The Rationale of the Maximum Principle

Prof. Luciano Nakabashi

A Variational View of the Control Problem

- To make things simple, it is assumed here that the control variable u is unconstrained, so that u^* is an interior solution.
- Moreover, the Hamiltonian function is assumed to be differentiable with respect to u, and the $\partial H/\partial u = 0$ condition can be invoked in place of the condition $\underset{u}{\mathrm{Max}}\, H$.
- We take the initial point to be fixed, but the terminal point is allowed to vary. The problem is then to

$$
\begin{array}{llll}
 & \text{Maximize} & \mathrm{V} = \int_0^T F(t, y, u)dt & \\
(5) & \text{Subject to} & \dot{y} = f(t, y, u) & \\
 & y(0) = y_0, & \mathrm{y}(T) = free & (T\ given)
\end{array}
$$

A Variational View of the Control Problem

- **Step I** - As the first step in the development of the maximum principle, let us incorporate the equation of motion into the objective functional, and then express the functional in terms of the Hamiltonian.
- If the variable y always obeys the equation of motion, then the quantity $[\boldsymbol{f}(\boldsymbol{t}, \boldsymbol{y}, \boldsymbol{u}) - \dot{\boldsymbol{y}}]$ will take a **zero value** for all t in the interval $[\mathbf{0}, \boldsymbol{T}]$.
- Thus, using the notion of Lagrange multipliers, we can form an expression $\boldsymbol{\lambda}(\boldsymbol{t})[\boldsymbol{f}(\boldsymbol{t}, \boldsymbol{y}, \boldsymbol{u}) - \dot{\boldsymbol{y}}]$ for each value of t, and still get **a zero value**.
- Summing $\lambda(t)[f(t, y, u) - \dot{y}]$ over t in the interval $[0, T]$ would still yield a total value of zero:

(6) $\int_0^T \lambda(t)[f(t, y, u) - \dot{y}]dt = 0$

A Variational View of the Control Problem

- For this reason, we can augment the old objective functional by the integral in (6) without affecting the solution. That is, we can work with the new objective functional

$$\mathcal{V} = \mathrm{V} + \int_0^T \lambda(t)[f(t, y, u) - \dot{y}]dt$$

(7) $$\mathcal{V} = \int_0^T F(t, y, u)dt + \int_0^T \lambda(t)[f(t, y, u) - \dot{y}]dt$$

$$\mathcal{V} = \int_0^T \{F(t, y, u) + \lambda(t)[f(t, y, u) - \dot{y}]\}dt$$

- Previously, we have defined the Hamiltonian function as

(8) $$H(t, y, u, \lambda) \equiv F(t, y, u) + \lambda(t)f(t, y, u)$$

A Variational View of the Control Problem

- Using (8) into (7):

(9) $\mathcal{V} = \int_0^T [H(t, y, u, \lambda) - \lambda(t)\dot{y}]dt$

$$\mathcal{V} = \int_0^T H(t, y, u, \lambda)dt - \int_0^T \lambda(t)\dot{y}dt$$

- Considering the integral by parts, as before:

(10) $\int_a^b v dw = vw|_{t=a}^{t=b} - \int_{t=a}^{t=b} w dv$

Consider that:

$\boldsymbol{v} = \boldsymbol{\lambda}(\boldsymbol{t})$ and $\boldsymbol{w} = \boldsymbol{y}(\boldsymbol{t})$, implying $\boldsymbol{dv} = \dot{\boldsymbol{\lambda}}(\boldsymbol{t})\boldsymbol{dt}$ and $\boldsymbol{dw} = \dot{\boldsymbol{y}}(\boldsymbol{t})\boldsymbol{dt}$:

$$\int_0^T \lambda(t)\dot{y}(t)dt = \lambda(t)y(t)|_0^T - \int_0^T y(t)\dot{\lambda}(t)dt$$

(11) $\int_0^T \lambda(t)\dot{y}(t)dt = \lambda(T)y_T - \lambda(0)y_0 - \int_0^T y(t)\dot{\lambda}(t)dt$

A Variational View of the Control Problem

- Using (11) into (9):

(12) $\mathcal{V} = \int_0^T H(t, y, u, \lambda)dt - \lambda(T)y_T + \lambda(0)y_0 + \int_0^T y(t)\dot{\lambda}(t)dt$

(13) $\mathcal{V} = \int_0^T \underbrace{\left[H(t, y, u, \lambda) + y(t)\dot{\lambda}(t)\right]}_{\Omega_1} dt - \underbrace{\lambda(T)y_T}_{\Omega_2} + \underbrace{\lambda(0)y_0}_{\Omega_3}$

- The $\mathcal{V}$ expression comprises three additive component terms, Ω_1, Ω_2, and Ω_3. Note that while the Ω_1 term, an integral, spans the entire planning period $[0, T]$, the Ω_2 term is exclusively concerned with the terminal time T, Ω_3 is concerned only with the initial time 0.
- **Step II -** The value of $\mathcal{V}$ depends on the time paths chosen for the three variables y, u, and λ, as well as the values chosen for T and y_T.

