $(x_1(t_s), x_2(t_s))$. These equations are plotted in figure 11. This poin tis

$$\begin{aligned} x_1(t_s) &= \frac{X_2^2 + 2X_1}{4} \\ x_2(t_s) &= -\sqrt{X_1 + \frac{1}{2}X_2^2} \end{aligned}$$

and the switching time t_s satisfies

$$t_s = X_2 + \sqrt{X_1 + \frac{1}{2}X_2^2} \tag{43}$$

and the final time is

$$t_f = t_s - x_2(t_s) = X_2 + 2\sqrt{X_1 + \frac{1}{2}X_2^2} \tag{44}$$

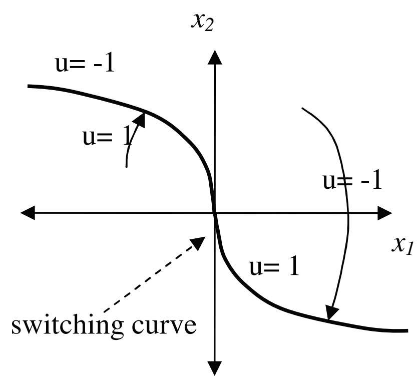

Figure 11: Positioning Problem - minimum time

The solution outlined above is the desired one provided the square root is real and t_s is positive. These conditions are satisfied if

$$\begin{aligned} X_1 > -\tfrac{1}{2}X_2^2 \text{ when } X_2 > 0 \\ X_1 > \tfrac{1}{2}X_2^2 \text{ when } X_2 \leq 0 \end{aligned}$$

The above conditions define a switching curve $x_2(x_1) = -\text{sgn}(x_1)\sqrt{2|x_1|}$. A shown in figure 11, this means that the initial states must lie to the right of the switching curve if we are to approach the home state with $u = 1$. If the initial state is on the curve, then there is no switching (i.e. $t_s = 0$). If the initial state is on the left hand side of the switching curve, then we approach the home state with $u = -1$.

In the above discussion, we didn't actually determine the co-states. But they can be determined as follows. Since the switch occurs when p_2 changes sign and when $t = t_s$, it follows that

$$B - At_s = 0$$

The value of the Hamiltonian H at $t = t_f$ must be zero so that

$$H(t_f) = 1 + B - At_f = 0$$

which implies that $B = At_s = At_f - 1$ and so $A = \frac{1}{t_f - t_s}$ and $B = \frac{t_s}{t_f - t_s}$. We can tehn use these to compute p_1 and p_2.

Minimum Fuel Consumption Problem: The preceding problem determined the time-optimal solution for the positioning problem. We now consider the problem of finding the control that minimizes fuel consumption subject to the constraint that the resulting trajectory satisfies the initial and terminal constraints. The cost functional for the fuel-optimal positioning problem is

$$J[u] = \int_0^{t_f} (k + |u|)dt \tag{45}$$

where k is a positive constant. We therefore have two components for the cost, one depending on the time and the other depending on the control effort (fuel). For simplicity we shall assume an initial state at rest at a distance X from the origin, so that $x_1(0) = X$ while $x_2(0) = x_1(t_f) = x_2(t_f) = 0$. The state equations are again those given in equation (39) and the control has unit bounded amplitude, $|u| \leq 1$.

The Hamiltonian of the fuel-optimal problem is

$$H(x, u, p) = -k - |u| + p_1 x_2 + p_2 u$$

where we've once again fixed the abnormal multiplier $p_0 = -1$. The co-state equations are unchanged from before in equation (40), so we again know that there exist constants, A and B such that

$$p_1^*(t) = A, \qquad p_2^*(t) = B - At \tag{46}$$

The novel feature of this problem appears when we try to determine the maximum of H. Applying the maximum principle, we see that

$$\begin{aligned} u^*(t) &= \arg \max_{u \in [-1,1]} H(x^*, u, p^*) \\ &= \arg \max_{u \in [-1,1]} (-k - |u| + p_1^*(t) x_2^*(t) + p_2^*(t) u) \\ &= \arg \max_{u \in [-1,1]} (-|u| + p_2^*(t) u) \end{aligned}$$

We now consider various cases depending on the actual sign of $u(t)$. Let $Q(x^*, u, p^*) = -|u| + p_2^* u$. If $u \geq 0$, then $Q = -u + p_2^* u$. This is maximized by $u = 0$ if $-1 + p_2^* < 0$ and $u = 1$ if $-1 + p_2^* > 0$. If $u \leq 0$, then $Q = u + p_2 u$, which is maximized by $u = 0$ if $1 + p_2 > 0$ and $u = -1$ if $1 + p_2 < 0$. Combining these cases, we see that

$$u^*(t) = \begin{cases} 1 & \text{if } p_2^*(t) > 1 \\ 0 & \text{if } -1 < p_2^*(t) < 1 \\ -1 & \text{if } p_2^*(t) < -1 \end{cases} \tag{47}$$

Previously, when the cost functional only depended on the terminal time, the control switched between the extreme values of -1 and $+1$. Now there are three possible settings of the control, the extreme values and the control switched off. With a cost function depending on the magnitude of u, there is a clear advantage to setting $u = 0$, if possible. But it should also be clear that at least some control must be used in order to satisfy the terminal constraints. As shown above the control is determined by the co-state $p_2^*(t)$. From the co-state equation we know $p_2^*(t) = B - At$, a linear function of t, which implies that the settings of the control must be $\{-1, 0+1\}$ or $\{+1, 0, -1\}$. Only part of either of these sequences may be needed in a particular case. It may be possible to reach the target without switching the control at all. It may be possible to switch from 0 to $+1$ are -1. But what can be stated definitely is that it is never possible to switch directly from -1 to $+1$ or back again in the optimal solution. Note that the determination of the control in equation (47) leaves its value arbitrary when $p_2^* = \pm 1$. In general this will only occur at discrete values of t. An exception would occur if p_2 is a constant. In this case Pontryagin's principle is of little help in determining the optimal control.

With the initial and final conditions, it is clear we cannot start with $u(0) = 0$, for we would then remain in the initial state. For a similar reason, we cannot end with $u = 0$. Also, an initial positive control would move the state away from the target and so we are led to conclude that the control must pass through the complete sequence $\{-1, 0, +1\}$. Let t_{s1} and t_{s2} denote the two switching times in $[0, t_f]$. The initial state trajectory is

$$
\begin{aligned}
x_1(t) &= X - \frac{1}{2}t^2 \\
x_2(t) &= -t
\end{aligned}
$$

for $t \in [0, t_{s1}]$ and the final segment of the state trajectory is

$$
\begin{aligned}
x_1(t) &= \frac{1}{2}(t_f - t)^2 \\
x_2(t) &= t - t_f
\end{aligned}
$$

for $t \in [t_{s_2}, t_f]$. These are equations for curves in the (x_1, x_2)-plane parameterized by t. We can remove the parameter to obtain the following equations for the state trajectories under the controls $u = \pm 1$,

$$
\begin{aligned}
x_1 + \frac{1}{2}x_2^2 &= X, \qquad \text{when } u = -1 \\
x_1 - \frac{1}{2}x_2^2 &= 0, \qquad \text{when } u = +1
\end{aligned}
$$

When $u = 0$, then we know that the resulting state trajectory is a straight line whose parameterized form is

$$
\begin{aligned}
x_1(t) &= b + at \\
x_2(t) &= a
\end{aligned}
$$

for $t \in [t_{s1}, t_{s2}]$. The associated trajectory curve is $x_2 = a$ (when $u = 0$). These three trajectory curves are plotted below in figure 12.

