10

The Linear Quadratic Regulator

10.1 Problem formulation

This chapter concerns optimal control of dynamical systems. Most of this development concerns linear models with a particularly simple notion of optimality. However, to understand the most basic concepts in optimal control, and not become lost in complex notation, it is most convenient to consider first the general model given in nonlinear state space form

$$\dot{x} = f(x, u, t), \qquad x(t_0) = x_0 \tag{10.1}$$

The cost V of a particular control input u is defined by the following integral

$$V(u) = \int_{t_0}^{t_1} \ell(x, u, \tau)\, d\tau + m(x(t_1))$$

where

(a) t_1 is the *final time* of the control problem.

(b) ℓ is a scalar-valued function of x, u, and t

(c) m is a function of x. It is called the *terminal penalty function*.

We assume that x_0, t_0, and t_1 are known, fixed values, and $x(t_1)$ is free. Our goal is to choose the control $u_{[t_0,t_1]}$ to minimize V.

A case which is typical in applications is where m and ℓ are quadratic functions of their arguments. For an LTV model, the system description and cost are then given by

$$\begin{aligned}\dot{x} &= A(t)x + B(t)u, \qquad x(t_0) = x_0 \\ V(u) &= \int_{t_0}^{t_1} (x^T Q(t) x + u^T R(t) u)\, dt + x^T(t_1) M x(t_1)\end{aligned} \tag{10.2}$$

where M, Q and R are positive semidefinite matrix-valued functions of time. These matrices can be chosen by the designer to obtain desirable closed loop response. The minimization of the quadratic cost V for a linear system is known as the *linear quadratic regulator* (LQR) problem.

We will study the LQR problem in detail, but first we develop some general results for the nonlinear state space model (10.1).

10.2 Hamilton-Jacobi-Bellman equations

The *value function* $V^\circ = V^\circ(x_0, t_0)$ is defined to be the minimum value of V over all controls. This is a function of the two variables x and t which can be written explicitly as

$$V^\circ(x,t) = \min_{u_{[t,t_1]}} \Big[\int_t^{t_1} \ell(x(\tau), u(\tau), \tau)\, d\tau + m(x(t_1)) \Big]$$

Under very general conditions, the value function satisfies a partial differential equation known as the *Hamilton-Jacobi-Bellman (HJB) equation*. To derive this result, let x and t be an arbitrary initial time and initial state, and let t_m be an intermediate time, $t < t_m < t_1$. Assuming that $x(\tau)$, $t \le \tau \le t_1$, is a solution to the state equations with $x(t) = x$, we must have

$$\begin{aligned} V^\circ(x,t) &= \min_{u_{[t,t_1]}} \Big[\int_t^{t_m} \ell(x(\tau), u(\tau), \tau)\, d\tau + \int_{t_m}^{t_1} \ell(x(\tau), u(\tau), \tau)\, d\tau + m(x(t_1)) \Big] \\ &= \min_{u_{[t,t_m]}} \Big[\int_t^{t_m} \ell(x(\tau), u(\tau), \tau)\, d\tau + \underbrace{\min_{u_{[t,t_1]}} \Big(\int_{t_m}^{t_1} \ell(x(\tau), u(\tau), \tau)\, d\tau + m(x(t_1)) \Big)}_{V^\circ(x(t_m), t_m)} \Big] \end{aligned}$$

This gives the functional equation

$$V^\circ(x,t) = \min_{u_{[t,t_m]}} \Big[\int_t^{t_m} \ell(x(\tau), u(\tau), \tau)\, d\tau + V^\circ(x(t_m), t_m) \Big] \tag{10.3}$$

As a consequence, the optimal control over the whole interval has the property illustrated in Figure 10.1: If the optimal trajectory passes through the state x_m at time $x(t_m)$ using the control u°, then the control $u^\circ_{[t_m, t_1]}$ must be optimal for the system starting at x_m at time t_m. If a better u^* existed on $[t_m, t_1]$, we would have chosen it. This concept is called the *principle of optimality*.