A Variational View of the Control Problem

- The choice of the $\lambda(t)$ path will produce no effect on the value of $\mathcal{V}$, so long as the equation of motion $\dot{y} = f(t, y, u)$ is strictly adhered to, that is, so long as

  (14) $\dot{y} = \frac{\partial H}{\partial \lambda}$ for all $t \in [0, T]$

- We impose (14) as a necessary condition for the maximization of $\mathcal{V}$, accounting for one of the three conditions of the maximum principle.
- **Step III** - We turn to the $u(t)$ path and its effect on the $y(t)$ path.
- If we have a known $u^*(t)$ path, and if we perturb the $u^*(t)$ path with a perturbing curve $p(t)$, we can generate "neighboring" control paths

  (15) $u(t) = u^*(t) + \epsilon p(t)$

A Variational View of the Control Problem

- The neighboring y paths can be written as

(16) $y(t) = y^*(t) + \epsilon q(t)$

- Furthermore, if T and y_T are variable, we also have

(17) $T = T^* + \epsilon \Delta T$ and $y_T = y_T^* + \epsilon \Delta y_T$, implying that:

(18) $\frac{dT}{d\epsilon} = \Delta T,$ and $\frac{dy_T}{d\epsilon} = \Delta y_T$

- In view of the u and y expressions in (15) and (16), we can express $\mathcal{V}$ in terms of ϵ, so that we can again apply the first-order condition $d\mathcal{V}/d\epsilon = 0$. The new version of $\mathcal{V}$ is

(19)

$$\mathcal{V} = \int_0^{T(\epsilon)} \left\{ H\left(t, \underbrace{y^* + \epsilon q(t)}_{y(t)}, \underbrace{u^*(t) + \epsilon p(t)}_{u(t)}, \lambda\right) + \dot{\lambda}(t)\left[\underbrace{y^* + \epsilon q(t)}_{y(t)}\right]\right\} dt - \lambda(T) y_T + \lambda(0) y_0$$

A Variational View of the Control Problem

Step IV - We now apply the condition $d\mathcal{V}/d\epsilon = 0$.

- In the differentiation process, the integral term yields the derivative:

(20) $$\frac{d\int_0^{T(\epsilon)} \ldots dt}{d\epsilon} = \int_0^{T(\epsilon)} \left[\left(\frac{\partial H}{\partial y}\frac{dy}{d\epsilon} + \frac{\partial H}{\partial u}\frac{du}{d\epsilon}\right) + \dot{\lambda}(t)q(t)\right] dt + \left[H + y\dot{\lambda}\right]_{t=T} \frac{dT}{d\epsilon}$$

(21) $$\frac{d\int_0^{T(\epsilon)} \ldots dt}{d\epsilon} = \int_0^{T(\epsilon)} \left\{\left[\frac{\partial H}{\partial y} q(t) + \frac{\partial H}{\partial u} p(t)\right] + \dot{\lambda}(t)q(t)\right\} dt + \left[H + y\dot{\lambda}\right]_{t=T} \frac{dT}{d\epsilon}$$

- And the derivative of the second term in (19) with respect to ϵ is, from the product rule,

(22) $$\frac{d\boldsymbol{\lambda}(\boldsymbol{T})\boldsymbol{y_T}}{d\epsilon} = \frac{\partial \lambda(T)}{\partial T}\frac{dT}{d\epsilon} y_T + \frac{dy_T}{d\epsilon}\lambda(T) = \dot{\lambda}(T)\Delta T y_T + \Delta y_T \lambda(T)$$

A Variational View of the Control Problem

- Since $y(0) = y_0$, and $t_0 = 0$ are given, the derivative of the third term in (19) with respect to ϵ drops out.
- Thus, when the sum of (21) and (22) is set equal to zero, the first-order condition emerges as

$$\frac{d\mathcal{V}}{d\epsilon} = \int_0^{T(\epsilon)} \left\{ \left[\frac{\partial H}{\partial y} q(t) + \frac{\partial H}{\partial u} p(t) \right] + \dot{\lambda} q(t) \right\} dt + \left[H + y\dot{\lambda} \right]_{t=T} \Delta T - \dot{\lambda}(T) \Delta T y_T - \Delta y_T \lambda(T) = 0$$

(23) $$\frac{d\mathcal{V}}{d\epsilon} = \int_0^{T(\epsilon)} \left[\left(\frac{\partial H}{\partial y} + \dot{\lambda} \right) q(t) + \frac{\partial H}{\partial u} p(t) \right] dt + [H]_{t=T} \Delta T - \lambda(T) \Delta y_T = 0$$

- The three components of this derivative relate to different arbitrary things.
- The integral contains arbitrary perturbing curves $p(t)$ and $q(t)$.
- The other two involve arbitrary ΔT and Δy_T, respectively. Consequently, each of the three must individually be set equal to zero in order to satisfy (23).