Suppose the switching values are t_{s1} and t_{s2} with $0 < t_{s1} < t_{s_2} < t_f$. Equating values of x_1 and x_2 at these switching times, to ensure that the state varies continuously yields the equations,

$$
\begin{aligned}
-t_{s1} &= a \\
t_{s2} - t_f &= a \\
X - \frac{1}{2}t_{s1}^2 &= b + at_{s1} \\
\frac{1}{2}(t_f - t_{s2})^2 &= b + at_{s2}
\end{aligned}
$$

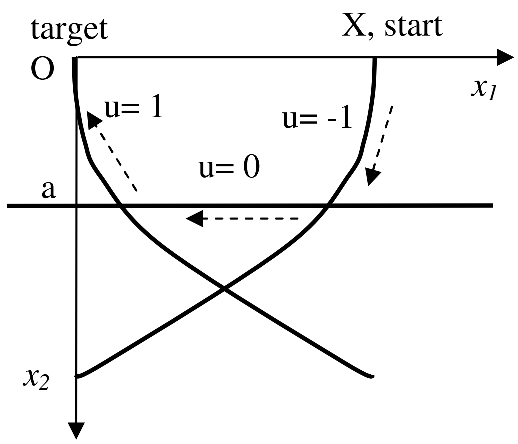

Figure 12: State trajectory curves of fuel-optimal positioning problem

from which we deduce that

$$t_{s1} + t_{s2} = t_f \tag{48}$$
$$t_{s1} t_{s2} = X \tag{49}$$

These equations determine the switching times provided we know the terminal time t_f.

To find t_f, we must consider the co-state variables. We know from equation 47 that the switchings occur when p_2 passes through the values -1 and $+1$. Therefore with $p_2 = B - At$, we see that

$$\begin{aligned} B - At_{s1} &= -1 \\ B - At_{s2} &= +1 \end{aligned}$$

so that

$$A = -\frac{2}{t_{s2} - t_{s1}} \tag{50}$$
$$B = -\frac{t_{s2} + t_{s1}}{t_{s2} - t_{s1}} = -\frac{t_f}{t_{s2} - t_{s1}} \tag{51}$$

These results determine the costates, but we still haven't determine t_f. The one result we haven't used is the fact that the Hamiltonian is constant along an optimal trajectory and in fact for problems in which the terminal time is free, this constant is zero. We can therefore use the relation

$$\begin{aligned} 0 &= H(x^*(t_f), u^*(t_f), p^*(t_f)) \\ &= -k - |u^*(t_f)| + p_1^*(t_f) x_2^*(t_f) + p_2^*(t_f) u^*(t_f) \\ &= -k - 1 + B - At_f \end{aligned}$$

to solve for t_f and see that

$$t_f = \frac{B - 1 - k}{A} \tag{52}$$

Combining equation (52) with our other equations for A (eq. (50)), B (eq. (51)), and the switching times (equations (48-49)), we find that

$$\begin{aligned}
t_{s2} - t_{s1} &= \frac{t_f}{k+1} \\
t_{s2} &= \frac{t_f(k+2)}{2(k+1)} \\
t_{s1} &= \frac{t_f k}{2(k+1)} \\
X &= \frac{k(k+2)}{4(k+1)^2} t_f^2
\end{aligned}$$

The optimal cost is given by

$$\begin{aligned}
J^* &= kt_f + t_{s1} + (t_f - t_{s2}) \\
&= t_f(k+1) - \frac{t_f}{k+1} \\
&= 2\sqrt{k(k+2)X}
\end{aligned}$$

If k is large, so the cost depends mainly on the time taken to reach the target and to a minor extent on the fuel consumed, the itnerval $t_{s2} - t_{s1}$ during which the engine is idle will be veryshort and we approach the time-otpimal solution we obtained earlier. If k is small, then the saving of fuel is the dominant consideration and the flat trajectory lies just underneath the x_1 axis. In the limit as k tends to zero, the cost can be made arbitrarily small, but the time to reach the target tends to infinity. If we attempted to solve the problem with $k = 0$, we would have found that the switching time $t_{s1} = 0$ so we would have to begin with the engines switched off. Since this implies we would remain at the initial state forever, there is no optimal solution, although we could make the cost arbitrarily small.

Minimum Time Harmonic Oscillator Problem: We now turn to the optimal control of a harmonic oscillator. As in the positioning problem, this leads to a bang-bang optimal control law. This problem often exhibits many more switches than the positioning problem. In the harmonic oscillator, the system oscillates about an equilibrium point. We represent the displacement of the object from its equilibrium by x_1 and its velocity by x_2. The state equations, in their simplest form, can be written as

$$\begin{aligned}
\dot{x}_1 &= x_2 \\
\dot{x}_2 &= -x_1 + u
\end{aligned}$$

where the control u is constrained to lie in the interval $[-1, +1]$. The initial state has the arbitrary value (X_1, X_2) and the objective is to reach the target state, $(0, 0)$, in the shortest possible time. So this is a problem with fixed terminal point and floating terminal time. The cost functional we seek to minimize is

$$J(u, t_f) = \int_0^{t_f} dt$$

so the Hamiltonian for this problem is

$$H(x, u, p) = -1 + p_1 x_2 + p_2(u - x_1)$$

We invoke the maximum principle to determine necessary conditions on the optimal control. In particular, the optimal control should satisfy,

$$\begin{aligned}
u^*(t) &= \arg \max_{u \in [-1,1]} H(x^*, u, p^*) \\
&= \arg \max_{u \in [-1,1]} (-1 + p_1^*(t) x_2^*(t) + p_2^*(t)(u - x_1^*(t))) \\
&= \arg \max_{u \in [-1,1]} p_2^*(t) u
\end{aligned}$$

If $p_2^*(t) < 0$, then this is maximized by taking $u = -1$ and if $p_2^*(t) > 0$, then the maximizing control is $u = 1$. We may therefore express the optimal control, u^* as

$$u^*(t) = \begin{cases} 1 & \text{when } p_2^*(t) > 0 \\ -1 & \text{when } p_2^*(t) < 0 \end{cases} \tag{53}$$

From the costate equation, we know that

$$\begin{bmatrix} \dot{p}_1 \\ \dot{p}_2 \end{bmatrix} = -\frac{\partial H}{\partial x} = \begin{bmatrix} +p_2 \\ -p_1 \end{bmatrix}$$

Since we are primarily concerned with p_2, we can combine these equations to obtain a second order differential equation

$$\ddot{p}_2 = -p_2$$

which has the general solution

$$p_2^*(t) = \beta \cos(t) - \alpha \sin(t) \tag{54}$$

where α and β are constants of integration. From equation 53 we see that

$$u^*(t) = -\operatorname{sgn}\left(\beta \cos(t) - \alpha \sin(t)\right)$$

As before, we see that the optimal control is a bang-bang type control law in which we switch back and forth between $+1$ and -1.