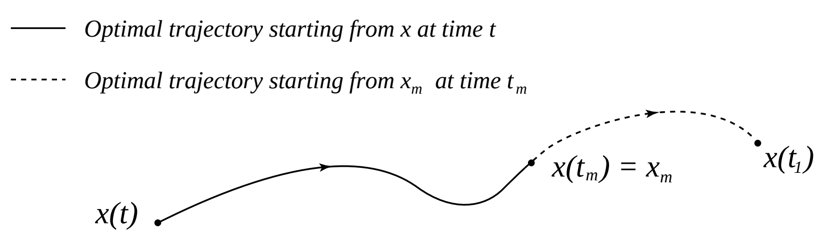

Fig. 10.1. If a better control existed on $[t_m, t_1]$, we would have chosen it.

By letting t_m approach t, we can derive a partial differential equation for the value function V°. Let Δt denote a small positive number, and define

$$\begin{aligned} t_m &= t + \Delta t \\ x_m &= x(t_m) = x(t + \Delta t) = x(t) + \Delta x \end{aligned}$$

Assuming that the value function is sufficiently smooth, we may perform a Taylor series expansion on V° using the optimality equation (10.3) to obtain

$$V^\circ(x,t) = \min_{u_{[t,t_m]}} \left\{ \ell(x(t),u(t),t)\Delta t + V^\circ(x,t) + \frac{\partial V^\circ}{\partial x}(x(t),t)\Delta x + \frac{\partial V^\circ}{\partial t}(x(t),t)\Delta t \right\}$$

Dividing through by Δt and recalling that $x(t)=x$ then gives

$$0 = \min_{u_{[t,t_m]}} \left\{ \ell(x,u(t),t)\frac{\Delta t}{\Delta t} + \frac{\partial V^\circ}{\partial x}(x,t)\frac{\Delta x}{\Delta t} + \frac{\partial V^\circ}{\partial t}(x,t)\frac{\Delta t}{\Delta t} \right\}$$

Letting $\Delta t \to 0$, the ratio $\Delta x/\Delta t$ can be replaced by a derivative to give

$$0 = \min_u \left[\ell(x,u,t) + \frac{\partial V^\circ}{\partial x}(x,t)\dot{x}(t) \right] + \frac{\partial V^\circ}{\partial t}$$

where

$$\frac{\partial V^\circ}{\partial x} = \left[\frac{\partial V^\circ}{\partial x_1} \cdots \frac{\partial V^\circ}{\partial x_n} \right] = (\nabla_x V^\circ)^T$$

Thus, we have obtained the following partial differential equation which the value function must satisfy if it is smooth. The resulting equation (10.4) is the *Hamilton Jacobi Bellman (HJB) equation*,

$$-\frac{\partial V^\circ}{\partial t}(x,t) = \min_u \left[\ell(x,u,t) + \frac{\partial V^\circ}{\partial x}(x,t) f(x,u,t) \right] \tag{10.4}$$

The terminal penalty term gives a boundary condition for the HJB equation

$$V^\circ(x(t_1),t_1) = m(x(t_1)).$$

The term in brackets in (10.4) is called the *Hamiltonian*, and is denoted H:

$$H(x,p,u,t) := \ell(x,u,t) + p^T f(x,u,t), \tag{10.5}$$

where $p = \nabla_x V^\circ$. We thus arrive at the following

Theorem 10.2.1. *If the value function V° has continuous partial derivatives, then it satisfies the following partial differential equation*

$$-\frac{\partial V^\circ}{\partial t}(x,t) = \min_u H(x, \nabla_x V^\circ(x,t), u, t),$$

and the optimal control $u^\circ(t)$ and corresponding state trajectory $x^\circ(t)$ must satisfy

$$\min_u H(x^\circ(t), \nabla_x V^\circ(x^\circ(t),t), u, t) = H(x^\circ(t), \nabla_x V^\circ(x^\circ(t),t), u^\circ(t), t). \tag{10.6}$$

□

An important consequence of Theorem 10.2.1 is that the optimal control can be written in state feedback form $u^\circ(t) = \bar{u}(x^\circ(t),t)$, where the function $\bar{u}$ is defined through the minimization in (10.6).

