The Pontryagin Maximum Principle

M403 Lecture Notes by Philip D. Loewen

Problem Statement. Given an initial point $(\tau, \xi) \in \mathbb{R} \times \mathbb{R}^n$, choose $T \in (\tau, +\infty)$, $u \in PWC([\tau, T]; \mathbb{R}^m)$, and $x \in PWS([\tau, T]; \mathbb{R}^n)$ to

$$\begin{aligned}
\text{minimize } & \ell_0(T, x(T)) + \int_\tau^T L_0(t, x(t), u(t))\, dt \\
\text{subject to } & \ell_j(T, x(T)) + \int_\tau^T L_j(t, x(t), u(t))\, dt = \gamma_j, \;\; j = 1, \ldots, M, \\
& \dot{x}(t) = f(t, x(t), u(t)) \text{ a.e. } t \in [\tau, T], \\
& u(t) \in U(t) \text{ a.e. } t \in [\tau, T], \\
& x(\tau) = \xi, \\
& (t, x(t)) \in G \;\forall t \in [\tau, T].
\end{aligned}$$

Hypotheses. The set $G \subseteq \mathbb{R} \times \mathbb{R}^n$ is open. The endpoint functions $\ell_j \colon G \to \mathbb{R}$ are continuously differentiable. The integrands $L_j \colon G \times \mathbb{R}^m \to \mathbb{R}$, and the right-hand side of the dynamic equation $f \colon G \times \mathbb{R}^m \to \mathbb{R}$ are continuous in all three variables (t, x, u), and so are their partial derivatives with respect to t and x. The control set $U(t)$ moves "measurably" with respect to t.

The Pre-Hamiltonian. For each scalar-valued constraint in the problem statement, we have one scalar Lagrange multiplier. The multipliers form a vector $\lambda = (\lambda_1, \ldots, \lambda_M)$ whose dimension equals the number of constraints. Together with a "normality indicator" $\lambda_0 \geq 0$ described below, the multipliers appear in the definition of the problem's pre-Hamiltonian:

$$H(t, x, p, u) = p^T f(t, x, u) - \lambda_0 L_0(t, x, u) - \sum_{j=1}^M \lambda_j L_j(t, x, u).$$

Theorem (Pontryagin Maximum Principle). *Suppose a control-state pair $(\widehat{u}, \widehat{x})$ gives the minimum in the problem described above; assume that $\widehat{u}$ is piecewise continuous. Then there exist a vector of Lagrange multipliers $(\lambda_0, \lambda) \in \mathbb{R} \times \mathbb{R}^M$ with $\lambda_0 \geq 0$ and a piecewise smooth function $p \colon [\tau, T] \to \mathbb{R}^n$ such that the function $\widehat{h}(t) \stackrel{\text{def}}{=} H(t, \widehat{x}(t), p(t), \widehat{u}(t))$ is piecewise smooth, and one has*

(a) *Adjoint equations:*
$$\dot{\widehat{h}}(t) = H_t(t, \widehat{x}(t), p(t), \widehat{u}(t)) \text{ a.e.,}$$
$$-\dot{p}(t) = H_x(t, \widehat{x}(t), p(t), \widehat{u}(t)) \text{ a.e.,}$$

(b) *State equations:*
$$\dot{\widehat{x}}(t) = H_p(t, \widehat{x}(t), p(t), \widehat{u}(t)) \text{ a.e.,}$$

(c) *Maximum condition:*
$$\widehat{u}(t) \in \arg\max_{v \in U(t)} H(t, \widehat{x}(t), p(t), v) \text{ a.e.,}$$

(d) *Transversality:*
$$\left(\widehat{h}(T),\, -p(T)\right) = \lambda_0 \nabla \ell_0(T, \widehat{x}(T)) + \sum_{j=1}^M \lambda_j \nabla \ell_j(T, \widehat{x}(T)),$$

(e) *Nontriviality: the function p and the vector (λ_0, λ) are not both zero.*

Extremality and Nontriviality. Notice that by setting $\lambda_0 = 0$, $\lambda = 0$, and $p(t) \equiv 0$, it is possible to satisfy conditions (a)–(d) for any control-state pair $(\widehat{u}, \widehat{x})$ whatsoever. Thus condition (e) is essential if we are to derive useful information from the theorem. A pair satisfying all five conditions is called **extremal.** Conceptually, a complete list of all extremal pairs in a given problem is guaranteed to contain the pair that actually gives the minimum (assuming such a pair exists); actually deciding which pair that is may require additional work.