In comparing this problem's costate to the positioning problem's costate, we see an obvious difference. Whereas the positioning problem's $p_2(t)$ had at most one switching instant, the costate for the harmonic problem has multiple switching instants. In order to solve our problem, we need to determine how many switching instants are required. Although we don't yet know α and β, we do know that the switching instants must occur at intervals of π. The first switch can occur at any time not greater than π after the first application of a control. The second switch, if needed, cannot occur before a time π has elapsed since the first switch, and it must occur then. Similarly, if any further time π has elapsed and the target has not yet been reached, a third switch must occur and so on.

From the state equation, it follows that

$$(x_1 - u)\dot{x}_1 + x_2 \dot{x}_2 = 0$$

so that when u is held constant, the trajectory is part of the curve

$$(x_1 - u)^2 + x_2^2 = c^2$$

for some constant c. Thus the optimal trajectory is made up of arcs of circles, with center at $(1, 0)$ when $u = 1$ and center $(-1, 0)$ when $u = -1$. If we write

$$\begin{aligned} x_1 - u &= c \cos\theta \\ x_2 &= c \sin\theta \end{aligned}$$

annd substitute these back into the state equations, we find that

$$\dot{\theta} = -1$$

Hence the circular paths are traced out in the clockwise direction and the time taken to move between two points on the same arc equals the angle between the two radii joining them to the center of the arc. It therefore follows that in the optimal solution,the maximum amount of any one circular arc that is used is a semicircle and when that has been used, the center of the next arc must switch to the alternate.

Instead of starting at the initial state and working forwards in time, it is more convenient to work backwards from the target and to consider all points from which the target can be reached by controls with an increasing number of switches. The origin lies on the circle with center $(1, 0)$ and radius 1, so it can be approached from negative values of x_2 using the control $u = 1$. We know that not more than a semicircle can be used without switching, so the points lying on this circle below the x_1 axis can be controlled by this choice of u. Let C_n denote the semicircle below the x_1 axis with radius 1 and center at $(n, 0)$ and let C_{-n} lie above the x_1 axis and have center at $(-n, 0)$ and radius 1. Points on C_1 can be controlled by $u = 1$ without any switch, and, by symmetry points lying on C_{-1} can be controlled by $u = -1$. The final approach to the target for the optimal trajectories from any initial states must be along one or the other of these two semicircles. These semicircles over which the final state is approached are plotted out in figure 13.

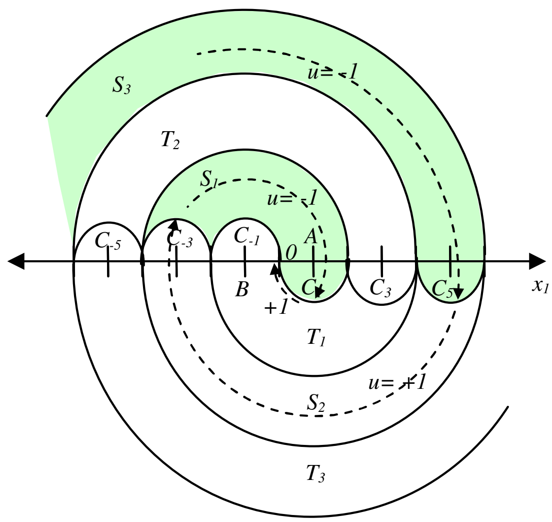

Figure 13: Optimal Trajectories for time-optimal harmonic oscillator problem

To simplify our discussion, let's focus on those paths whose last section lie on C_1. To reach C_1, we must be using the opposite control $u = -1$, so the path will be an arc with center $B = (-1, 0)$. The maximum possible arc length is a semicircle, so that the end of the arc that passes through a point P on C_1 will be the opposite end of the diameter of the circle with center B through P. These end points lie on the reflection of C_1 in B that is they will lie on C_{-3}. Hence all points in the region S_1 of figure 13 can be controlled to the origin by the control sequence $\{-1, 1\}$ with one switch. The boundaries of S_1 are C_1, C_{-1}, C_{-3}, and the semicircle with center B and radius 3. The reflection of C_{-3} about $A = (1, 0)$ is C_5 and points lying in S_2 require the control sequence $\{1, -1, 1\}$ to reach the origin optimally. The boundaries of S_2 are C_{-3}, C_t and the semicircles with center A and radii 3 and 5. Proceeding in this manner, we cover half the plane by regions of this type; for any point lying in one of these regions a sequence of controls ending with $+1$ must be used. The trajectory for a point Z in S_4 is also shown in figure 13. The other half of the plane is covered by regions T_1, T_2, etc. for which the optimal control ends in -1.

Minimum Fuel Harmonic Oscillator Problem: We now examine optimal control of the harmonic oscillator under a minimum-energy cost functional of the form

$$J[u] = \int_0^{t_f} |u| dt$$

We suppose the system is initially at rest a distance X from the equilibrium position and we assume that we again wish to get "home". The required initial and final states, therefore, are

$$\begin{aligned} x_1(0) &= X_1, \qquad & x_2(0) &= X_2 \\ x_1(t_f) &= 0, \qquad & x_2(t_f) &= 0 \end{aligned}$$

The Hamiltonian for this problem is

$$H(x, p, u) = -|u| + p_1 x_2 + p_2(-x_1 + u)$$

We apply the maximum principle to see that u^* satisfies,

$$\begin{aligned} u^*(t) &= \arg \max_{u \in [-1,1]} H(x^*, u, p^*) \\ &= \arg \max_{u \in [-1,1]} (-|u| + p_1^*(t) x_2^*(t) + p_2^*(t)(-x_1^*(t) + u)) \\ &= \arg \max_{u \in [-1,1]} (-|u| + p_2^*(t) u) \end{aligned}$$

The maximum of H is achieved if we choose

$$u^*(t) = \begin{cases} 1 & \text{if } p_2^*(t) > 1 \\ 0 & \text{if } |p_2(t)| < 1 \\ -1 & \text{if } p_2^*(t) < -1 \end{cases}$$

The costate equations are

$$\begin{bmatrix} \dot{p}_1 \\ \dot{p}_2 \end{bmatrix} = -\frac{\partial H}{\partial x} = \begin{bmatrix} -p_2 \\ p_1 \end{bmatrix}$$

and the points $(p_1(t), p_2(t))$ lie on a circle centered at $(0,0)$ with radius R. If $R \le 1$, we must have $u = 0$ for all t. It is impossible to reach the target state using this control, so this must imply that the control follows the sequence $\{1, 0, -1, 0, 1, \ldots\}$. Since the path of the costate, $p(t)$, is traced out at a unit angular rate, and since switchings occur whenever p_2 passes through the values ± 1, it follows that the controls ± 1 are applies for equal times α, say, and the zero control is applied for equal times $\pi - \alpha$, where $0 < \alpha < \pi$. The corresponding trajectories are composed of arcs of circles in the (x_1, x_2) plane; the centers of these circles are at $(1,0)$ when $u = 1$, at $(0,0)$ when $u = 0$, and at $(-1,0)$ when $u = -1$. The solution for a general initial state can be found by piecing together a succession of circular arcs.

