

Pontryagin principle of maximum; time-optimal control

Constrained control, bang-bang control

*Zdeněk Hurák
April 4, 2018*

THE techniques of calculus of variations introduced in the previous lecture significantly extended our set of tools for solving optimal control problems—instead of optimizing over (finite) sequences of real numbers we are now able to optimize over functions. Nonetheless, the methods were also subject to severe restrictions. For example, the property that the optimal controls maximize the Hamiltonian were checked by testing the derivative of the Hamiltonian, which only makes sense if 1.) the Hamiltonian is differentiable with respect to control and 2.) the optimal control is inside the set of allowable controls (vanishing derivative of Hamiltonian on the boundary is not a necessary condition for the Hamiltonian to achieve a maximum value there). It also appears that the classes of perturbations characterized by 1-norm and even 0-norm are not rich enough for practical applications. Just consider switched (on-off) controls which differ in the times of switching. For all these reasons, some more advanced mathematical machinery has been developed. Unfortunately, it calls for a different and a bit more advanced mathematics than what we used in the calculus of variations. The mathematics here is indeed rather involved, therefore we will only state the most important result—the powerful Pontryagin’s principle of maximum. Although it sort of replaces some of the previous results (and you might start complaining why on earth we spent time with calculus of variations), having been introduced to the calculus of variations-style of reasoning we are now certainly well-equipped to digest at least the very statement of the powerful result by Pontryagin. If you are interested (and courageous), for a proof see the book [1] or elsewhere.

1 Pontryagin’s principle of maximum

We have already seen in the calculus of variations that the Hamiltonian when evaluated along the extremal, has the property that

$$H_{y'} = 0. \quad (1)$$

Combined with the second-order necessary condition of minimum

$$L_{y'y'} \geq 0, \quad (2)$$

we concluded that Hamiltonian is not only stationary along the extremal; it is actually maximized since

$$H_{y'y'} = -L_{y'y'} \leq 0. \quad (3)$$

This result can be written as

$$H(x, y^*, (y^*)', p^*) \geq H(x, y^*, y', p^*). \quad (4)$$

This is a major observation and would never be obtained without viewing y' as a separate variable (see [2] for an insightful discussion). After a notational transition to the optimal control setting and considering the augmented Lagrangian

$$L^{\text{aug}}(t, \mathbf{x}, \mathbf{u}, \dot{\mathbf{x}}, \boldsymbol{\lambda}) = L(t, \mathbf{x}, \mathbf{u}) + \boldsymbol{\lambda}^T \cdot (\dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}, \mathbf{u}, t)), \quad (5)$$

Hamiltonian can be written as

$$\begin{aligned} H(t, \mathbf{x}, \mathbf{u}, \dot{\mathbf{x}}, \boldsymbol{\lambda}) &= \left(\underbrace{L^{\text{aug}}}_{\boldsymbol{\lambda}} \right)^T \cdot \dot{\mathbf{x}} - L^{\text{aug}}, \\ &= \boldsymbol{\lambda}^T \cdot \dot{\mathbf{x}} - L(t, \mathbf{x}, \mathbf{u}) - \boldsymbol{\lambda}^T \cdot \dot{\mathbf{x}} + \boldsymbol{\lambda}^T \cdot \mathbf{f}(\mathbf{x}, \mathbf{u}, t), \\ &= \boldsymbol{\lambda}^T \cdot \mathbf{f}(\mathbf{x}, \mathbf{u}, t) - L(t, \mathbf{x}, \mathbf{u}). \end{aligned}$$

Thanks to the fact that $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, t)$, the Hamiltonian can be considered as a function of $t, \mathbf{x}, \mathbf{u}$ and $\boldsymbol{\lambda}$:

$$H(t, \mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}). \quad (6)$$

An important observation is that \mathbf{u} now plays a similar role to that of y' in calculus of variations. See more on this in the book.

