Optimal Control: The Maximum Principle

Prof. Luciano Nakabashi

Optimal Control: The Maximum Principle

- The calculus of variations, the classical method for tackling problems of dynamic optimization, like the ordinary calculus, requires for its applicability the differentiability of the functions that enter in the problem.
- More importantly, only interior solutions can be handled.
- A more modern development that can deal with nonclassical features such as corner solutions, is found in optimal control theory.
- As its name implies, the optimal-control formulation of a dynamic optimization problem focuses upon one or more control variables that serve as the instrument of optimization.

Optimal Control: The Maximum Principle

- Unlike the calculus of variations, where the goal is to find the optimal time path for a state variable y, optimal control theory has as its foremost aim the determination of the optimal time path for a control variable, u.
- Once the optimal control path, $u^*(t)$, has been found, it is possible to find the optimal state path, $y^*(t)$, that corresponds to it.
- A control variable possesses the following two properties:

1) It is something that is subject to our discretionary choice.
2) The choice of the control variable impinges upon the state variable. Consequently, the control variable is like a steering mechanism which we can maneuver so as to "drive" the state variable to various positions at any time t.

The Simplest Problem of Optimal Control

- To keep the introductory framework simple, we first consider a problem with a single state variable y and a single control variable u.
- Any chosen control path $u(t)$ will imply an associated state path $y(t)$.
- Our task is to choose the optimal admissible control path $u^*(t)$ which, along with the associated optimal admissible state path $y^*(t)$, optimize the objective functional over a time interval $[0, T]$.
- A control path does not have to be continuous in order to become admissible; it only needs to be piecewise continuous.
- This means that it is allowed to contain jump discontinuities, although no discontinuities that involve an infinite value of $u(t)$.
- The state path $y(t)$ have to be continuous throughout $[0, T]$. But it is permissible to have a finite number of sharp points, or corners.