As a specific example, consider the initial state $X_1 = 0$ and $X_2 = 2$. The control sequence $\{-1, 1\}$ steers this point to the origin via the point $(1, -1)$ in a total time π, with cost also equal to π. This is, in fact, the strategy for the time-optimal problem. The sequence $\{0, 1\}$ of controls, however, also reaches the target state, with the siwtch occuring when the state is $(2, 0)$. The cost is again equal to π and the terminal time is $\frac{3}{2}\pi$, but this cannot be optimal since the total time during which the pair of controls is used is greater than π. If we now try the sequence of controls $\{-1, 0, 1\}$ and piece together the corresponding arcs, we find that the first control switch must be made at time $2\alpha - \frac{1}{2}\pi$, the second at 2α, and the target is reached at time $2\alpha + \frac{1}{2}\pi$, where $\alpha = \tan^{-1} 2$. The optimal cost equals the total time during which a non-zero control is being used and this is equal to 2α. Note that this is less than the time-optimal cost, which was equal to π, but the terminal time is greater than π as anticipated. This optimal state trajectory is shown in figure 14.

It must be remembered that Pontryagin's principle only provides necessary conditions for the optimal solution. It may be that there are many solutions that satisfy the maximum principle. If this is the case, we must find all of the soolutions and find the one that gives minimal cost. In the present example, there

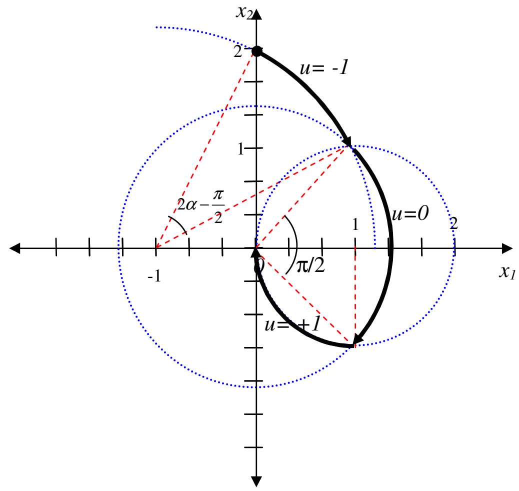

Figure 14: Minimum-fuel state trajectory for harmonic oscillator problem with two switches

are many solutions satisfying the maximum principle. Suppose the optimal trajectory from the given initial point first crosses the x_1 axis at the point $x_1 = X$. Now we apply the following cycle of controls: $u = 0$ for a time β, $u = -1$ for a time $\pi - 2\beta$,, and $u = 0$ for a further time β. On solving the state equations we find that after the total time π we reach the point $(x_1, x_2) = (2\cos\beta - X, 0)$ and the cost incurred is equal to $\pi - 2\beta$. If we repeat this cycle for a further time π, but with the sign of u reversed, we reach the point $(x_1, x_2) = (X - 4\cos\beta, 0)$ at a toal cost equal to $2(\pi - 2\beta)$. It follows thhat we can reach the target after n such cycles, provided we choose β so that $\cos\beta = X/2n$, and that the total cost along the trajectory from $(X, 0)$ to $(0, 0)$ is given by

$$J_n = n\left[\pi - 2\cos^{-1}\left(\frac{X}{2n}\right)\right] = 2n\sin^{-1}\left(\frac{X}{2n}\right)$$

Although the total time increases with n, the cost decreases as shown in the plot on the righthand side of figure 15. We conclude therefore that there is NO optimal solution, since the set of values of J_n is bounded below by X, but is never equal to X.

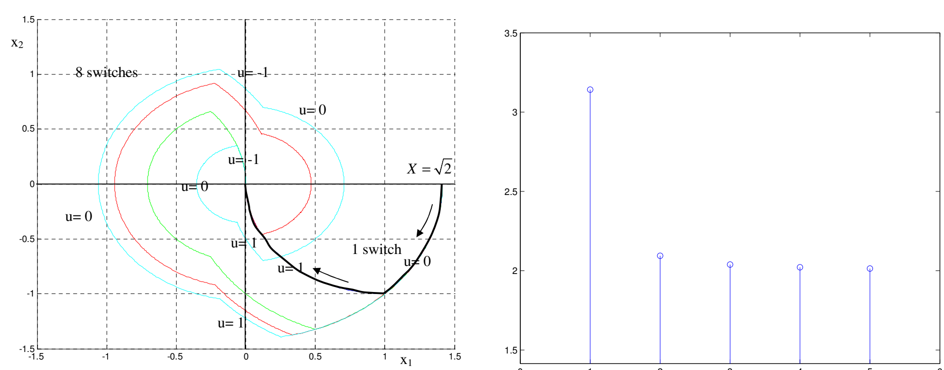

Figure 15: Minimum-fuel harmonic oscillator problem

The previous example shows that unless we are careful, there may exist no “optimal” solution, even though we can find a sequence of controls that approach this optimal solution arbitrarily closely. What is the optimal solution? The following trajectory in figure 16was computed using Matlab for a large number (25) switches. Note that essentially, what the control system is doing is remaining in the coast state for as long as possible and then turns on the control for a brief period for a small interval about the $x_1 = 0$ axis. As the number of switches increases, we expect this interval to shrink and the "control" jumps to get smaller. This control strategy takes an increasing amount of time to run, but in the limit the control effort is minimized. We can think of our control trajectories with finite number of switches as an approximation to this "limiting" control.

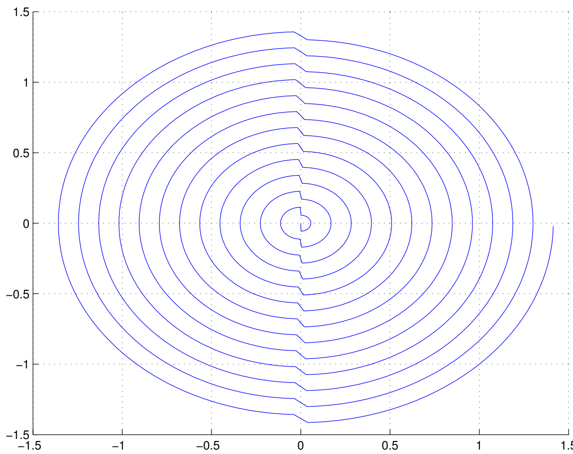

Figure 16: Minimum-fuel control with 25 switches

9.1 - Bang-Bang Principle - Preliminaries

The next couple lectures present an alternative proof the maximum principal for the linear time optimal control problem. The techniques used in this proof are much different than the extensions of the variational methods employed by Pontryagin and his colleagues. Rather than adopting a calculus of variations view, we take advantage of some powerful results in functional analysis. These results allow us to derive the maximum principle as a consequence of the sequential compactness of the unit sphere in the $\mathcal{L}_\infty$ space of essentially bounded functions. This lecture presents those results from functional analysis that will be used later.