If we now label the set of all allowable controls as \mathcal{U} , the result on maximization of Hamiltonian can then be written as

$$\boxed{H(t, \mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*) \geq H(t, \mathbf{x}^*, \mathbf{u}, \boldsymbol{\lambda}^*) \quad \forall \mathbf{u} \in \mathcal{U},} \quad (7)$$

or, equivalently as

$$\boxed{\mathbf{u}^* = \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} H(t, \mathbf{x}^*, \mathbf{u}, \boldsymbol{\lambda}^*),} \quad (8)$$

The essence of the celebrated Pontryagin's principle is that actually the above condition is the necessary condition of optimality. The fact that

$$H_{\mathbf{u}} = 0 \quad (9)$$

is just a consequence in the situation when $H_{\mathbf{u}}$ exists and the set of allowable controls u is not bounded. Let us emphasize the fact that the control u comes from some bounded set U by writing the Pontryagin's principle as

Theorem 1 (Pontryagin's principle of maximum). *For a given system and a given optimal criterion, let $\mathbf{u}^* \in \mathcal{U}$ be an optimal control, then there is a variable called costate which together with the state satisfies the Hamilton canonical equations*

$$\dot{\mathbf{x}} = \nabla_{\boldsymbol{\lambda}} H, \quad (10)$$

$$\dot{\boldsymbol{\lambda}} = -\nabla_{\mathbf{x}} H, \quad (11)$$

where

$$H(t, \mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = \boldsymbol{\lambda}^T(t) \cdot \mathbf{f}(\mathbf{x}, \mathbf{u}, t) - L(t, \mathbf{x}, \mathbf{u}) \quad (12)$$

and

$$H(t, \mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*) \geq H(t, \mathbf{x}^*, \mathbf{u}, \boldsymbol{\lambda}^*), \quad \mathbf{u} \in \mathcal{U}. \quad (13)$$

Moreover, *the corresponding boundary conditions must hold.*

As a matter of fact, Hamiltonian here is defined as

$$H(t, \mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = \boldsymbol{\lambda}^T(t) \cdot \mathbf{f}(\mathbf{x}, \mathbf{u}, t) - \lambda_0 L(t, \mathbf{x}, \mathbf{u}), \quad (14)$$

which allows for degenerate situations by setting $\lambda_0 = 0$.

We could certainly rederive our previous results on LQ-optimal control with fixed and free final states. Nonetheless, this would be an overkill unless we want to explore the bounded controls case. There is another scenario, where the Pontryagin's principle is immediately needed, and that is the *minimum-time problem*. The task is to bring the system to a given state in the shortest time possible. Apparently, with no bounds on controls, the time can be shrunk to zero (the control signal assuming the shape of a Dirac impulse). Therefore, bounds need to be imposed on the magnitudes of control signals in order to compute realistic outcomes.

In order to investigate this situation, we must first know how our necessary conditions change if we relax the final time, that is, the final time becomes one of the optimization variables.

2 Necessary conditions for a free final time

Setting the final time free means that we want to use the final time as yet another parameter for optimization.

2.1 Free final time and free final state

Example 2.1. *Maybe by late submission of your solution to a homework assignment in a general course, you can have a net gain in grading since the penalty for late submission may be more than compensated by a much higher grading for the increased quality of the solution. Time may be a parameter. Well, not so in our course :-)*

2.1.1 Calculus of variations setting

Let us return back to the calculus of variations setting with their notation. The problem we are going to solve is visualized in Fig. 1.




Figure 1: Optimizing over curves with one of their end point on a curve.

This trick is, that stretching or shrinking of the interval of the independent variable is done by perturbing the stationary value of the right end b of the interval with the same α as we use to perturb the functions y and y' . That is, b is perturbed by $\Delta b = \alpha \Delta x$ and the perturbed cost functional is then

$$J(y^* + \alpha \eta) = \int_a^{b+\alpha \Delta x} L(x, y^* + \alpha \eta, (y^*)' + \alpha \eta') dx. \quad (15)$$

Note that we have a minor technical problem here since y^* is only defined on the interval $[a, b]$. But there is an easy fix: we will define a continuation of the function even to the right of b in the form of a linear approximation given by the derivative of y^* at b . We will exploit it in a while.

Now, in order to find a variation δJ , we can either proceed by fitting the Taylor's expansion of the above perturbed cost function to the general Taylor's expansion and identifying the first-order term in α . Alternatively (well, in fact, equivalently), we can use the already stated fact that

$$\delta J = \left. \frac{d}{d\alpha} J(y^* + \alpha \eta) \right|_{\alpha=0} \alpha. \quad (16)$$