Example 10.2.1. Consider the simple integrator model, with the polynomial cost criterion

$$\dot{x} = u \qquad V(u) = \int_0^T [u^2 + x^4]\, dt$$

Here we have $f(x,u,t) = u$, $\ell(x,u,t) = u^2 + x^4$, and $m(x,t) \equiv 0$. The Hamiltonian is thus

$$H(x,p,u,t) = pu + u^2 + x^4,$$

and the HJB equation becomes

$$-\frac{\partial V^\circ}{\partial t} = \min_u \{\frac{\partial V^\circ}{\partial x} u + u^2 + x^4\}.$$

Minimizing with respect to u gives

$$u^\circ = -\tfrac{1}{2}\frac{\partial V^\circ}{\partial x}(x,t),$$

which is a form of state feedback. The closed loop system has the appealing form

$$\dot{x}^\circ(t) = -\tfrac{1}{2}\frac{\partial V^\circ}{\partial x}(x^\circ(t),t)$$

This equation shows that the control forces the state to move in the direction in which the "cost to go" V° decreases.

Substituting the formula for u° back into the HJB equation gives the PDE

$$-\frac{\partial V^\circ}{\partial t}(x,t) = -\frac{1}{4}\left(\frac{\partial V^\circ}{\partial x}(x,t)\right)^2 + x^4,$$

with the boundary condition $V^\circ(x,T) = 0$. This is as far as we can go, since we do not have available methods to solve a PDE of this form. If a solution is required, it may be found numerically. However, a simpler set of equations is obtained in the limit as $T \to \infty$. This simpler problem is treated in Exercise 1 below. □

We now leave the general nonlinear model and concentrate on linear systems with quadratic cost. We will return to the more general problem in Chapter 11.

10.3 A solution to the LQR problem

For the remainder of this chapter we consider the LQR problem whose system description and cost are given in (10.2). For this control problem, x_0, t_0 and t_1 are given, and $x(t_1)$ is free. To ensure that this problem has a solution we assume that R is strictly positive definite ($R > 0$).

To begin, we now compute the optimal control u° by solving the HJB partial differential equation. The Hamiltonian for this problem is given by

$$H(x,p,u,t) = \ell + p^T f = x^T Q x + u^T R u + p^T(Ax + Bu)$$

To minimize H with respect to u we compute the derivative

11

Introduction to the Minimum Principle

We now return to the general nonlinear state space model, with general cost criterion of the form

$$\begin{aligned} \dot{x}(t) &= f(x(t), u(t), t), \quad x(t_0) = x_0 \in \mathbb{R}^n \\ V(u) &= \int_{t_0}^{t_1} \ell(x(t), u(t), t)\, dt + m(x(t_1)) \end{aligned} \tag{11.1}$$

We have already shown under certain conditions that if the input u minimizes $V(u)$, then a partial differential equation known as the HJB equation must be satisfied. One consequence of this result is that the optimal control can be written in state feedback form through the derivative of the value function V°. The biggest drawback to this result is that one must solve a PDE which in general can be very complex. Even for the relatively simple LQR problem, the solution of the HJB equations required significant ingenuity.

The minimum principle is again a method for obtaining necessary conditions under which a control u is optimal. This result is based upon the solution to an ordinary differential equation in $2n$ dimensions. Because it is an ordinary differential equation, in many instances it may be solved even though the HJB equation is intractable. Unfortunately, this simplicity comes with a price. The solution u° to this ordinary differential equation is in open-loop form, rather than state space form. The two view points, both the Minimum Principle and the HJB equations, each have value in nonlinear optimization, and neither approach can tell the whole story.

