Calculus of Variations, Lecture 7, 3/12/2013

Our treatment of optimal control has thus far focused on the "value function" + the assoc Hamilton-Jacobi-Bellman eqn.

Today we discuss an alternative approach, based on Pontryagin's Maximum Principle (PMP). It is of interest for several reasons:

1) PMP sometimes leads easily to explicit solutions.

2) In high dimensions it may be much more efficient than solving the HJ eqn

3) It provides a very concrete interpretation for the characteristics of the HJ eqn

Two references:

- Leslie Hocking, Optimal Control: An Introduction to the Theory with Applications, Oxford Univ Press 1991 (a really good book, with lots of examples)

- A. Dixit, Optimization in Economic Theory, Oxford Univ Press 1990 (a pretty good book, with a focus on economics examples naturally.)

(The book by Gamkrelidze on my syllabus is more advanced; its focus is how to "relax" an optimal control problem when the optimal behavior involves making $\alpha(s)$ oscillate in time.)

The main idea: consider the "min arrival time problem"

$$u(x) = \min \{ \text{time to arrive at } \partial D, \text{ starting from } x \in D + \text{travelling with speed} \le 1 \}$$

for which the HJ eqn is $|\nabla u| = 1$ in D, $u = 0$ at ∂D; solution formula is

$$u(x) = \min_{z \in \partial D} dist(x, z).$$

There are good algs for solving the HJ eqn; basically one must build the level sets one by one starting at $u = 0$ (see eg J. Sethian's book Level Set Methods + Fast Marching Methods, Cambridge Univ Press 1999)

But using the PMP in this setting amounts to using the solution formula, ie searching for $z \in \partial D$ st $dist(x, z)$ is smallest. It's a 1D problem instead of a 2D calculation.

Let's derive the PMP formalism, by a min/max argument, focusing on the finite-horizon problem

$$u(x,t) = \max_{\alpha} \int_t^T h(y(s), \alpha(s))\, ds + g(y(T))$$

with state eqn

$$y_s = f(y(s), \alpha(s)) \text{ for } t < s < T, \quad y(t) = x.$$

We write it as a max/min, using a new vector-valued unknown $\pi(s)$ (taking values in $\mathbb{R}^n$ if $\vec{y}(t) \in \mathbb{R}^n$)

$$u(x,t) = \max_{\substack{\alpha \\ y(t)=x}} \min_{\pi(s)} \left\{ \int_t^T \pi(s) \cdot \left[f(y(s), \alpha(s)) - \frac{dy}{ds} \right] ds + \int_t^T h(y(s), \alpha(s))\, ds + g(y(T)) \right\}$$

where now $y(t) = x$ is the only constraint in the maximization (the min over $\pi(s)$ enforces the state eqn). Assuming max min = min max, we get (after integn by parts)

$$u(x,t) = \min_{\pi(s)} \max_{\substack{y(t)=x \\ \alpha \in A}} \left\{ \int_t^T \left[y \cdot \frac{d\pi}{ds} + f(y,\alpha)\cdot\pi + h(y,\alpha) \right] ds + g(y(T)) + x\cdot\pi(t) - y(T)\cdot\pi(T) \right\}$$

Now max over α: if $\pi(s)$ + $y(s)$ are fixed, best $\alpha(s)$ maximizes $f\cdot\pi + h$, So

$$u(x,t) = \min_{\pi(s)} \max_{y(t)=x} \left\{ \int_t^T \left[y \cdot \frac{d\pi}{ds} + H(\pi,y) \right] ds + g(y(T)) + x\cdot\pi(t) - y(T)\cdot\pi(T) \right\}$$

with

(*) $$H(p,x) = \max_{\alpha \in A} \left[p\cdot f(x,\alpha) + h(x,\alpha) \right]$$

(Note: we saw earlier that $u_t + H(\nabla u, x) = 0$ using the <u>same</u> function H!)

Note this corollary to (*):

(1) $$\frac{\partial H}{\partial p} = f(x,\alpha_0)$$ where α_0 is optimal for the defn (*)

In fact: let $\alpha_0 = \alpha_0(p,x)$ be the best α as a fn of $p + x$; Then

$$H(p,x) = p \cdot f(x, \alpha_0(p,x)) + h(x, \alpha_0(p,x))$$

$$\Rightarrow \quad \frac{\partial H}{\partial p} = f(x, \alpha_0(p,x)) + \left(p \cdot \frac{\partial f}{\partial \alpha} + \frac{\partial h}{\partial \alpha}\right)\Bigg|_{\alpha=\alpha_0} \frac{\partial \alpha_0}{\partial p}$$

$$= f(x, \alpha_0(p,x))$$

since the 2nd term vanishes (because α_0 is optimal).

Also note: if we choose the control at time s to be $\alpha_0(\pi(s), y(s))$ then

$$(2) \qquad \frac{dy}{ds} = \nabla_p H(\pi(s), y(s))$$

using (1) and the state eqn $dy/ds = f$.