In order to find this derivative, we have to observe that the variable with respect to which we are differentiating is included in the upper bound of the integral. Therefore we cannot just change the order of differentiation and integration. This situation is handled by the well-known *Leibniz rule for differentiation under the integral sign*. Look it up yourself in the full generality. In our case it leads to

$$\left. \frac{d}{d\alpha} J(y^* + \alpha \eta) \right|_{\alpha=0} = \int_a^b \left(L_y - \frac{d}{dx} L_{y'} \right) \eta(x) dx + L_{y'}|_b \eta(b) + L|_b \Delta x, \quad (17)$$

which after multiplication by α gives the variation of the functional

$$\boxed{\delta J = \int_a^b \left(L_y - \frac{d}{dx} L_{y'} \right) \delta y(x) dx + L_{y'}|_b \delta y(b) + L|_b \underbrace{\Delta x \alpha}_{\Delta b}}, \quad (18)$$

where the first two terms on the right are already known to us. The only new is the third term. The reasoning now is pretty much the same as it was in the fixed-interval free-end case. We argue that among the variations δy there are also those that vanish at b , hence the conditions must be satisfied even if the last two terms are zero. But then the integral must be zero, which gives rise to the familiar Euler-Lagrange equation. The last two terms must together be zero and it does not hurt to rewrite them in a complete notation to dispell any ambiguity

$$\boxed{\left. \frac{\partial L(x, y(x), y'(x))}{\partial y'} \right|_{x=b} \delta y(b) + L(x, y(x), y'(x))|_{x=b} \Delta b = 0}. \quad (19)$$

Now, in order to get some more insight into the above condition, the relation between the participating objects can be further explored. We will do it using the Fig. 1 but we will augment it a bit with a few labels, see Fig. 2 below.

Note that we have included a brand new label here, namely, dy_f for the perturbation of the value of the function $y(\cdot)$ at the end of the interval (taking into consideration that the length of the interval can change as well). We can now write

$$y^*(b + \Delta b) + \delta y(b + \Delta b) = y^*(b) + dy_f, \quad (20)$$




Figure 2: Optimizing over curves with one end of the interval of the independent variable x set free and relaxing also the value of the function there.

which after approximating each term with the first two terms of its Taylor expansion gives

$$\cancel{y^*(b)} + y^{*'}(b)\Delta b + \delta y(b) + \cancel{\delta'(b)\Delta b} = \cancel{y^*(b)} + dy_f. \quad (21)$$

Note that the third product on the left can be neglected since it contains two terms that are both of order one in the perturbation variable α . In other words, we approximate $\delta y(b + \Delta b)$ by $\delta y(b)$. In addition, the term $y^*(b)$ can be subtracted from both sides. From what remains after these cancellations, we can conclude that

$$y^{*'}(b)\Delta b + \delta y(b) = dy_f, \quad (22)$$

or, equivalently,

$$\delta y(b) = dy_f - y^{*'}(b)\Delta b. \quad (23)$$

We will now substitute this into the general form of the boundary equation in (19)

$$L_{y'}(b, y(b), y'(b)) \cdot [dy_f - y^{*'}(b)\Delta b] + L(b, y(b), y'(b))\Delta b = 0. \quad (24)$$

Collecting now the terms with the two *independent* perturbation variables dy_f and Δb , we reformat the above expression into

$$\boxed{L_{y'}(b, y(b), y'(b))dy_f + [L(b, y(b), y'(b)) - L_{y'}(b, y(b), y'(b))y^{*'}(b)]\Delta b = 0.} \quad (25)$$

Now, since dy_f and Δb are assumed independent, the corresponding terms must be simultaneously and independently equal zero, that is,

$$L_{y'}(b, y(b), y'(b)) = 0, \quad (26)$$

$$L(b, y(b), y'(b)) - L_{y'}(b, y(b), y'(b))y^{*'}(b) = 0. \quad (27)$$

Note that the first condition actually constitutes n scalar conditions whereas the second one is just a scalar condition itself, hence, $n + 1$ boundary conditions.