Scaling; Normality. If $(\widehat{u}, \widehat{x})$ is an extremal control-state pair with associated Lagrange multipliers (λ_0, λ) and adjoint function p, then for any constant $\gamma > 0$, conditions (a)–(e) also hold for the Lagrange multipliers $(\gamma\lambda_0, \gamma\lambda)$ and adjoint function γp. A standard choice of γ is λ_0^{-1}, which is valid whenever $\lambda_0 > 0$. Thus the phrase "$\lambda_0 \geq 0$" in the theorem statement can be replaced by the simple alternative "$\lambda_0 = 0$ or $\lambda_0 = 1$". An extremal $(\widehat{u}, \widehat{x})$ is called **abnormal** if is possible to satisfy (a)–(e) with $\lambda_0 = 0$, and **normal** otherwise. As the terminology suggests, abnormal extremals are strange and rare: when $\lambda_0 = 0$, the necessary conditions do not seem to involve the ingredients ℓ_0 and L_0 of the function we originally set out to minimize! A standard approach to problem-solving is first to set $\lambda_0 = 0$ and try to deduce that $p \equiv 0$ and $\lambda = 0$: if these consequences follow, there can be no abnormal extremals, and one can confidently proceed with the solution assuming $\lambda_0 = 1$.

Equation Counting. Taking the original constraints of the problem together with the extremality conditions (a)–(e), we have a system of equations for the three unknown functions $\widehat{x}$ in $\mathbb{R}^n$, $\widehat{u}$ in $\mathbb{R}^m$, and p in $\mathbb{R}^n$, the unknown final time T, and the M-vector λ of Lagrange multipliers. To determine p and $\widehat{x}$ we have differential equations of the appropriate sizes in (a) and (b), involving a parameter $\widehat{u}$ that is in principle available from (c). The initial condition on $\widehat{x}$ provides n constants of integration. The M constraints in the problem statement can be used to find the M Lagrange multipliers λ, and the transversality condition (d) contains $n+1$ additional equations that provide n constants of integration for the function p and one last condition useful for specifying the terminal time T.

Special Cases

The stated theorem covers an enormous range of applications. Some of the possibilities are sketched below. In all cases, the adjoint equation (a), the state equation (b), and the maximum condition (c) must hold. The difference usually centres on the form of the transversality condition (d) and the strength of a suitably modified nontriviality assertion (e). The transversality conditions are the control-theoretic counterparts of the "natural boundary conditions" in the calculus of variations: when the problem statement leaves some component of the primal vector $(T, x(T))$ unconstrained, the transversality condition says something useful about the corresponding component of the dual vector $(h(T), -p(T))$. But for those components that are fixed in the problem statement, the transversality condition gives no hints about the corresponding dual components. Thus, for example, in a fixed time problem one gets no information about $h(T)$, while in a fixed-endpoint problem one gets no information about $p(T)$.

By contrast, in a problem with completely free time (and no final-time dependence in the objective function) one learns that $h(T) = 0$, and in a problem with a completely free right endpoint $x(T)$ (and no final-state dependence in ℓ_0), the missing boundary condition on x is replaced by the condition $p(T) = 0$ on the costate.

1. Autonomous problems. If the dynamic function f and the integrands L_j $(j = 0, \ldots, M)$ have no explicit time-dependence, then $H_t \equiv 0$ and the first adjoint equation implies that the function $\widehat{h}(t)$ is essentially constant. In problems with fixed final time (see item 2 below), this is redundant information which can nonetheless be useful in problem-solving: compare the Second Weierstrass-Erdmann condition in the calculus of variations. In problems with free final time, the information gleaned from $\widehat{h}$ is essential for determining when to stop—whether or not one is in the autonomous case.

2. Fixed-time problems. Suppose the problem is stated exactly as above, except that the final time $T = \widehat{T} > \tau$ is prescribed in advance. Here it does not make sense to allow explicit T-dependence in the endpoint functions ℓ_j, so the M side constraints must have the form

$$\ell_j(x(\widehat{T})) + \int_\tau^{\widehat{T}} L_j(t, x(t), u(t))\, dt = \gamma_j, \;\; j = 1, \ldots, M.$$