The system under study is a linear time-invariant system whose state, $x(\cdot) : \mathbb{R}^+ \to \mathbb{R}^n$, satisfies the initial value problem,

$$\dot{x}(t) \quad = \quad Ax(t) + Bu(t), \quad x(0) = x_0$$

where $u(\cdot) : [0, T] \to [-1, 1]^m$ is a piecewise continuous function in $\hat{\mathcal{C}}$. Note that the range of this function is an m-dimensional hypercube $[-1, 1]^m$. We'll denote the set of such functions as $\mathcal{U}$.

The question examined in this section is whether there exists a piecewise continuous control $u(\cdot) : [0, T] \to [-1, 1]^m$ that takes the system from its initial state x_0 to the origin. The main finding will be that the space of such controls contains a *bang-bang* control; namely a control whose components are either $+1$ or -1 for all time.

As indicated above, our issue concerns the controllability/reachability of the origin from an initial state x_0. While this is studied by many first year graduate students in control; there is little harm in repeating the definitions to make this discussion self-contained. The *reachable set from time t* is denoted as $\mathcal{C}(t)$. It consists of all initial states x_0 for which there exists a control in $\mathcal{U}$ such that $x(t) = 0$. The *reachable set* $\mathcal{C}$ of the system is simply the union of all reachable sets at t; i.e., $\mathcal{C} = \bigcup_{t \geq 0} \mathcal{C}(t)$.

Since the system is linear and time-invariant, we know that

$$\begin{aligned} x(t) &= \Phi(t)x_0 + \int_0^t \Phi(t-s)Bu(s)ds \\ &= \Phi(t)x_0 + \Phi(t)\int_0^t \Phi^{-1}(s)Bu(s)ds \end{aligned}$$

where $\Phi(t) = e^{At}$ is the matrix exponential function of matrix A. Since this is a continuous-time system, we know Φ is invertible. Furthermore we know that if $x_0 \in \mathcal{C}(t)$ (reachable at t set), then $x(t) = 0$. So the above equation can be rewritten as

$$x_0 = -\int_0^t \Phi^{-1}(s)Bu(s)ds \tag{55}$$

Equation (55) can be seen as a necessary and sufficient characterizing of the control u that moves the system state from x_0 to the origin. This equation, therefore, also becomes a necessary and sufficient characterization of the set $\mathcal{C}(t)$.

We say the control $u \in \mathcal{U}$ is *bang-bang* if for all $t \geq 0$ and each index $i = 1, 2, \ldots, m$, we know $|u_i(t)| = 1$. Our main results today concerns the existence of bang-bang controls. To prove this theorem, however, we need to first introduce some useful tools from functional analysis.

We're confining our attention to the complete linear space of essentially bounded functions. In particular let $x(\cdot) : \mathbb{R}^+ \to \mathbb{R}^n$ be a measurable function. The $\mathcal{L}_\infty$ norm of x is defined as

$$\|x\|_{\mathcal{L}_\infty} = \text{ess} \sup_{0 \leq s \leq t} |x(s)|$$

where $|x|$ is the Euclidean norm of a vector $x \in \mathbb{R}^n$ and $\text{ess}\sup\Omega$ is the supremum of all points in Ω excluding a set of measure zero (the essential supremum of the set Ω). The linear space of all essentially bounded functions (i.e. functions with bounded $\mathcal{L}_\infty$ norm) is denoted as $\mathcal{L}_\infty$. This is a completed normed linear space (i.e. Banach space) of functions.

In the following we'll need the following notion of *weak-** convergence. Let $u_n \in \mathcal{L}_\infty$ be an infinite sequence of functions ($n = 1, 2, \ldots, \infty$). Let u be some other function in $\mathcal{L}_\infty$. We say u_n converges to u in the weak-* sense (denoted as $u_n \overset{*}{\to} u$) if and only if

$$\int_0^t u_n(s)v(s)ds \to \int_0^t u(s)v(s)s \tag{56}$$

as $n \to \infty$ for all $v(\cdot) : [0, t] \to \mathbb{R}^n$ such that $\int_0^t |v(s)|ds < \infty$ (absolutely integrable over $[0, t]$). The following theorem will be useful to us.

Alaoglu's Theorem: The unit sphere in $\mathcal{L}_\infty$ is compact in the weak-* topology.

Recall that there are at least two ways of defining when a set (space) is *compact*. A countable notion of compactness is defined as the existence of a finite subcover. This is equivalent to what is often referred to as *sequential compactness*. A set is sequentially compact if any sequence in the set has a convergent subsequence. Our statement of Alaoglu's theorem makes use of the definition of sequential compactness and the convergence of the sequences is defined with respect to the weak-* topology shown in equation (56).

The other main result we need is the Krein-Milman Theorem. To properly state this theorem, however, we first need to define the notion of an *extreme point*. Given a convex set K, a point $z \in K$ is said to be an *extreme point* if there do not exist distinct $x, y \in K$ such that

$$z = \lambda x + (1 - \lambda)y$$

with $\lambda \in (0, 1)$. Essentially an extreme point is a "corner" of the convex set K.

Krein-Milman Theorem: Let K be a convex non-empty subset of $\mathcal{L}_\infty$ which is weak-* compact. Then K has at least one extreme point.

With these functional analysis preliminaries out of the way, we can now prove the main theorem of this section. This theorem proves that the set of reachable controls in $\mathcal{U}$ contain bang-bang controls.

Theorem: Let $t > 0$ and suppose $x_0 \in \mathcal{C}(t)$ for the LTI system

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0$$

where $u \in \mathcal{U}$. Then there exists a bang-bang control $u^* \in \mathcal{U}$ which steers x_0 to 0 at time t.

Proof: Let $x_0 \in \mathcal{C}(t)$ and consider the set

$$K = \{u \in \mathcal{U} \,:\, u \text{ steers } x_0 \text{ to } 0 \text{ at time } t\}$$

K is, therefore, the set of all admissible controls that can reach the origin from x_0 in time t. We make the following assertions about K. First that K is *convex* and second that it is *compact* with respect to the weak-* topology. These two assertions allow us to use the Krein-Milman theorem to establish the existence of a control in K that is an extreme point. We then prove that this extreme point is a bang-bang control.

First let's establish that K is convex. Recall that if $u_1 \in K$ and $u_2 \in K$, then

$$\begin{aligned} x_0 &= -\int_0^t \Phi^{-1}(s)Bu_1(s)ds \\ x_0 &= -\int_0^t \Phi^{-1}(s)Bu_2(s)ds \end{aligned}$$

Note that for any $\lambda \in [0,1]$ that

$$\begin{aligned} x_0 &= \lambda x_0 + (1-\lambda)x_0 \\ &= -\lambda \int_0^t \Phi^{-1}(s)Bu_1(s)ds - (1-\lambda)\int_0^t \Phi^{-1}(s)Bu_2(s)ds \\ &= -\int_0^t \Phi^{-1}(s)B\left[\lambda u_1(s) + (1-\lambda)u_2(s)\right]ds \end{aligned}$$

This is sufficient to emply that $\lambda u_1 + (1-\lambda)u_2 \in K$ and so K is convex.