2.1.2 Optimal control setting

Let's now switch back to the optimal control setting with t as the independent variable. Recall that the optimal control problem is

$$\min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot), t_f} \left[\phi(\mathbf{x}(t_f)) + \int_{t_i}^{t_f} L(\mathbf{x}, \mathbf{u}, t) dt \right]. \quad (28)$$

subject to

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}, \mathbf{u}, t), \quad \mathbf{x}(t_i) = \mathbf{r}_i. \quad (29)$$

We have already seen that the integrand of the augmented cost function now contains not only the term that corresponds to the Lagrange multiplier but also the term that penalizes the state at the final time, that is,

$$L^{\text{aug}}(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}, t) = L(\mathbf{x}, \mathbf{u}, t) + \underbrace{\frac{\partial \phi}{\partial t} + (\nabla_{\mathbf{x}} \phi)^{\text{T}} \frac{d\mathbf{x}}{dt}}_{\frac{d\phi(\mathbf{x}(t), t)}{dt} \Big|_{t=t_f}} + \boldsymbol{\lambda}(\dot{\mathbf{x}} - f(\mathbf{x}, \mathbf{u}, t)) \quad (30)$$

We then rewrite the boundary conditions (25) as

$$\boxed{(\nabla_{\mathbf{x}} \phi + \boldsymbol{\lambda}) \Big|_{t=t_f}^{\text{T}} d\mathbf{x}_f + \left(L + \frac{\partial \phi}{\partial t} - \boldsymbol{\lambda}^{\text{T}} f(\mathbf{x}, \mathbf{u}, t) \right) \Big|_{t=t_f} dt_f.} \quad (31)$$

Since here we assume that the final time and the state at the final time are independent, this single conditions breaks down into two boundary conditions¹

$$\nabla_{\mathbf{x}} \phi(x(t_f), t_f) + \boldsymbol{\lambda}(t_f) = 0 \quad (32)$$

$$L(x(t_f), u(t_f), t_f) + \frac{\partial \phi(\mathbf{x}(t_f), t_f)}{\partial t} - \boldsymbol{\lambda}^{\text{T}}(t_f) f(\mathbf{x}(t_f), \mathbf{u}(t_f), t_f) = 0. \quad (33)$$

The first one is actually representing n scalar conditions, the second one is just a single scalar condition. Hence, altogether we have $n + 1$ boundary conditions.

Let's try to get some more insight into this. Let's assume now that the term penalizing the state at the final time does not explicitly depend on time, that is, $\frac{\partial \phi}{\partial t} = 0$. Then the boundary condition modifies to

$$(\nabla_{\mathbf{x}} \phi + \boldsymbol{\lambda}) \Big|_{t=t_f}^{\text{T}} d\mathbf{x}_f + \left(L - \boldsymbol{\lambda}^{\text{T}} f(\mathbf{x}, \mathbf{u}, t) \right) \Big|_{t=t_f} dt_f, \quad (34)$$

which can be rewritten as

$$(\nabla_{\mathbf{x}} \phi + \boldsymbol{\lambda}) \Big|_{t=t_f}^{\text{T}} d\mathbf{x}_f - H(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}, t) \Big|_{t=t_f} dt_f, \quad (35)$$

which, in turn, enforces the scalar boundary condition (on top of those other n conditions)

$$\boxed{H(\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), t) \Big|_{t=t_f} dt_f = 0.} \quad (36)$$

This is an observation that is worth memorizing—for a free final time optimal control problem, Hamiltonian vanishes at the end of the time interval.

Let's now add one more observation. We could have mentioned it even in the previous lecture since it is a general property of a Hamiltonian evaluated along the

¹Note that here we commit the common abuse of notation in writing the functions to be differentiated as explicitly dependent on t_f such as in $\frac{\partial \phi(\mathbf{x}(t_f), t_f)}{\partial t}$. Instead, we should perhaps keep writing it as $\frac{\partial \phi(\mathbf{x}(t), t)}{\partial t} \Big|_{t=t_f}$ but it is tiring and the formulas look cluttered.

optimal solution—the total derivative of a Hamiltonian (evaluated along the solution) with respect to time is equal to its partial derivative with respect to time:

$$\frac{dH}{dt} = \underbrace{\frac{\partial H}{\partial \mathbf{x}} \frac{d\mathbf{x}}{dt}}_{-\dot{\lambda}} + \underbrace{\frac{\partial H}{\partial \boldsymbol{\lambda}} \frac{d\boldsymbol{\lambda}}{dt}}_{\dot{\mathbf{x}}} + \underbrace{\frac{\partial H}{\partial \mathbf{u}} \frac{d\mathbf{u}}{dt}}_0 + \frac{\partial H}{\partial t} = \frac{\partial H}{\partial t}. \quad (37)$$

Now, if neither the system equations nor the optimal control cost function depend *explicitly* on time, that is, if $\frac{\partial H}{\partial t} = 0$, the Hamiltonian remains constant along the optimal solution (trajectory), that is,

$$\boxed{H(\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\lambda}(t)) = \text{const.} \quad \forall t} \quad (38)$$

Combined with the previous result (boundary value of H at the end of the free time interval is zero), we obtain the powerful conclusion that the Hamiltonian evaluated along the optimal trajectory is always zero in the free final time scenario:

$$\boxed{H(\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\lambda}(t)) = 0 \quad \forall t} \quad (39)$$

This is a pretty insightful piece of information. Since some (numerical) techniques for optimal control are based on iterative minimization of a Hamiltonian, here we already know the minimum value.