Figure 1

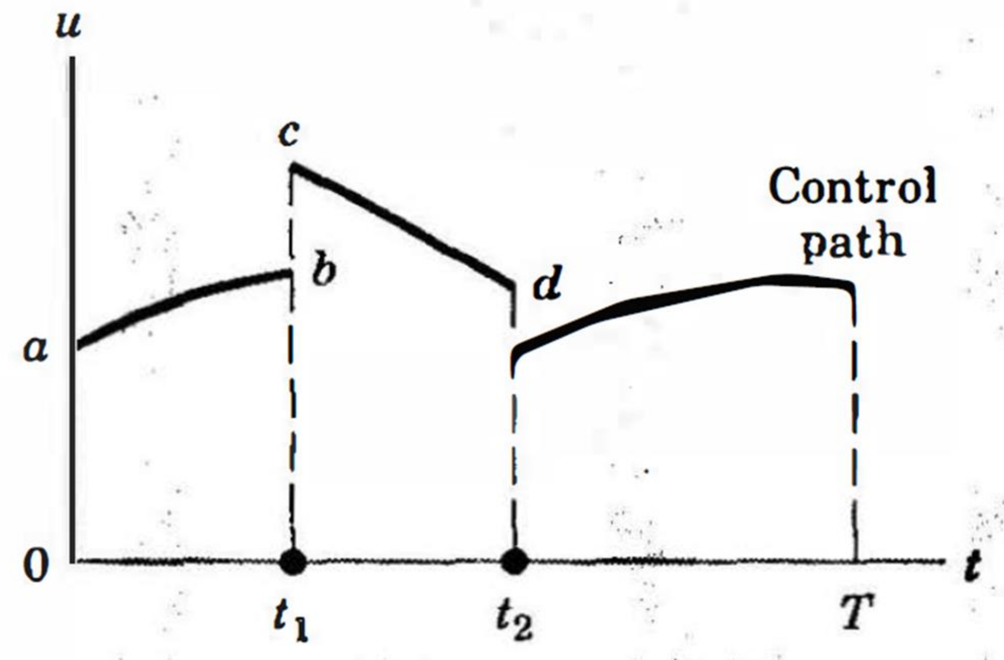

Figure 1a

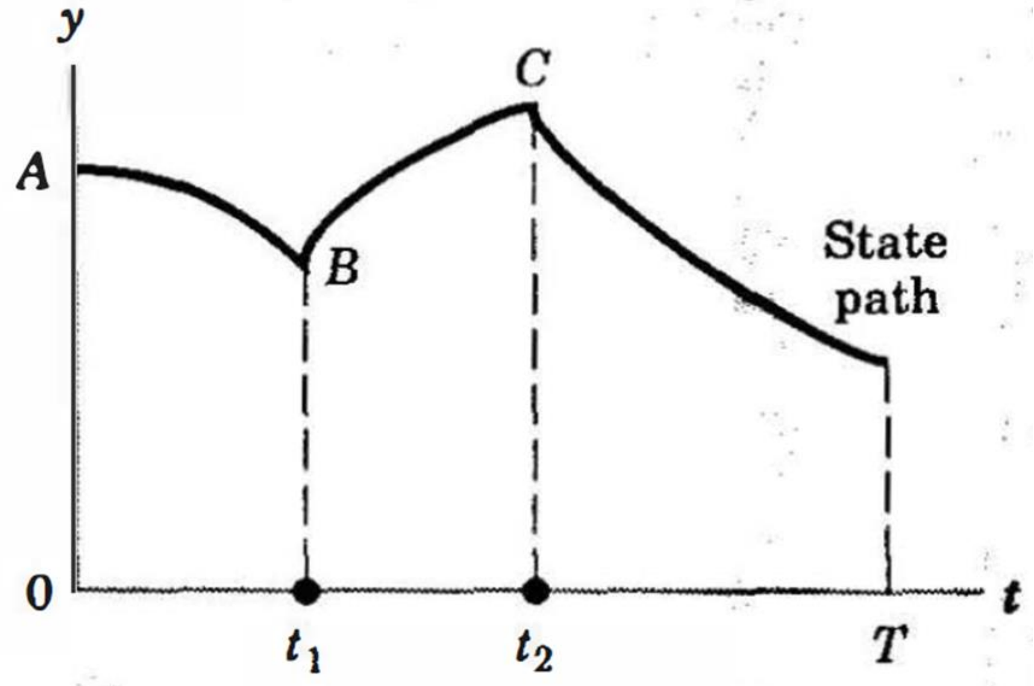

Figure 1b

The Simplest Problem of Optimal Control

- Another feature of importance is that optimal control theory is capable of directly handling a constraint on the control variable u, such as the restriction $u(t) \in \mathcal{U}$ for all $t \in [0, T]$, where $\mathcal{U}$ denotes some bounded control set.
- The control set can in fact be a closed, convex set. The fact that $\mathcal{U}$ can be a closed set means that corner solutions can be admitted.
- Finally, the simplest problem in optimal control theory, unlike in the calculus of variations, has a free terminal state (vertical terminal line) rather than a fixed terminal point.
- In the development of the fundamental first-order condition known as the maximum principle, we shall invoke the notion of an arbitrary Δu. Any arbitrary Δu must, however, imply an associated Δy.

The Simplest Problem

- Based on the preceding discussion, we may state the simplest problem of optimal control as:

Maximize ou minimize $\quad \mathrm{V} = \int_0^T F[t, y, u]dt$

(1) Subject to $\quad \dot{y} = f(t, y, u)$

$y(0) = A, \quad \mathrm{y}(T) = free \quad (A, T\ given)$

and $u(t) \in \mathcal{U}$ for all $t \in [0, T]$

- Here, as in the subsequent discussion, we shall deal exclusively with the maximization problem.
- In (1), the objective functional still takes the form of a definite integral, but the integrand function F no longer contains a y' argument. Instead, there is a new argument u.

The Simplest Problem

- The presence of the control variable u necessitates a linkage between u and y, to tell us how u will specifically affect the course taken by the state variable y.
- This information is provided by the equation $\dot{y} = f(t, y, u)$ [*equation of motion*], where the dotted symbol y, denoting the time derivative dy/dt, is an alternative notation to the y'.
- For some chosen policy at $t = 0$, say, $u_1(0)$, this equation will yield a specific value for y, say $y_1(0)$, which entails a specific direction the y variable is to move.
- What this equation does, therefore, is to provide the mechanism whereby our choice of the control u can be translated into a specific pattern of movement of the state variable y.