We now prove K is compact (weak-*). Consider an infinite sequence of functions $u_n \in K$ for $n = 1, 2, \ldots, \infty$. Since this sequence is also in the unit sphere (with respect to the $\mathcal{L}_\infty$-norm) we know by Alaoglu's theorem that there exists a convergent subsequence $\{u_{n_k}\}_{k=1}^\infty$ that converges to a function $u \in K$ in the weak-* sense. In other words, $u_{n_k} \xrightarrow{*} u$. However, since $u_{n_k} \in K$, we know it must satisfy

$$x_0 = -\int_0^t \Phi^{-1}(s)Bu_{n_k}(s)ds$$

In the limit as $k \to \infty$ we know that

$$-\int_0^T \Phi^{-1}(s)Bu_{n_k}(s)ds \to x_0 = -\int_0^t \Phi^{-1}(s)Bu(s)ds$$

Clearly $v(s) = \Phi^{-1}(s)B$ is absolutely integrable over the finite interval $[0,t]$. So we can conclude that this convergence of u_{n_k} to u is in the weak sense and hence the limit point $u \in K$. Since the limit is in K we know that K is sequentially compact.

Knowing that K is convex and compact means that the hypotheses of the Krein-Milman theorem are satisfied. We know, therefore, that there exists an extreme point $u^* \in K$. We now prove that this extreme point is a bang-bang control. This will be done through contradiction.

In particular, suppose that u^* is an extreme point in K but is not bang-bang. Then there exists an index $i \in \{1, 2, \ldots, m\}$ and a time interval $I \subset [0, t]$ such that $|u_i^*(s)| < 1$ for all $s \in E$. For any $\epsilon > 0$, we can find an interval $F \subset E$ such that $|u_i^*(s)| \leq 1 - \epsilon$ for all $s \in F$. We may choose any real-valued continuous function where $v(s) = 0$ for $s \notin F$ and $v(s) \neq 0$ for all $s \in F$ such that

$$\int_F \Phi^{-1}(s) B v(s) ds = 0$$

We then define two new functions

$$\begin{aligned} u_1 &= u^* + \epsilon v \\ u_2 &= u^* - \epsilon v \end{aligned}$$

We then note that for u_1,

$$\begin{aligned} -\int_0^t \Phi^{-1}(s) B u_1(s) ds &= -\int_0^t \Phi^{-1}(s) B u^*(s) - \epsilon \int_0^t \Phi^{-1}(s) B v(s) ds \\ &= x_0 - \epsilon \int_0^t \Phi^{-1}(s) B v(s) ds = x_0 \end{aligned}$$

which means that $u_1 \in K$.

We can prove in a similar way that $u_2 \in K$ also. So we know know u_1, u_2, and u^* are in K and that u^* is an extreme point. But clearly from the definition of u_1 and u_2 we see that

$$u^* = \frac{1}{2} u_1 + \frac{1}{2} u_2$$

which means that u^* cannot be an extreme point of K. We have therefore obtained a contradiction that was generated by us assuming u^* was an extreme point but not necessarily bang-bang. Therefore if u^* is extreme it must also be a bang-bang control. □

9.2 - Maximum Principle (revisited)

We use the results from last lecture to establish that the linear time optimal (LTO) problem is solved by a bang-bang control. We then use this result to reprove Pontryagin's Maximum Principle for the LTO problem. This derivation provides an alternative view of the maximum principle that is somewhat more streamlined than the original proof developed by Pontraygin and his colleagues. Part of the reason for the simplicity lies in our restriction to the linear time-optimal problem.

Consider the linear dynamical system

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0$$

where $u \in \mathcal{U}$ is a piecewise continuous function takes values in the supremum norm's unit sphere, $[-1, 1]^m$. We want to minimize

$$J(u) = -\int_0^T ds = -T$$

where T is the first time under control u when $x(s) = 0$. The LTO problem is to find the $u^* \in \mathcal{U}$ such that

$$J(u^*) \geq J(u), \quad \forall u \in \mathcal{U}$$

This maximum is $T^* = -J(u^*)$ is the minimum needed to reach the origin. Note that this problem is a fixed endpoint free-time problem (identical to what we considered earlier except our restriction to linear systems)

The first theorem claims that the solution to the LTO problem is a bang-bang control.

> **Theorem:** If the system is controllable (i.e., $\mathcal{C} = \mathbb{R}^n$) then there exists a bang-bang control $u^* \in \mathcal{U}$ solving the LTO problem.

Proof: Let

$$T^* = \inf \{\tau \,:\, x_0 \in \mathcal{C}(\tau)\}$$

We want to show that $x_0 \in \mathcal{C}(T^*)$, since this would mean there exists a control in $\mathcal{U}$ that derives the initial state to the origin.

Choose an infinite sequence of times

$$t_1 \geq t_2 \geq t_3 \geq \cdots$$

so that $t_n \to T^*$ and

$$x_0 \in \mathcal{C}(t_n)$$

for each t_n in the sequence.

Since $x_0 \in \mathcal{C}(t_n)$, there exists a control $u_n \in \mathcal{U}$ such that

$$x_0 = -\int_0^{t_n} \Phi^{-1}(s) B u_n(s) ds$$

for all $n = 1, 2, \ldots, \infty$. By Alaoglu's theorem, there exists a subsequence u_{n_k} and a control $u^* \in \mathcal{U}$ such that $u_{n_k} \xrightarrow{*} u^*$ as $n_k \to \infty$. This implies that for each n_k

$$x_0 = -\int_0^{t_{n_k}} \Phi^{-1}(s) B u_{n_k}(s) ds$$

If we take the limit of the u_{n_k} we see that

$$-\int_0^{t_{n_k}} \Phi^{-1}(s) B u_{n_k}(s) ds \to -\int_0^{t_{n_k}} \Phi^{-1}(s) B u^*(s) ds = x_0$$

And if we then take the limit of the limits of integration as $t_{n_k} \to T^*$, we obtain

$$-\int_0^{t_{n_k}} \Phi^{-1}(s) B u_{n_k}(s) ds \to -\int_0^{T^*} \Phi^{-1}(s) B u^*(s) ds = x_0$$

which implies that $x_0 \in \mathcal{C}(T^*)$ and so u^* is optimal since T^* is the minimum time. Moreover, our earlier theorem (last lecture) showed that this control must be bang bang. $\square$

Pontryagin's maximum Principle (PMP) is a necessary characterization of this bang-bang control for the LTO problem. We now use our new tools to prove the maximum principle for this special problem. You will notice various similarities to what we did earlier. Before marching ahead with the proof, let's define the set $K(t, x_0)$

$$\begin{aligned} K(t, x_0) &= \text{reachable set from } x_0 \text{ by time } t \\ &= \{x_1 \in \mathbb{R}^n \,:\, \exists u \in \mathcal{U} \text{ such that } x(t) = x_1 \text{ and } x(0) = x_0\} \end{aligned}$$

From our usual results in linear systems theory we know that $x_1 \in K(t, x_0)$ if and only if

$$x_1 = \Phi(t)x_0 + \Phi(t)\int_0^t \Phi^{-1}(s)Bu(s)ds$$

for some $u \in \mathcal{U}$. We now prove the following result about K.