Remark on notation In the previous lecture (notes) we already discussed the unfortunate discrepancy in the definitions of Hamiltonian in the literature. Perhaps there is no need to come back to this topic because you are now aware of the problem, but I will do it anyway. My only motivation is to have the formulas at hand.

Recall that the ambiguity starts with the definition of the augmented Lagrangian. I could have easily written instead of (30) the following

$$\hat{L}^{\text{aug}}(\mathbf{x}, \mathbf{u}, \hat{\boldsymbol{\lambda}}, t) = L(\mathbf{x}, \mathbf{u}, t) + \underbrace{\frac{\partial \phi}{\partial t} + (\nabla_{\mathbf{x}} \phi)^T \frac{d\mathbf{x}}{dt}}_{\left. \frac{d\phi(\mathbf{x}(t), t)}{dt} \right|_{t=t_f}} + \hat{\boldsymbol{\lambda}}^T (f(\mathbf{x}, \mathbf{u}, t) - \dot{\mathbf{x}}). \quad (40)$$

The boundary condition would then modify to

$$\boxed{(\nabla_{\mathbf{x}} \phi - \hat{\boldsymbol{\lambda}})^T \Big|_{t=t_f} d\mathbf{x}_f + \left(L + \frac{\partial \phi}{\partial t} + \hat{\boldsymbol{\lambda}}^T f(\mathbf{x}, \mathbf{u}, t) \right) \Big|_{t=t_f} dt_f,} \quad (41)$$

which can be rewritten in the case of $\frac{\partial \phi(\mathbf{x}(t_f), t_f)}{\partial t} = 0$ and using the alternative definition of the Hamiltonian $\hat{H} = L + \hat{\boldsymbol{\lambda}}^T \mathbf{f}$ as

$$(\nabla_{\mathbf{x}} \phi - \hat{\boldsymbol{\lambda}})^T \Big|_{t=t_f} d\mathbf{x}_f + \hat{H}(\mathbf{x}, \mathbf{u}, \hat{\boldsymbol{\lambda}}, t) \Big|_{t=t_f} dt_f, \quad (42)$$

2.2 Free final time but the final state on a prescribed curve

2.2.1 Calculus of variations setting

We will now investigate a special case when the final value of the solution $y(x)$ is to be on the curve described by $\psi(x)$, that is

$$y^*(b + \Delta b) + \delta y(b + \Delta b) = \psi(b + \Delta b). \quad (43)$$




Figure 3: Free final time on curve

This corresponds to the situation depicted in Fig. 3.

We already discussed the terms on the left. What is new here is the term on the right. It can also be approximated by the first two terms in Taylor's expansion

$$\psi(b + \Delta b) = \psi(b) + \psi'(b)\Delta b. \quad (44)$$

Therefore, we can expand (43) into

$$\cancel{y^*(b)} + (y^*)'(b)\Delta b + \delta y(b) = \underbrace{\cancel{y^*(b)}}_{\psi(b)} + \psi'(b)\Delta b, \quad (45)$$

from which we can express $\delta y(b)$ as

$$\delta y(b) = \psi'(b)\Delta b - (y^*)'(b)\Delta b \quad (46)$$

and substitute to the boundary condition (19), which after cancelling the common Δb term yields

$$\boxed{L_{y'}(b, y(b), y'(b)) \cdot [\psi'(b) - y'(b)] + L(b, y(b), y'(b)) = 0.} \quad (47)$$

This is just one scalar boundary conditions. But the original n conditions that the state that $y(x) = \psi(x)$ at the right end of the interval must be added. Altogether, we have $n + 1$ boundary conditions.

Anyway, the above single equation is called *transversality condition* for the reason to be illuminated by the next example.