To put this fixed-time problem into the framework discussed above, where the final time is subject to choice, simply imagine a free-time problem involving the additional constraint $0 = \ell_{M+1}(T, x(T))$, where $\ell_{M+1}(t, x) = t - \widehat{T}$. The PMP as stated provides for an arc p, a normality indicator $\lambda_0 \geq 0$, and a vector $(\lambda_1, \ldots, \lambda_M, \lambda_{M+1})$, **not all zero,** such that conditions (a)–(c) hold, along with the transversality condition (d):

$$\left(\widehat{h}(\widehat{T}),\, -p(\widehat{T})\right) = \lambda_0 \left(0,\, \nabla\ell_0(\widehat{x}(\widehat{T}))\right) + \sum_{j=1}^{M} \lambda_j \left(0,\, \nabla\ell_j(\widehat{x}(\widehat{T}))\right) + \lambda_{M+1}\,(1,\, 0)\,.$$

This vector equation encodes two conditions. The first component gives the scalar equation

$$h(\widehat{T}) = \lambda_{M+1},$$

while the second gives the vector equation

$$(\mathrm{d}')\;\; -p(\widehat{T}) = \lambda_0 \nabla\ell_0(\widehat{x}(\widehat{T})) + \sum_{j=1}^{M} \lambda_j \nabla\ell_j(\widehat{x}(\widehat{T})).$$

Thus we get no information about the value of $h(\widehat{T})$ in a fixed-time problem. Moreover, if all of the multipliers except for λ_{M+1}—namely, λ_0, $(\lambda_1, \ldots, \lambda_M)$, and p—are zero, then it follows that $\widehat{h}(t) \equiv 0$ and consequently the first component equation above implies $\lambda_{M+1} = 0$ also, which violates the nontriviality condition. Thus we can state the nontriviality condition in terms of multipliers excluding λ_{M+1}.

Summary. In the fixed-time problem stated here, every minimizing (control,state)-pair $(\widehat{u}, x)$ must satisfy conditions (a)–(c) of the PMP for some adjoint function $p(t)$ and some multipliers $\lambda_0 \geq 0$ and $(\lambda_1, \ldots, \lambda_M)$. The appropriate transversality condition is (d′); moreover, at least one element of the list p, λ_0, $(\lambda_1, \ldots, \lambda_M)$ must be nonzero.

3. Fixed Endpoint Problems. Consider a problem formulated just like the one stated above, except that it includes an additional constraint that $x(T) = \eta$ for some fixed vector $\eta \in \mathbb{R}^n$. To embed this problem in the one just stated, we simply introduce n additional scalar constraints—one for each component of the endpoint condition:

$$0 = \ell_{M+j}(T, x(T)), \text{ where } \ell_{M+j}(t, x) = x_j - \eta_j, \quad j = 1, \ldots, n.$$

Then any minimizing (control,state)-pair $(\widehat{u}, x)$ must have associated a costate arc p and a vector of multipliers $(\lambda_0, \lambda_1, \ldots, \lambda_M, \lambda_{M+1}, \ldots, \lambda_{M+n})$, **not both zero,** such that (a)–(c) all hold, and the transversality condition (d) states

$$\left(\widehat{h}(T),\ -p(T)\right) = \lambda_0 \nabla \ell_0(T, \widehat{x}(T)) + \sum_{j=1}^{M} \lambda_j \nabla \ell_j(T, \widehat{x}(T)) + \sum_{j=1}^{n} \lambda_{M+j} \left(0,\ \mathbf{e}_{M+j}\right).$$

Since the Lagrange multipliers $\lambda_{M+1}, \ldots, \lambda_{M+n}$ are unknown, this equation gives us no useful information about the vector $p(T)$, and so the only piece of information potentially worth keeping is from the first component:

$$(\mathrm{d}') \;\; \widehat{h}(T) = \left[\lambda_0 \frac{\partial \ell_0}{\partial t} + \sum_{j=1}^{M} \frac{\partial \ell_j}{\partial t} \right]_{(T, \widehat{x}(T))}.$$

(And, as noted above, even this must be discarded in the case where the problem also has a preassigned final time.) As for the nontriviality condition, suppose that the arc p, the normality indicator λ_0, and the initial Lagrange multipliers $(\lambda_1, \ldots, \lambda_M)$ are all zero. In this case, the second component-equation in the full transversality condition written above implies that

$$0 = \sum_{j=1}^{n} \lambda_{M+j} \left(0,\ \mathbf{e}_{M+j}\right) = (\lambda_{M+1}, \ldots, \lambda_{M+n}),$$

and hence that *all* the multipliers in our application of PMP vanish. This contradicts the nontriviality condition (e), so we can make the stronger nontriviality assertion in the statement below.

Summary. If the problem stated above is modified by adding the terminal constraint $x(T) = \eta$, then there exist an adjoint function $p(t)$, a normality indicator $\lambda_0 \geq 0$, and a Lagrange multiplier vector $(\lambda_1, \ldots, \lambda_M)$, **not all zero,** such that conditions (a)–(c) of PMP hold, while (d) is replaced by (d′).