Theorem: $K(t, x_0)$ is convex and closed.

Proof: The proof is very similar to what we did in the last lecture. Let $x_1, x_2 \in K(t, x_0)$, then there exists $u_1, u_2 \in \mathcal{U}$ such that

$$\begin{aligned} x_1 &= \Phi(t)x_0 + \Phi(t)\int_0^t \Phi^{-1}(s)Bu_1(s)ds \\ x_2 &= \Phi(t)x_0 + \Phi(t)\int_0^t \Phi^{-1}(s)Bu_2(s)ds \end{aligned}$$

Let $0 \leq \lambda \leq 1$ and form the convex combination of x_1 and x_2 to obtain

$$\lambda x_1 + (1-\lambda)x_2 = \Phi(t)x_0 + \Phi(t)\int_0^t \Phi^{-1}(s)B\left(\lambda u_1(s) + (1-\lambda)u_2(s)\right)ds$$

Note that because $\lambda u_1(s) + (1-\lambda)u_2(s) \in [-1, 1]^m$, we can conclude that $\lambda x_1 + (1-\lambda)x_2 \in K(t, x_0)$ and so $K(t, x_0)$ is convex.

To show that this set is closed (i.e. it contains its limit points), let's first consider an infinite sequence $x_k \in K(t, x_0)$ for $k = 1, 2, \ldots, \infty$ where $x_k \to y$ as $k \to \infty$. We must show that $y \in K(t, x_0)$.

This is done using Alaoglu's theorem. Again since $x_k \in K(t, x_0)$ there exist $u_k \in \mathcal{U}$ such that

$$x_k = \Phi(t)x_0 + \Phi(t)\int_0^t \Phi^{-1}(s)Bu_k(s)ds \tag{57}$$

for each k. By Alaoglu's theorem, there exists a convergent subsequence $u_{k_j} \xrightarrow{*} u$ where $u \in \mathcal{U}$. So letting $k = k_j$ in equation (57)and taking the limit as $k_j \to \infty$ we see that

$$y = \Phi(t)x_0 + \Phi(t)\int_0^t \Phi^{-1}(s)Bu(s)ds$$

which shows $y \in K(t, x_0)$ and so $K(t, x_0)$ is also a closed set. $\square$

We are now in a position to prove PMP for LTO.

PMP-LTO: If u^* is an optimal control for the LTO problem, then there exists a nonzero vector $h \in \mathbb{R}^n$ such that

$$h^T\Phi^{-1}(t)Bu^*(t) = \max_{u \in [-1,1]^m}\left\{h^T\Phi^{-1}(t)Bu\right\}$$

for all $0 \leq t \leq T^* = \inf\{t : x_0 \in \mathcal{C}(t)\}$.

Remark: This doesn't like our earlier statement of the PMP theorem. It is an intermediate form that depends heavily on the fact that our system is linear and time-invariant. We will see later that this is equivalent to our earlier maximum principle.

Proof: Let $\partial K(T^*, x_0)$ be the boundary of $K(T^*, x_0)$. Since T^* is the minimum time it takes to reach the origin, we can clearly see $0 \in \partial K(T^*, x_0)$. Since $K(T^*, x_0)$ is convex, there exists a supporting hyperplane to $K(T^*, x_0)$ at 0. In other words, we can find a vector $g \in \mathbb{R}^n$ such that $g^T x_1 \leq 0$ for all $x_1 \in K(T^*, x_0)$.

Now if $x_1 \in K(T^*, x_0)$, then there exists $u \in \mathcal{U}$ such that

$$x_1 = \Phi(T^*)x_0 + \Phi(T^*) \int_0^{T^*} \Phi^{-1}(s)Bu(s)ds$$

Also since $0 \in K(T^*, x_0)$ we see that

$$0 = \Phi(T^*)x_0 + \Phi(T^*) \int_0^{T^*} \Phi^{-1}(s)Bu^*(s)es$$

Since $g^T x \leq 0$ this implis

$$g^T \left(\Phi(T^*)x_0 + \Phi(T^*) \int_0^T \Phi^{-1}(s)Bu(s)ds \right) \leq 0 = g^T \left(\Phi(T^*)x_0 + \Phi(T^*) \int_0^T \Phi^{-1}(s)Bu^*(s)ds \right)$$

Let's define

$$h^T = g^T \Phi(T^*) \in \mathbb{R}^n$$

so that the prior integral expression simplifies to

$$\begin{aligned} \int_0^{T^*} h^T \Phi^{-1}(s)Bu(s)ds &\leq \int_0^{T^*} h^T \Phi^{-1}(s)Bu^*(s)ds \\ \Rightarrow \int_0^{T^*} h^T \Phi^{-1}(s)B\left(u^*(s) - u(s)\right)ds &\geq 0 \end{aligned} \tag{58}$$

for all $u \in \mathcal{U}$.

Assume the PMP does not hold. Then there exists a subset $E \subset [0, T^*]$ such that

$$h^T \Phi^{-1}(s)Bu^*(s) < \max_{u \in [-1,1]^m} \left\{ h^T \Phi^{-1}(s)Bu \right\}$$

for all $s \in E$. So let's design a new control (reminiscent of our earlier needle perturbation)

$$\hat{u}(s) = \begin{cases} u^(s) & s \notin E \\ \underline{u}(s) & s \in E \end{cases}$$

where $\underline{u}(s)$ is selected so that

$$h^T \Phi^{-1} B\underline{u}(s) = \max_{u \in [-1,1]^m} \left\{ h^T \Phi^{-1}(s)Bu \right\}$$

With this choice we see that

$$\int_E h^T \Phi^{-1}(s)B(u^*(s) - \hat{u}(s))ds < 0 \Rightarrow \int_0^{T^*} \Phi^{-1}(s)B(u^*(s) - \hat{u}(s))ds < 0$$

This clearly contradicts what we established in equation (58) and so the PMP must hold. □

Note that the above theorem does not have the form we obtained earlier for the maximum principle since the maximum principle was not stated in terms of the Hamiltonian,

$$H(x, u, p) = (Ax + Bu)^T p$$

We now show how to recover our original statement of the maximum principle (with the exception of the transversality conditions)