A Special Case

- As a special case, consider the problem where the choice of u is unconstrained, and where the equation of motion takes the form $\dot{y} = u$. Then, the optimal control problem becomes:

Maximize $\quad \mathrm{V} = \int_0^T F[t, y, u]dt$

(2) Subject to $\quad \dot{y} = u$

$y(0) = A, \qquad \mathrm{y}(T) = free \qquad (A, T\ given)$

- By substituting the equation of motion into the integrand function:

Maximize $\quad \mathrm{V} = \int_0^T F[t, y, \dot{y}]dt$

(2’) Subject to $\quad y(0) = A, \qquad \mathrm{y}(T) = free \qquad (A, T\ given)$

- The problem of the calculus of variations with a vertical terminal line.

The Costate Variable and the Hamiltonian Function

- The most important result in optimal control theory-a first-order necessary condition-is known as the maximum principle.
- The statement of the maximum principle involves the concepts of the Hamiltonian function and costate variable.
- Three types of variables were already presented in the problem statement (1): t (time), y (state), and u (control).
- It turns out that in the solution process, yet another type of variable will emerge. It is called the costate variable (or auxiliary variable), to be denoted by λ.

The Costate Variable and the Hamiltonian Function

- A costate variable is akin to a Lagrange multiplier and, as such, it is in the nature of a valuation variable, measuring the shadow price of an associated state variable.
- The vehicle through which the costate variable gains entry into the optimal control problem is the Hamiltonian function , or simply the Hamiltonian, which figures very prominently in the solution process.
- Denoted by H, the Hamiltonian is defined as

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

- Since H consists of the integrand function F plus the product of the costate variable and the function f, it itself should naturally be a function with four arguments: $t, y, u, and\ \lambda$.

The Maximum Principle

- The maximum principle involves two first-order differential equations in the state variable y and the costate variable λ.
- Besides, there is a requirement that the Hamiltonian be maximized with respect to the control variable u at every point of time.
- For the problem in (1), and with the Hamiltonian defined in (3), the maximum principle conditions are

$$\underset{u}{\text{Max}}\, H(t, y, u, \lambda) \qquad \text{for all } t \in [0, T]$$

(4) Subject to

$$\dot{y} = \frac{\partial H}{\partial \lambda} \qquad \text{[equation of motion for } y\text{]}$$

$$\dot{\lambda} = -\frac{\partial H}{\partial y} \qquad \text{[equation of motion for } \lambda\text{]}$$

$$\lambda(T) = 0 \qquad \text{[transversality condition]}$$

The Maximum Principle

- The symbol $\underset{u}{\mathrm{Max}}\, H(t, y, u, \lambda)$ means that the Hamiltonian is to be maximized with respect to u alone as the choice variable.
- An equivalent way of expressing this condition is

(5) $H(t, y, u^*, \lambda) \geq H(t, y, u, \lambda)$ for all $t \in [0, T]$

- where u^* is the optimal control, and u is any other control value.
- It is this requirement of maximizing H with respect to u that gives rise to the name " the maximum principle.“
- In Fig. 2, we have drawn three curves, each indicating a possible plot of the Hamiltonian H against the control variable u at a specific point of time, for specific values of y and λ.

Figure 2

- The control region is assumed to be the closed interval $[a, c]$.
- For curve 1, which is differentiable with respect to u, the maximum of H occurs at $u = b$, an interior point of the control region $\mathcal{U}$. In this case, $\partial H/\partial u = 0$ could serve to identify the optimal control at that point of time.
- But if curve 2 is the relevant curve, then the control in $\mathcal{U}$ maximizes H is $u = c$, a boundary point of $\mathcal{U}$. Thus the condition $\partial H/\partial u = 0$ does not apply, even though the curve is differentiable.
- And in the case of curve 3, with the
- Hamiltonian linear in u, the maximum of H occurs at $u = a$, another boundary point, and $\partial H/\partial u = 0$ is again inapplicable.