Example 2.2. *To get an insight, consider again the minimum distance problem. This time we want to find the shortest distance from a point to a curve given by $\phi(x)$. The answer is intuitive, but let us see what our rigorous tools offer here. The EL equation stays intact, therefore we know that the shortest path is a line. It starts at $(a, 0)$ but in order to determine its end, we need to invoke the other boundary condition. Remember that the Lagrangian is*

$$L = \sqrt{1 + (y')^2} dx \quad (48)$$

and

$$L_{y'} = \frac{y'}{\sqrt{1 + (y')^2}} dx. \quad (49)$$

The transversality condition boils down to

$$1 + y'(b)\psi'(b) = 0, \quad (50)$$

which can also be visualized using vectors in the plane

$$\begin{bmatrix} 1 & y'(b) \end{bmatrix} \cdot \begin{bmatrix} 1 \\ \psi'(b) \end{bmatrix} = 0. \quad (51)$$

The interpretation of this result is that our desired curve y hits the target curve ϕ in a perpendicular (transverse) direction.

Understanding the boundary conditions is crucial. Let us have yet another look at the result just derived. It can be written as

$$\boxed{L_{y'}\psi'|_b - H|_b = 0.} \quad (52)$$

It follows that for a free length of the interval and fixed value of the variable at the end of the interval, in which

$$\psi(x) = c, \quad c \in \mathbb{R}, \quad (53)$$

the transversality condition simplifies to

$$\boxed{H(b) = 0.} \quad (54)$$

2.2.2 Optimal control setting

Once again, let's recall that the optimal control problem is

$$\min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot), t_f} \left[\phi(\mathbf{x}(t_f)) + \int_{t_i}^{t_f} L(\mathbf{x}, \mathbf{u}, t) dt \right]. \quad (55)$$

subject to

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}, \mathbf{u}, t), \quad (56)$$

$$\mathbf{x}(t_i) = \mathbf{r}_i \quad (57)$$

$$\mathbf{x}(t_f) = \boldsymbol{\psi}(t_f). \quad (58)$$

Translating the above derived transversality condition from the domain (and notation) of calculus of variations into the optimal control setting gives

$$\boxed{(\nabla_{\mathbf{x}}\phi + \boldsymbol{\lambda})^T \dot{\boldsymbol{\psi}}(t) + L + \frac{\partial \phi}{\partial t} - \boldsymbol{\lambda}^T f(\mathbf{x}, \mathbf{u}, t) \Big|_{t=t_f} = 0.} \quad (59)$$

Of course, as usual, on top of this single condition, the $2n$ boundary conditions shown above must be added.

3 Time-optimal control for a linear system—bang-bang control

The task of bringing the system from a given state to some given final state (either a single state or a set of states) can be formulated by setting

$$L = 1, \quad (60)$$

which turns the cost functional to

$$J = \int_{t_i}^{t_f} 1 dt = (t_f - t_i). \quad (61)$$

Let us solve the task for an LTI system

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}, \quad t_i = 0, \quad \mathbf{x}(t_i) = \mathbf{r}_0 \quad (62)$$

for which we set the final desired state as

$$\mathbf{x}(t_f) = \mathbf{0}. \quad (63)$$

As already discussed, this only makes sense if we impose some bounds on the control, therefore

$$|u_i(t)| \leq 1 \quad \forall i. \quad (64)$$

The necessary conditions can be build immediately by forming Hamiltonian as

$$H = \boldsymbol{\lambda}^T \cdot (\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}) - 1 \quad (65)$$

and substituting into the Hamilton canonical equations

$$\dot{\mathbf{x}} = \nabla_{\mathbf{x}} H = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}, \quad (66)$$

$$\dot{\boldsymbol{\lambda}} = -\nabla_{\mathbf{x}} H = -\mathbf{A}^T \boldsymbol{\lambda}. \quad (67)$$

plus the Pontryagin's statement about maximization of H with respect to \mathbf{u}

$$H(t, \mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*) \geq H(t, \mathbf{x}^*, \mathbf{u}, \boldsymbol{\lambda}^*), \quad u_i(t) \in [-1, 1] \quad \forall t. \quad (68)$$

Application of Pontryagin's principle gives

$$(\boldsymbol{\lambda}^*)^T \cdot (\mathbf{A}\mathbf{x}^* + \mathbf{B}\mathbf{u}^*) - 1 \geq (\boldsymbol{\lambda}^*)^T \cdot (\mathbf{A}\mathbf{x}^* + \mathbf{B}\mathbf{u}) - 1, \quad u_i \in [-1, 1]. \quad (69)$$