A Variational View of the Control Problem

- By setting the integral component equal to zero, we can deduce two conditions:

(24) $\int_0^{T(\epsilon)} \left[\left(\frac{\partial H}{\partial y} + \dot{\lambda} \right) q(t) + \frac{\partial H}{\partial u} p(t) \right] dt = 0$

(25) $\frac{\partial H}{\partial y} + \dot{\lambda} = 0 \Rightarrow \dot{\lambda} = -\frac{\partial H}{\partial y}$ and $\frac{\partial H}{\partial u} = 0$

- The first term in (25) gives us the equation of motion for the costate variable λ (or the costate equation for short).
- And the second term represents a weaker version of the "$\underset{u}{\mathrm{Max}}\, H$ " condition.
- Weaker in the sense that it is predicated on the u assumption that H is differentiable with respect to u and there is an interior solution.

A Variational View of the Control Problem

- Since the simplest problem has a fixed T and free y_T, the ΔT term in (23) is automatically equal to zero, but Δy_T is not.
- In order to make the $\lambda(T)\Delta y_T$ expression vanish, we must impose the restriction

  (26) $\lambda(T) = 0$
- This explains the transversality condition in (4).
- Note that although the $\lambda(T)$ path was earlier, in **Step II**, brushed aside as having no effect on the value of the objective functional, it has now made an impressive comeback.
- In order for the maximum principle to work, the $\boldsymbol{\lambda(T)}$ **path** is not to be arbitrarily chosen, but is required to **follow a prescribed equation of motion**, and it must end with a **terminal value of zero** with a free terminal state.

Fixed Terminal Point

- The reason why the problem with a **fixed terminal point** (with both the terminal state and the terminal time fixed) does not qualify as the **"simplest" problem** in optimal control theory is that the specification of a fixed terminal point entails a complication in the notion of an **"arbitrary" perturbing curve $\boldsymbol{p(t)}$ for the control variable $\boldsymbol{u}$**.
- If the perturbation of the $\boldsymbol{u^*(t)}$ **path is supposed to generate** through the equation of motion $\dot{y} = f(t, y, u)$ a corresponding **perturbation in the $\boldsymbol{y^*(t)}$ path** that **has to end at a preset terminal state**, then the choice of the perturbing curve $p(t)$ **is not truly arbitrary**.
- The question then arises as to whether we can still legitimately deduce the condition $\frac{\partial H}{\partial u} = 0$ from (23).

Fixed Terminal Point

- Fortunately, **the validity of the maximum principle is not affected by this compromise in the arbitrariness of $\boldsymbol{p(t)}$**.
- For simplicity, however, we shall not go into details to demonstrate this point.
- For our purposes, it suffices to state that, with a fixed terminal point, the transversality is replaced by the condition:

  (27) $y(T) = y_T \qquad\qquad (T, y_T\ given)$

Horizontal Terminal Line (Fixed Endpoint Problem)

- If the problem has a horizontal terminal line (with a free terminal time but a fixed "endpoint," meaning a fixed terminal state), then y_T is fixed $(\Delta \boldsymbol{y_T} = \mathbf{0})$, but T is not (ΔT is arbitrary).
- From the second and third component terms in **(23)** $[\boldsymbol{H}]_{\boldsymbol{t=T}}\Delta \boldsymbol{T} - \boldsymbol{\lambda}(\boldsymbol{T})\Delta \boldsymbol{y_T} = 0$, it is easy to see that the transversality condition for this case is

  (28) $[H]_{t=T} = 0$
- The Hamiltonian function must attain a zero value at the optimal terminal time. But there is no restriction on the value of λ at time T.