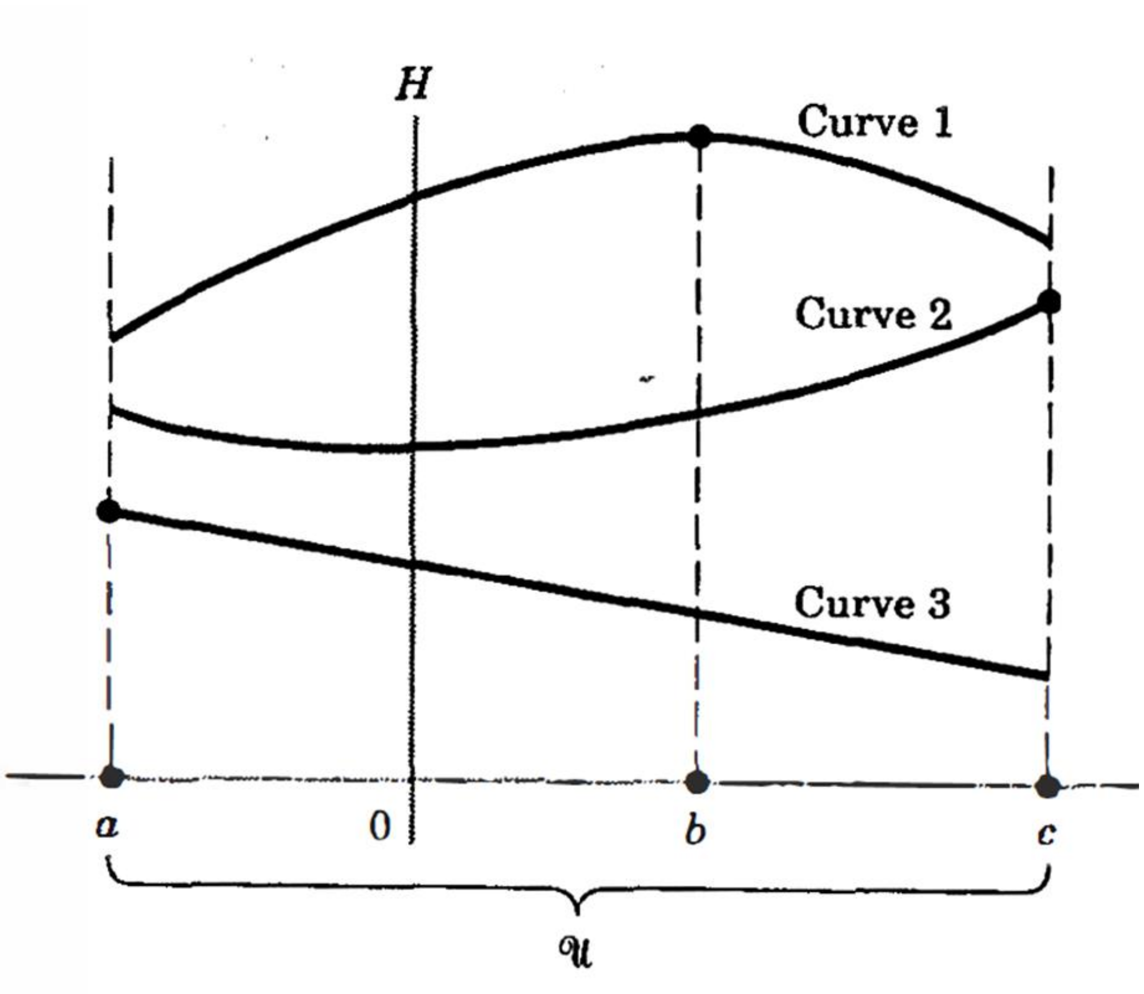

The Maximum Principle

- The condition $\dot{y} = \frac{\partial H}{\partial \lambda}$ is nothing but a restatement of the equation of motion for the state variable originally specified in (1).
- To express $\dot{y}$ as a partial derivative of H with respect to the costate variable λ is for showing the symmetry between this equation of motion and that for the costate variable.
- Note, however, that in the latter equation of motion, $\dot{\lambda}$ is the negative of the partial derivative of H with respect to y $\left(\dot{\lambda} = -\frac{\partial H}{\partial y}\right)$.
- Together, the two equations of motion are referred to collectively as the Hamiltonian system, or the canonical system (meaning the "standard“ system o f differential equations) for the given problem.