Cancelling the identical terms on both sides we are left with

$$(\boldsymbol{\lambda}^*)^T \cdot \mathbf{B}\mathbf{u}^* \geq (\boldsymbol{\lambda}^*)^T \cdot \mathbf{B}\mathbf{u}, \quad u_i \in [-1, 1]. \quad (70)$$

It turns out that if this inequality is to hold then with the \mathbf{u} arbitrary on the left (within the bounds), the only way to guarantee the validity is to have

$$\mathbf{u}^* = \text{sgn}((\boldsymbol{\lambda}^*)^T \cdot \mathbf{B}), \quad (71)$$

where the signum function is applied elementwise. Clearly the optimal control is *switching*—it only assumes values of 1 or -1. This is visualized in Fig. 4 for a scalar case (the \mathbf{B} matrix has only a single column).

Well, in fact to support this claim, it must be rigorously excluded that the argument of the signum function, the so-called switching function can assume zero value for longer then just a time instant (although repeatedly). Check this by yourself in [1] (or its online version). Search for *normality conditions*.




Figure 4: Switching function and an optimal control derived from it.

3.1 Time-optimal control for a double integrator system

Let us analyze the situation for a double integrator. This corresponds to a system described by the second Newton's law. For a normalized mass the state space model is

$$\begin{bmatrix} \dot{y} \\ \dot{v} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} y \\ v \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u. \quad (72)$$

The switching function is obviously $\lambda_2(t)$ and an optimal control is given by

$$u(t) = \text{sgn} \lambda_2(t). \quad (73)$$

We do not know $\lambda_2(t)$. In order to get it, we may need to solve the costate equations. Indeed, we can solve them independently of the state equations since it is decoupled from them

$$\begin{bmatrix} \dot{\lambda}_1 \\ \dot{\lambda}_2 \end{bmatrix} = - \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix}, \quad (74)$$

from which it follows that

$$\lambda_1(t) = c_1 \quad (75)$$

and

$$\lambda_2(t) = c_1 t + c_2. \quad (76)$$

for some constants c_1 and c_2 . To determine the constants, we will have to bring the boundary conditions finally into the game. The condition that $H(t_f) = 0$ gives

$$\lambda_2(t_f)u(t_f) = 1. \quad (77)$$

We can now sketch possible profiles of the switching function. A few characteristic versions are in Fig. 5

What we have learnt is that the costate λ_2 would go through zero at most once during the whole control interval. Therefore we will have at most one switching of the control signal. This is a valuable observation.

We are approaching the final stage of the derivations. So far we have learnt that we can only consider $u(t) = 1$ and $u(t) = -1$. The state equations can be easily integrated to get

$$v(t) = v(0) + ut, \quad y(t) = y(0) + v(0)t + \frac{1}{2}ut^2$$




Figure 5: Possible evolutions of the costate in time-optimal control.

To visualize this in $y - v$ domain, express t from the first and substitute into the second equation

$$u(y - y(0)) = v(0)(v - v(0)) + \frac{1}{2}(v - v(0))^2,$$

which is a family of parabolas parameterized by $(y(0), v(0))$. These are visualized in Fig. 6.

There is a distinguished curve in the figure, which is composed of two branches. It is special in that for all the states starting on this curve, the system is brought to the origin for a corresponding setting of the control (and no further switching). This curve, called *switching curve* can be expressed as

$$y = \begin{cases} \frac{1}{2}v^2 & v < 0 \\ -\frac{1}{2}v^2 & v > 0 \end{cases}$$

or

$$y = -\frac{1}{2}v|v|$$

The final step can be done using this figure. Point your finger anywhere in the plane. Follow the state trajectory that emanates from the particular point for which you can get to the origin with at maximum 1 switching. Clearly the strategy is to set u such that it brings us to the switching curve (the red one in the figure) and then follow it (after switching). That is it. This control strategy can be written as

$$u(t) = \begin{cases} -1 & \text{if } y(t) > -\frac{1}{2}v(t)|v(t)| \text{ or if } y(t) = -\frac{1}{2}v(t)|v(t)| \text{ and } y < 0 \\ 1 & \text{if } y(t) < -\frac{1}{2}v(t)|v(t)| \text{ or if } y(t) = -\frac{1}{2}v(t)|v(t)| \text{ and } y > 0 \end{cases},$$

which can be written in a compact form as

$$\boxed{u(t) = \text{sign} \left(-\frac{1}{2}v(t)|v(t)| - y(t) \right)}. \tag{78}$$

A simulation scheme in Simulink is in Fig. 7. and the expected simulated optimal response is in Fig. 8.

