

Calculus of Variations

Summer Term 2017

Lecture 16

Universität des Saarlandes

16 June 2017

Purpose of Lesson:

- To introduce Pontryagin's Maximum Principle (PMP)
- To discuss several PMP examples

Pontryagin's Maximum Principle

Modern optimal control theory often starts from the PMP. It is a simple, concise condition for an optimal control.

General control problem

Minimize functional

$$J[\mathbf{x}, \mathbf{u}] = \int_{t_0}^{t_1} F_0(t, \mathbf{x}, \mathbf{u}) dt$$

subject to constraints $\dot{\mathbf{x}} = \mathbf{F}(t, \mathbf{x}, \mathbf{u})$, or more fully,

$$\dot{x}_i = F_i(t, \mathbf{x}, \mathbf{u})$$

- notice no dependence on $\dot{\mathbf{x}}$ in F_0
 - this differs from many CoV problems
- no dependence on $\dot{\mathbf{x}}$ in F_i because we rearrange the equations so that derivatives are on the LHS.

Pontryagin's Maximum Principle (PMP)

Let $\mathbf{u}(t)$ be an admissible control vector that transfers (t_0, \mathbf{x}_0) to a target $(t_1, \mathbf{x}(t_1))$. Let $\mathbf{x}(t)$ be the trajectory corresponding to $\mathbf{u}(t)$.

In order that $\mathbf{u}(t)$ be optimal, it is necessary that there exists $\mathbf{p}(t) = (p_1(t), p_2(t), \dots, p_n(t))$ and a constant scalar p_0 such that

- \mathbf{p} and \mathbf{x} are the solution to the canonical system

$$\dot{\mathbf{x}} = \frac{\partial \mathbb{H}}{\partial \mathbf{p}} \quad \text{and} \quad \dot{\mathbf{p}} = -\frac{\partial \mathbb{H}}{\partial \mathbf{x}}$$

- where the Hamiltonian is $\mathbb{H} = \sum_{i=0}^n p_i F_i$ with $p_0 = -1$
- $\mathbb{H}(t, \mathbf{x}, \mathbf{u}, \mathbf{p}) \geq \mathbb{H}(t, \mathbf{x}, \hat{\mathbf{u}}, \mathbf{p})$ for all alternate controls $\hat{\mathbf{u}}$
- all boundary conditions are satisfied

PMP proof sketch-1

Consider the general problem: minimize functional

$$J[\mathbf{x}, \mathbf{u}] = \int_{t_0}^{t_1} F_0(t, \mathbf{x}, \mathbf{u}) dt$$

subject to constraints

$$\dot{x}_i = F_i(t, \mathbf{x}, \mathbf{u}).$$

We can incorporate the constraints into the functional using the Lagrange multipliers λ_i , e.g.

$$\begin{aligned} \widehat{J} &= \int_{t_0}^{t_1} L(t, \mathbf{x}, \dot{\mathbf{x}}, \mathbf{u}) dt \\ &= \int_{t_0}^{t_1} F_0(t, \mathbf{x}, \mathbf{u}) dt + \sum_{i=1}^n \lambda_i(t) [\dot{x}_i - F_i(t, \mathbf{x}, \mathbf{u})] dt \end{aligned}$$

PMP proof sketch-2

Given such a function we get (by definition)

$$p_i = \frac{\partial L}{\partial \dot{x}_i} = \lambda_i.$$

So we can identify the Lagrange multipliers λ_i with the **generalized momentum** terms p_i

- 1 the p_i are known in economics literature as **marginal valuation** of x_i or the **shadow prices**
- 2 shows how much a unit increment in x at time t contributes to the optimal objective functional \hat{J}
- 3 the p_i are known in control as **co-state variables** (sometimes written as z_i).

PMP proof sketch-3

By definition the Hamiltonian is

$$\begin{aligned}
 \mathbb{H}(t, \mathbf{x}, \mathbf{p}, \mathbf{u}) &= \sum_{i=1}^n p_i \dot{x}_i - L(t, \mathbf{x}, \dot{\mathbf{x}}, \mathbf{p}, \mathbf{u}) \\
 &= \sum_{i=1}^n p_i \dot{x}_i - F_0(t, \mathbf{x}, \mathbf{u}) - \sum_{i=1}^n \lambda_i(t) [\dot{x}_i - F_i(t, \mathbf{x}, \mathbf{u})] \\
 &= -F_0(t, \mathbf{x}, \mathbf{u}) + \sum_{i=1}^n p_i F_i(t, \mathbf{x}, \mathbf{u})
 \end{aligned}$$

because $\lambda_i = p_i$, so the \dot{x}_i terms cancel. The final result is just the Hamiltonian as defined in the PMP.

PMP proof sketch-4

From previous slide the Hamiltonian can be written

$$\mathbb{H}(t, \mathbf{x}, \mathbf{p}, \mathbf{u}) = -F_0(t, \mathbf{x}, \mathbf{u}) + \sum_{i=1}^n p_i F_i(t, \mathbf{x}, \mathbf{u})$$

which is the Hamiltonian defined in the PMP. Then the canonical E-L equations (Hamilton's equations) are

$$\frac{\partial \mathbb{H}}{\partial p_i} = \frac{dx_i}{dt} \quad \text{and} \quad \frac{\partial \mathbb{H}}{\partial x_i} = -\frac{dp_i}{dt}.$$

Note that the equations $\frac{\partial \mathbb{H}}{\partial p_i} = \frac{dx_i}{dt}$ just revert to

$$F_i(t, \mathbf{x}, \mathbf{u}) = \dot{x}_i$$

which are just the system equations.

PMP proof sketch-5

Finally, note that Hamilton's equations above only relate x_i and p_i . What about equations for u_i ?