4. Minimum-time problems. Consider the problem of minimum time to hit a fixed target—say the origin. If the starting point is fixed, this has the form just discussed with cost functions $\ell_0(t, x) = t$, $L_0(t, x, u) = 0$ and no additional constraints (so $M = 0$). In this case the pre-Hamiltonian is simply $H(t, x, p, u) = p^T f(t, x, u)$, and the PMP asserts that any optimal pair $(\widehat{u}, x)$ must come equipped with some costate arc p and normality indicator λ_0 such that

(a) $\dot{\widehat{h}}(t) = p(t)^T f_t(t, \widehat{x}(t), \widehat{u}(t))$; $-\dot{p}(t) = \left[p(t)^T f_x(t, \widehat{x}(t), \widehat{u}(t))\right]^T = \widehat{f}_x(t)^T p(t)$,

(b) $\dot{\widehat{x}}(t) = f(t, \widehat{x}(t), \widehat{u}(t))$,

(c) $\widehat{u}(t) \in \arg\max_{v \in U(t)} p(t)^T f(t, \widehat{x}(t), v)$,

(d) $\widehat{h}(T) = \lambda_0$,

(e) $\lambda_0 + |p(T)| > 0$.

Now if $\lambda_0 = 0$, then (e) says $p(T) \neq 0$. On the other hand, if $\lambda_0 > 0$, then (d) says $0 < h(T) = p(T)^T f(T, \widehat{x}(T), \widehat{u}(T))$, which again implies $p(T) \neq 0$. Thus it is legitimate to replace condition (e) with the stronger-looking condition $p(T) \neq 0$. With this adjustment, conditions (a)–(e) compare favourably with the necessary conditions we proved for the linear time-optimality problem. The differential equations for p and x and the maximum condition are all identical, as is the modified nontriviality condition (e). The only new information here pertains to the behaviour of the pre-Hamiltonian function $\widehat{h}$ along the optimal trajectory—something that is interesting, but that also follows easily from the other conditions in the linear case.

5. The Lagrange Multiplier Rule. Given a convex open set $\Omega \subseteq \mathbb{R}^n$ and a collection of smooth functions $g_0, g_1, \ldots, g_M \colon \Omega \to \mathbb{R}$,

$$\text{minimize } g_0(\eta) \text{ subject to } g_j(\eta) = 0,\ j = 1, \ldots, M,\ \eta \in \Omega.$$

Suppose the minimum occurs at some point $\widehat{\eta}$. Without loss of generality, assume $0 \in \Omega$. Set $(\tau, \xi) = (0, 0)$, $L_0 = L_j = 0$ for $j = 1, \ldots, M$, $U = \mathbb{R}^m$, and $f(t, x, u) = u$. Define $\ell_j(t, x) = g_j(x)$, and introduce one last constraint to fix the final time at $T = 1$: $\ell_{M+1}(t, x) = t - 1$. Then the control-state pair $(\widehat{u}(t), \widehat{x}(t)) = (\widehat{\eta}, \widehat{\eta}t)$ solves the dynamic problem above, so the PMP applies. We have $H(t, x, p, u) = p^T u$: this is independent of both t and x, so the adjoint equations imply that the functions $\widehat{h}$ and p in the theorem statement must both be constant. The maximum condition (c) can hold only if the indicated maximum has the value zero, and if this arises because $p = 0$. It follows then that $h = 0$ as well. Now the PMP provides a constant λ_0 equal to either 0 or 1, and Lagrange multipliers $\lambda_1, \ldots, \lambda_M, \lambda_{M+1}$ such that we have the transversality condition

$$(0, 0) = (\widehat{h}, -p) = (0, \lambda_0 \nabla g_0(\widehat{\eta})) + \sum_{j=1}^{M} (0, \lambda_j \nabla g_j(\widehat{\eta})) + \lambda_{M+1}(1, 0).$$

This gives $0 = \lambda_{M+1}$ and

$$0 = \lambda_0 \nabla g_0(\widehat{\eta}) + \sum_{j=1}^{M} \lambda_j \nabla g_j(\widehat{\eta}).$$

The upshot is that if $\widehat{\eta}$ solves the stated finite-dimensional problem, then there exist a constant λ_0 equal to either 0 or 1 and a vector $\lambda = (\lambda_1, \ldots, \lambda_M)$, not both zero, such that $(*)$ holds. This is the Lagrange multiplier rule for finite-dimensional minimization.