Terminal Curve

- In case a terminal curve $y_T = \phi(T)$ governs the selection of the terminal point.
- Then $\Delta \boldsymbol{T}$ **and** $\Delta \boldsymbol{y_T}$ **are not both arbitrary**, but are **linked** to each other by the relation $\Delta \boldsymbol{y_T} = \boldsymbol{\phi}'(\boldsymbol{T})\Delta \boldsymbol{T}$.
- Using this to eliminate Δy_T, we can combine the last two terms in (23) into a single expression involving ΔT only:

(29) $[H]_{t=T}\Delta T - \lambda(T)\Delta y_T = [H]_{t=T}\Delta T - \lambda(T)\phi'(T)\Delta T = [H - \lambda\phi']_{t=T}\Delta T$

- It follows that, for an arbitrary ΔT, **the transversality condition** should be

(30) $[H - \lambda\phi']_{t=T} = 0$

Truncated Vertical Terminal Line

- Now consider the problem in which the **terminal time $\boldsymbol{T}$ is fixed**, but the terminal state is free to vary, only **subject to $\boldsymbol{y_T \geq y_{min}}$**, where y_{min} denotes a given minimum permissible level of y.
- Only two types of outcome are possible in the optimal solution:

1. $y_T^* > y_{min}$;
2. $y_T^* = y_{min}$

- In the first outcome, the terminal restriction is automatically satisfied. Thus, the transversality condition for the problem with a regular vertical terminal line would apply:

  (31) $\lambda(T) = 0$ for $y_T^* > y_{min}$

Truncated Vertical Terminal Line

- In the other outcome, $y_T^* = y_{min}$, since the terminal restriction is binding, the admissible neighboring y paths consist only of those that have terminal states $y_T \geq y_{min}$.
- If we evaluate (16) $y(t) = y^*(t) + \epsilon q(t)$ at $t = T$ and let $y_T^* = y_{min}$:

(32) $y_T = y_{min} + \epsilon q(T)$

- Assuming that $\boldsymbol{q(T) > 0}$ on the perturbing curve $q(t)$, the requirement $\boldsymbol{y_T \geq y_{min}}$ **would dictate that** $\boldsymbol{\epsilon \geq 0}$**.**
- By the Kuhn-Tucker conditions, the nonnegativity of ϵ **would alter the first-order condition** $\boldsymbol{d\mathcal{V}/d\epsilon = 0}$ **to** $\boldsymbol{d\mathcal{V}/d\epsilon \leq 0}$ for our maximization problem.

Truncated Vertical Terminal Line

- It follows that (23) would yield an inequality transversality condition

  (34) $[H]_{t=T}\Delta T - \lambda(T)\Delta y_T \leq 0$

- Since $\Delta T = 0$, the transversality condition becomes

  (35) $-\lambda(T)\Delta y_T \leq 0$

At the same time, we can see from (17) $y_T = y_T^* + \epsilon\Delta y_T$, **given $\boldsymbol{\epsilon} \geq \mathbf{0}$**, the requirement of $\boldsymbol{y_T} \geq \boldsymbol{y_{min}}$ $(y_T^* = y_{min})$ implies $\Delta\boldsymbol{y_T} \geq \mathbf{0}$.

- Thus the preceding inequality transversality condition reduces to

  (36) $\lambda(T) \geq 0$ for $y_T^* = y_{min}$

Truncated Vertical Terminal Line

- Combining (31) and (36) and omitting the * symbol, we can write a single summary statement of the transversality condition as follows:

(37) $\lambda(T) \geq 0; \qquad y_T \geq y_{min}; \qquad (y_T - y_{min})\lambda(T) = 0$

- Note that the last part of this statement represents the familiar complementary-slackness condition from the Kuhn-Tucker conditions. The application of (37) is not as complicated as the condition may appear. We can try the $\lambda(T) = 0$ condition first, and check whether the resulting y_T^* value satisfies the terminal restriction $y_T^* \geq y_{min}$.
- If it does, the problem is solved. If not, we then set $y_T^* = y_{min}$ in order to satisfy the complementary-slackness condition, and treat the problem as one with a given terminal point.

Truncated Horizontal Terminal Line

- Let the terminal state be fixed, but the terminal time T be allowed to vary subject to the restriction that $T^* \leq T_{max}$, where T_{max} is the maximum permissible value of T.
- Then we either have

1. $T^* < T_{max}$;
2. $T^* = T_{max}$

in the optimal solution. In the first outcome, the terminal restriction turns out to be nonbinding, and the transversality condition for the problem with a regular horizontal terminal line would still hold:

(38) $[H]_{t=T}\Delta T - \lambda(T)\Delta y_T = 0 \implies [H]_{t=T} = 0$ for $T^* < T_{max}$

since $\Delta y_T = 0$ and ΔT is arbitrary.

Truncated Horizontal Terminal Line

- But if $T^* = T_{max}$, then by implication all the admissible neighboring y paths must have terminal time $T \leq T_{max}$.
- By analogous reasoning to that leading to the result (36) for the truncated vertical terminal line, it is possible to establish the transversality condition

  (39) $[H]_{t=T} \geq 0$ for $T^* = T_{max}$
- By combining (38) and (39) and omitting the * symbol, we obtain the following summary statement of the transversality condition:

  (40) $[H]_{t=T} \geq 0$ for $T \leq T_{max}$ $(T - T_{max})[H]_{t=T} = 0$