We return now to the min max expression for $u(x,t)$, and evaluate the inner max. Since the only constraint is $y(t) = x$, $y(s)$ is is arbitrary for $s > t$, so $y \cdot \frac{d\pi}{ds} + H(\pi, y)$ must be maximized in y for each s. Therefore

(3) $$\frac{d\pi}{ds} = -\frac{\partial H}{\partial y} .$$

Arguing similarly in optimizing over $y(T)$ we get

(4) $$\pi(T) = \nabla g(y(T)).$$

Collecting (1) - (4) we get Pontryagin's Max Prin for this class of problems:

- at each s, the control $\alpha(s)$ should be the value α_0 such that $\pi \cdot f(y, \alpha_0) + h(y, \alpha_0)$ is max (with $\pi = \pi(s)$ and $y = y(s)$).

- evolution of $y(s)$ + $\pi(s)$ is given by

$$(**) \begin{cases} dy/ds = \nabla_\pi H(\pi(s), y(s)) \\ d\pi/ds = -\nabla_y H(\pi(s), y(s)) \end{cases}$$

for $t < s < T$

- initial cond'n for $y(s)$ is known: $y(t) = x$. But endpt cond'n for $\pi(s)$ is not known; rather, it is determined implicitly by the final-time condition $\pi(T) = \nabla g(y(T))$.

Since H is being evaluated at $(\pi(s), y(s))$ and $u_t + H(\nabla u, x) = 0$, it's natural to guess that $\nabla u(y(s),s) = \pi(s)$ along the solution. This is true; the proof rests on the fact that the ODE's for $y(s)$ + $\pi(s)$ are in fact the characteristic eqns for our HJ eqn. (Exercise: prove that $\nabla u(y(s), s) = \pi(s)$.)

Can the PMP be used to solve problems? Yes, but it's not usually trivial. The problem is that we don't know an initial condn for π (ie the value of $\pi(t)$) except by searching for a choice that gives $\pi(T) = \nabla g(y(T))$. Typically this must be done by performing some kind of search or "shooting" scheme.

Here's a simple example:

$$\min_{\alpha(s)} \int_0^T \tfrac{1}{2}\alpha^2(s)\, ds + \tfrac{1}{2} y^2(T)$$

where $y_s = y + \alpha$, $y(0) = x$ (with $x, y, \alpha \in \mathbb{R}$). Evidently: we prefer $y(T)$ to be small, which requires nonzero α since $\alpha = 0 \Rightarrow$ exp growth; but we also prefer $\int \alpha^2$ to be small. (This pbm involves minimization not maximization; the PMP is as above, but where we maximized over α before, we must now minimize). Evidently

$$g(y) = \frac{1}{2} y^2 \qquad h(x,\alpha) = \frac{1}{2}\alpha^2$$

$$f(y,\alpha) = y + \alpha$$

so

$$H(p,x) = \min_{\alpha} \; p(x+\alpha) + \frac{1}{2}\alpha^2 = px - \frac{1}{2}p^2$$

and the PMP gives

$$\frac{dy}{ds} = y(s) - \pi(s) \qquad y(0) = x$$

$$\frac{d\pi}{ds} = -\pi(s) \qquad \pi(T) = y(T).$$

It's clear we should be able to find the soln by taking $\pi(0) = \pi_0$ as an unknown, solving for $\pi(s)$, $y(s)$, then asking which π_0 gives $\pi(T) = y(T)$. (Exercise: do this.)

Suggested exercises:

(1) If you have studied the method of characteristics (as it applies to a 1st order pde like our HJB eqn for the value function), then use it to

show that when $y(s)$ + $\pi(s)$ solve the ODE's of Pontryagin's max principle, $\nabla u(y(s), s) = \pi(s)$ for all $t < s < T$.

(2) Complete the example begun on pp 7.7-7.8. The value function for that optimal control problem can be determined by the method of Lectures 5-6 Exercise 4 (ie the example is another "linear quadratic regulator" problem). Is this fact of any use in solving the PMP?

(3) Our other favorite class of examples was the "minimum arrival time" problem

$$\min_{\alpha(s)\in A} \{ 1^{st} \text{ time } y(s) \text{ reaches } \Gamma \}$$

where Γ is a specified "target set" and the state equation is

$$dy/ds = f(y(s), \alpha(s)) \;, \quad y(0) = x \,.$$

What should the Pontryagin Max Principle say in this case? (Check the case $x \in D$, $\Gamma = \partial D$, $f(y(s), \alpha(s)) = \alpha(s)$, $A = \{ |\alpha| \le 1 \}$ to be sure your answer is reasonable.)

(4) Consider the "optimal consumption" example we

solved in §5-6, with $\rho=0$ for simplicity:

$$\max_{a(s)} \int_t^T a^\gamma(s)\,ds$$

where the state eqn is

$$\frac{dy}{ds} = ry - a \qquad t<s<T, \quad y(t)=x.$$

(Here $0<\gamma<1$ and $r>0$ are fixed constants.)

a) What ode's + end conditions does the PMP give in this case? How is the optimal control $a(s)$ related to $\pi(s)$?

b) Show that $a(s) = Ce^{rs/(1-\gamma)}$ for some constant C.

[To find a complete, explicit solution this way requires solving a nonlinear eqn for the value of $\pi(t)$ that gives $y(T)=0$. This is messy, so I'm not suggesting that you do it by hand.]

The books by Hocking + Dixit (both on reserve now) have lots more examples.