In the plots you can find a confirmation of the fact that we derived rigorously—the fact that there will be at most one switch in the control signal...Oops... This is




Figure 6: Typical trajectories for both $u(t) = 1$ and $u(t) = -1$. Red is the switching curve.




Figure 7: The structure of the time-optimal controller.

actually not quite what we see in the plot above, is it? We can see two switches in the control signals. The first one happens at about 2.2 s, but the second happens close to 3.5 s. In fact, what you would experience if you run the code in Simulink is an error statement that “At time 3.449489782192745, simulation hits (1000) consecutive zero crossings.” and the simulation will be finished. Obviously, what is going on is that the simulator is tempted to include not just two but in fact a huge number of switches in the control signal as it approaches the origin. This is quite characteristic of bang-bang control—a phenomenon called *chattering*. In this particular example you may decide to ignore it since both state variables are already close enough to the origin and you may want to declare the control task as finished². Generally, this chattering phenomenon needs to be handled somehow. Any suggestion how to reduce it?

4 Further reading

This lecture was prepared using [1], in particular chapters 3 (application of calculus of variations to general problem of optimal control) and chapter 4 (Pontryagin’s princi-

²Remember that we still consider a control over finite time interval, even though its length is a tunable parameter. Hence, after reaching the end of the interval, the task is over.

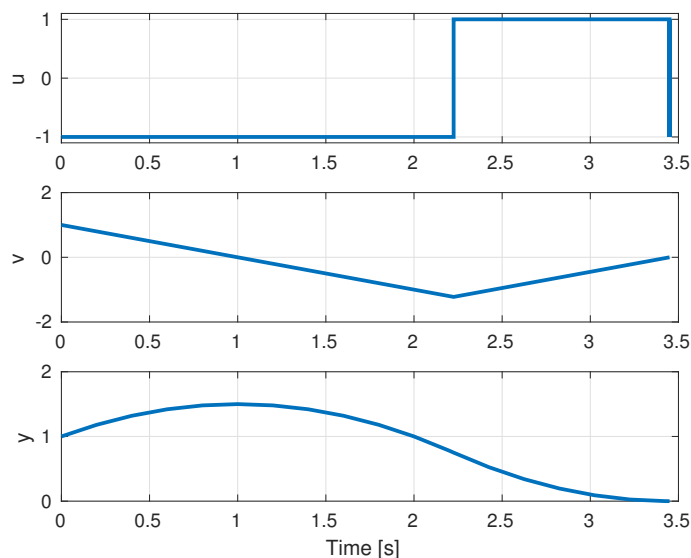


Figure 8: Time optimal response obtained from numerical simulation in Simulink.

ple). We did not talk about the proof of Pontryagin’s principle at the lecture and we do not even command the students to go through the proof in the book. Understanding the result, its roots in calculus of variations and how it removes the deficiencies of the calculus of variations based results will suffice for our purposes.

The transition from the calculus of variations to the optimal control, especially when it comes to the definition of Hamiltonian, is somewhat tricky. Unfortunately, it is not discussed satisfactorily in the literature. Even Liberzon leaves it as an (unsolved) exercise (3.5 and 3.6) to the student. Other major textbooks avoid the topic altogether. The only treatment can be found in the famous journal paper [2], in particular the section “The first fork in the road: Hamilton” on page 39. The issue is so delicate that they even propose to distinguish the two types of Hamiltonian by referring to one as *control Hamiltonian*.

The time-optimal control for linear systems, in particular bang-bang control for a double integrator is described in section 4.4.1 and 4.4.2. The material is quite standard and can be found in many books and lecture notes. What is not covered, however, is the fact that without any adjustment, the bang bang control is very troublesome from an implementation viewpoint. A dedicated research thread has evolved, especially driven by the needs of hard disk drive industry, which is called (*a*)*proximate time-optimal control (PTOS)*. Many dozens of papers can be found with this keyword in the title.

References

- [1] Daniel Liberzon. *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, December 2011.
- [2] H.J. Sussmann and J.C. Willems. 300 years of optimal control: from the brachystochrone to the maximum principle. *IEEE Control Systems*, 17(3):32–44, 1997.