Take the conjugate variable to be z_i , and we get (by definition) that

$$z_i = \frac{\partial L}{\partial \dot{u}_i} = 0$$

and the second of Hamilton's equations is therefore

$$\frac{\partial \mathbb{H}}{\partial u_i} = -\frac{dz_i}{dt} = 0$$

which suggests a stationary point of \mathbb{H} WRT u_i .

In fact we look for a maximum (and note this may happen on the bounds of u_i).

PMP Example: plant growth

Example 13.1 (Plant growth-1)

Plant growth problem:

- market gardener wants to plants to grow to a fixed height 2 within a fixed window of time $[0, 1]$
- can supplement natural growth with lights (at night)
- growth rate dictates

$$\dot{x} = 1 + u$$

- cost of lights

$$J[u] = \int_0^1 \frac{1}{2} u^2 dt$$

PMP Example: plant growth

Example 13.1 (Plant growth-2)

Minimize

$$J[u] = \int_0^1 \frac{1}{2} u^2 dt$$

subject to $x(0) = 0$ and $x(1) = 2$ and

$$\dot{x} = F_1(t, x, u) = 1 + u.$$

Hamiltonian is

$$\begin{aligned} \mathbb{H} &= -F_0(t, x, u) + pF_1(t, x, u) \\ &= -\frac{1}{2}u^2 + p(1 + u). \end{aligned}$$

PMP Example: plant growth

Example 13.1 (Plant growth-3)

Hamiltonian is

$$\mathbb{H} = -\frac{1}{2}u^2 + p(1 + u).$$

Canonical equations

$$\begin{aligned} \frac{\partial \mathbb{H}}{\partial p} &= \frac{dx}{dt} & \text{and} & & \frac{\partial \mathbb{H}}{\partial x} &= -\frac{dp}{dt} \\ \downarrow & & & & \downarrow & \\ 1 + u &= \dot{x} & & & 0 &= -\dot{p} \end{aligned}$$

LHS \Rightarrow system DE

RHS $\Rightarrow \dot{p} = 0$ means that $p = c_1$ where c_1 is a constant.

PMP Example: plant growth

Example 13.1 (Plant growth-4)

Maximum principle requires \mathbb{H} be a maximum, for which

$$\frac{\partial \mathbb{H}}{\partial u} = -u + p = 0.$$

So $u = p$, and $\dot{x} = 1 + u$ so

$$\dot{x} = (1 + c_1) t + c_2.$$

The solution which satisfies $x(0) = 0$ and $x(1) = 2$ is

$$x = 2t.$$

So $u = c_1 = 1$, and the optimal cost is $\frac{1}{2}$.

PMP and natural boundary conditions

Typically we fix t_0 and $\mathbf{x}(t_0)$, but often the right-hand boundary condition is not fixed, so we need natural boundary conditions.

Here, they differ from traditional CoV problems in two respects:

- The terminal cost ϕ
- The function F_0 is not explicitly dependent on $\dot{\mathbf{x}}$.

The resulting natural boundary conditions are

$$\sum_i \left(\frac{\partial \phi}{\partial x_i} + p_i \right) \delta x_i \Big|_{t=t_1} + \left(\frac{\partial \phi}{\partial t} - \mathbb{H} \right) \delta t \Big|_{t=t_1} = 0$$

for all allowed δx_i and δt .

PMP and natural boundary conditions

The resulting natural boundary condition is

$$\sum_i \left(\frac{\partial \phi}{\partial x_i} + p_i \right) \delta x_i \Big|_{t=t_1} + \left(\frac{\partial \phi}{\partial t} - \mathbb{H} \right) \delta t \Big|_{t=t_1} = 0.$$

Special cases

- when t_1 is fixed and $\mathbf{x}(t_1)$ is completely free we get

$$\left(\frac{\partial \phi}{\partial x_i} + p_i \right) \delta x_i \Big|_{t=t_1} = 0, \quad \forall i$$

- when $\mathbf{x}(t_1)$ is fixed, $\delta x_i = 0$, and we get

$$\left(\frac{\partial \phi}{\partial t} - \mathbb{H} \right) \delta t \Big|_{t=t_1} = 0.$$

Example: stimulated plant growth

Example 13.2 (Stimulated plant growth-1)

Plant growth problem:

- market gardener wants to plants to grow as much as possible within a fixed window of time $[0, 1]$
- supplement natural growth with lights as before
- growth rate dictates $\dot{x} = 1 + u$
- cost of lights

$$J[u] = \int_0^1 \frac{1}{2} u^2(t) dt$$

- value of crop is proportional to the height

$$\phi(t_1, \mathbf{x}(t_1)) = x(t_1).$$

Plant growth problem statement

Example 13.2 (Stimulated plant growth-2)

Write as a minimization problem

$$J[x, u] = -x(t_1) + \int_0^1 \frac{1}{2} u^2 dt$$

subject to $x(0) = 0$, and

$$\dot{x} = 1 + u.$$

- the terminal cost doesn't affect the shape of the solution
- but we need a natural end-point condition for t_1 .

Plant growth: natural BC

Example 13.2 (Stimulated plant growth-3)

The problem is solved as before, but we write the natural boundary condition at $x = t_1$ as

$$\left(\frac{\partial \phi}{\partial x_i} + p_i \right) \Big|_{t=t_1} = 0, \quad \forall i$$

which reduces to

$$-1 + p|_{t=t_1} = 0.$$

Given p is constant, this sets $p(t) = 1$, and hence the control $u = 1$ (as before).

Autonomous problems

Autonomous problems have no explicit dependence on t .

- time invariance symmetry
- hence \mathbb{H} is constant along the optimal trajectory
- if the end-time is free (and the terminal cost is zero) then the transversality conditions ensure $\mathbb{H} = 0$ along the optimal trajectory.