Theorem: Let u^* be the optimal control for the LTO problem with associated state trajectory x^*. Then there exists a function $p^* : [0, T] \to \mathbb{R}^n$ such that

$$\begin{aligned} \dot{x}^*(t) &= H_p(x^*, u^*, p^*) \\ \dot{p}^*(t) &= -H_x(x^*, u^*, p^*) \end{aligned}$$

with

$$H(x^*, u^*, p^*) = \max_{u \in [-1,1]^m} H(x^*, u, p^*)$$

Proof: Select the vector h from the proof in the previous theorem and consider the adjoint system

$$\dot{p}^*(t) = -A^T p^*(t), \quad p^*(0) = h$$

Then we see that

$$p^*(t) = \left(e^{-At}\right)^T h$$

which implies that

$$(p^*(t))^T = h^T \Phi^{-1}(t)$$

Inserting this into the equations for our previous maximum principle recover the maximization of the Hamiltonian and a quick check of the dynamical systems confirm that x^* and p^* satisfy Hamitlon's canonical equations. $\square$

This lecture has provided an alternate proof for the maximum principle for the LTO problem. We say that for the LTO case, the optimal solution is always a bang-bang control. We saw that the proof could be greatly simplified, particularly, with regard to establishing the existing of the optimal control. This was accomplished using some tools from functional analysis such as Alaoglu's theorem and the Krein-Milman theorem.

Bang-bang controls are interesting because they reduce the number of control decisions to a finite set. This can greatly simplify the search for optimal controls. While we saw that such bang-bang controls always appear in LTO problems. One might wonder under what conditions a similar set of bang-bang controls might appear in more general nonlinear systems. The next lecture examines the existence of bang-bang controls for a special class of nonlinear controls that are affine in the controls. We also discuss singular control problems in the next lecture.

9.3 - Bang-Bang Nonlinear Controls and Singular Optimal Control Problems

The previous lectures showed that the strong extremals of the linear time-optimal problem where $u(t) \in [-1, 1]^m$ is a bang-bang control. An important part of this proof rested on the linear nature of the system. We now try to extend the Bang-Bang Property to nonlinear systems.

We start with an example. Consider the planar system,

$$\begin{aligned} \dot{x}_1 &= x_2^2 - 1 \\ \dot{x}_2 &= u \end{aligned}$$

with the constraint that $u(t) \in [-1, 1]$. We want to find a control that transfers the system state from $x_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ to the origin in *minimum time*. The unique optimal control for this problem is $u \equiv 0$. This is not a bang-bang control. So we already know that, in general, bang-bang controls may not exist for the time optimal control of nonlinear systems. The issue is whether or not we can identify a set of nonlinear systems for which bang-bang controls do exist.

Let's focus on a system that is affine in the controls,

$$\dot{x} = f(x) + g(x)u$$

The Hamiltonian for this system's time-optimal problem is

$$H(x, u, p, p_0) = \langle p, f(x) + g(x)u \rangle + p_0$$

Applying the maximum principle we see that

$$u^*(t) = \arg \max_{u \in [-1,1]} \langle p^*(t), g(x^*(t)) \rangle u$$

We define a *switching function*

$$\phi(t) = \langle p^*(t), g(x^*(t)) \rangle$$

and quickly deduce that the control maximizing the Hamiltonian is

$$u^*(t) = \begin{cases} 1 & \phi(t) > 0 \\ -1 & \phi(t) < 0 \\ ? & \phi(t) = 0 \end{cases}$$

So we see the control is bang-bang provided $\phi(t)$ does not equal zero for a closed interval $[s_0, s_1] \in [0, T^*]$. It may be zero at an isolated point, but not over a closed interval. Since f and g are assumed to be analytic, we'll see that a condition requiring $\phi(t) \neq 0$ over an interval is that all of its derivatives don't vanish over the interval. What we're interested in doing is identifying a condition on the vector fields f and g such that this *singular* case does not occur.

To do this, we need to examine the derivatives of the switching function ϕ. Let's first look at the first derivative,

$$\begin{aligned} \dot{\phi}(t) &= \langle \dot{p}^*(t), g(x^*(t)) \rangle + \langle p^*(t), g_x(x^*(t)) \dot{x}^*(t) \rangle \\ &= -\left.\langle (f_x)^T p^*, g \rangle\right|_* - \left.\langle (g_x)^T p^*, g \rangle\right|_* u^* \\ &\quad + \langle p^*, g_x f \rangle|_* + \langle p^*, g_x g \rangle| * u^* \\ &= \langle p^*, g_x f - f_x g \rangle|_* \\ &= \langle p^*(t), [f, g](x^*(t)) \rangle \end{aligned}$$

where $[f, g]$ is the Lie bracket of the two vector fields f and g.

For $\phi(t) = 0$ over an interval, we would also require $\dot{\phi}(t) = 0$ over that interval. The first conditions means that

$$0 = \phi(t) = \langle p^*, g(x^*) \rangle$$

The second condition on $\dot{\phi}$ requires

$$0 = \dot{\phi}(t) = \langle p^*(t), [f, g](x^*(t)) \rangle$$

Since p^* must be orthogonal to both g and $[f, g]$, we see this as being equivalent to the vector fields g and $[f, g]$ being linearly dependent. In other words, if the bang-bang property to hold in this class of nonlinear systems, we would expect g and $[f, g]$ to be linearly independent vector fields.

What we've just established is a necessary condition that must be satisfied if a bang-bang control exists for the time-optimal nonlinear system

$$\dot{x} = f(x) + g(x)u \tag{59}$$

To establish a stronger result concerning the existence of bang-bang controls for such systems, we actually need all of the derivatives of ϕ to go to zero, plus an additional restriction on the iterated Lie brackets generated by the vector fields f and g. This theorem is stated below.

We say the system in equation (59) is *analytic* if the state space of this system is an analytic manifold and the vector fields f and g are analytic (derivatives of all orders exist).

Let $[X, Y]$ be the Lie bracket of vector fields X and Y. Let $\mathrm{ad}X$ be the operator that assigns to each vector field Y, the vector field $[X, Y]$. We can apply the ad operator in an iterative manner so that

$$\begin{aligned}
\mathrm{ad}X &= [X, Y] \\
\mathrm{ad}^2X(Y) &= [X, \mathrm{ad}X] = [X, [X, Y]] \\
&\vdots \\
(\mathrm{ad}X)^n(Y) &= [X, \mathrm{ad}^{n-1}X] = [X, [\cdots, [X, Y]]\cdots]
\end{aligned}$$

Let $x \in M$ and let $m > 0$ be an integer. We say that the system in equation (59) satisfies condition $\Delta_{x,m}$ if and only if there exist analytic functions α_i and β with $|\beta(x)| < 1$ for all $u \in \mathcal{U}$ such that

$$[g, (\mathrm{ad}f)^m(g)] = \sum_{i=0}^{m}(\mathrm{ad}f)^i(g) + \beta(\mathrm{ad}f)^{m+1}(g)$$

We say that the system satisfies condition Δ if condition $\Delta_{x,m}$ is satisfied for all $x \in M$ and all m.

The main theorem regarding existence of bang-bang controls for such systems requires that the system satisfies condition Δ. The proof for this theorem may be found in a 1979 paper by Hector Sussman. Systems where $\phi(t) = 0$ over an interval may be referred to as *singular systems*.