Example

- Find the curve with the shortest distance from a given point P to a given straight line L. Assume that L is a vertical line.
- The F function can be written as $(1+u^2)^{1/2}$, with u in place of $\dot{y}$.
- To convert the distance-minimization problem to one of maximization, we must attach a minus sign to the integrand.

$$Max\ \text{V} = \int_0^T -(1+u^2)^{1/2}dt$$

(4) Subject to $\quad \dot{y} = u \quad$ [equation of motion for y]

$$y(0) = A; \quad y(T) = free \quad [A \text{ and } T \text{ given}]$$

- Note that the control variable is not constrained, so the optimal control will be an interior solution.

Example

- Step I - We begin by writing the Hamiltonian function

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

$$H = -(1 + u^2)^{1/2} + \lambda(t) u$$

- Observing that H is differentiable and nonlinear, we can apply the first-order condition $\partial H / \partial u = 0$ to get

$$\frac{\partial H}{\partial u} = -\frac{1}{2}(1 + u^2)^{-1/2} 2u + \lambda = 0$$

$$\Rightarrow \ (1 + u^2)^{-1/2} u = \lambda \qquad \Rightarrow \ u^2 = \lambda^2 \ (1 + u^2)$$

$$\Rightarrow \ u^2 = \lambda^2 + \lambda^2 u^2 \qquad \Rightarrow \ u^2 (1 - \lambda^2) = \lambda^2$$

$$\Rightarrow \ u = \frac{\lambda}{(1-\lambda^2)^{1/2}} \qquad \Rightarrow \ u = \lambda \ (1 - \lambda^2)^{-1/2}$$

Example

- Further differentiation of $\partial H/\partial u$ using the product rule yields

$\frac{\partial^2 H}{\partial u^2} = \frac{1}{4}(1+u^2)^{-3/2}4u^2 - (1+u^2)^{-1/2} = (1+u^2)^{-3/2}u^2 - (1+u^2)^{-1/2} = u^2(1+u^2)^{-3/2} - (1+u^2)(1+u^2)^{-3/2} = -(1+u^2)^{-3/2} \leq 0$

- Thus $u = \lambda(1-\lambda^2)^{-1/2}$ maximize H. Since this equation expresses u in terms of λ, we must look for a solution for λ.
- Step II - To do that, we resort to the equation o f motion for the costate variable $\dot{\lambda} = -\frac{\partial H}{\partial y}$. But since $H = -(1+u^2)^{1/2} + \lambda(t)$ is independent of y, we have:

$\dot{\lambda} = -\frac{\partial H}{\partial y} = 0 \qquad\qquad \lambda(t) = constant$

Example

- Conveniently, the transversality condition $\lambda(T) = 0$ is sufficient for definitizing the constant. For if λ is a constant, then its value at $t = T$ also its value for all t. Thus,

  $\lambda^*(t) = 0$ for all $t \in [0, T]$

- Since $u = \lambda\,(1 - \lambda^2)^{-1/2}$

  $u^*(t) = 0$ for all $t \in [0, T]$

- Step III - From the equation of motion $\dot{y} = u$, we are now able to write

  $\dot{y} = 0 \quad \Rightarrow y(t) = constant$

Moreover, the initial condition $y(0) = A$ enables us to definitize this constant and write

Figure 3 - Vertical terminal line

$y^*(t) = A \qquad$ for all $t \in [0, T]$.

- This $y^*(t)$ path, illustrated in Fig. 3, is a horizontal straight line.
- Alternatively, it may be viewed as a path orthogonal to the vertical terminal line.