11.1 Minimum Principle and the HJB equations

There is no space in a book this size or a course at this level to give a complete proof of the Minimum Principle. We can however give some heuristic arguments to make the result seem plausible, and to gain some insight. We initially consider the optimization problem (11.1) where x_0, t_0, t_1 are fixed, and $x(t_1)$ is free. Our first approach is through the HJB equation

$$\begin{aligned}
-\frac{\partial}{\partial t}V^\circ(x,t) &= \min_u(\ell(x,u,t) + \frac{\partial V^\circ}{\partial x}f(x,u,t)) \\
u^\circ(t) &= \arg\min_u(\ell(x,u,t) + \frac{\partial V^\circ}{\partial x}f(x,u,t))
\end{aligned}$$

For x, t fixed, let $\bar{u}(x,t)$ denote the value of u which attains the minimum above, so that $u^\circ = \bar{u}(x^\circ(t),t)$. In the derivation below we assume that $\bar{u}$ is a smooth function of x. This assumption is false in many models, which indicates that another derivation of the Minimum Principle which does not rely on the HJB equation is required to create a general theory. With $\bar{u}$ so defined, the HJB equation becomes

$$-\frac{\partial}{\partial t}V^\circ(x,t) = \ell(x,\bar{u}(x,t),t) + \frac{\partial V^\circ}{\partial x}(x,t)f(x,\bar{u}(x,t),t).$$

Taking partial derivatives of both sides with respect to x gives the term $\frac{\partial V^\circ}{\partial x}$ on both sides of the resulting equation:

$$\begin{aligned}
-\frac{\partial^2}{\partial x \partial t}V^\circ(x,t) = {} & \frac{\partial \ell}{\partial x}(x,\bar{u}(x,t),t) \\
& + \frac{\partial^2 V^\circ}{\partial^2 x}(x,t)f(x,\bar{u}(x,t),t) \\
& + \frac{\partial V^\circ}{\partial x}(x,t)\frac{\partial f}{\partial x}(x,\bar{u}(x,t),t) \\
& + \underbrace{\frac{\partial}{\partial u}\left(\ell + \frac{\partial V^\circ}{\partial x}f\right)(x,\bar{u}(x,t),t)}_{\textit{derivative vanishes at } \bar{u}}\frac{\partial \bar{u}(x,t)}{\partial x}
\end{aligned}$$

As indicated above, since $\bar{u}$ is the unconstrained minimum of $H = \ell + \frac{\partial V^\circ}{\partial x}f$, the partial derivative with respect to u must vanish.

This PDE holds for any state x and time t - they are treated as independent variables. Consider now the optimal trajectory $x^\circ(t)$ with optimal input $u^\circ(t) = \bar{u}(x^\circ(t),t)$. By a simple substitution we obtain

$$\begin{aligned}
0 = \frac{\partial^2}{\partial x \partial t}V^\circ(x^\circ(t),t) \quad &+ \quad \frac{\partial \ell}{\partial x}(x^\circ(t),u^\circ(t),t) \\
&+ \quad \frac{\partial^2 V^\circ}{\partial x^2}(x^\circ(t),t)f(x^\circ(t),u^\circ(t),t) \\
&+ \quad \frac{\partial V^\circ}{\partial x}(x^\circ(t),t)\frac{\partial f}{\partial x}(x^\circ(t),u^\circ(t),t). \qquad (11.2)
\end{aligned}$$

We can now convert this PDE into an ODE. The main trick is to define

$$p(t) := \frac{\partial V^\circ}{\partial x}(x^\circ(t),t)$$

This function is interpreted as the *sensitivity of the cost with respect to current state*, and takes values in $\mathbb{R}^n$. The derivative of p with respect to t can be computed as follows:

$$\begin{aligned}\frac{d}{dt}p(t) &= \frac{\partial^2 V^\circ}{\partial x^2}(x^\circ(t),t)\dot{x}^\circ(t) + \frac{\partial^2 V^\circ}{\partial x \partial t}(x^\circ(t),t) \\ &= \frac{\partial^2 V^\circ}{\partial x^2}(x^\circ(t),t) f(x^\circ(t),u^\circ(t),t) + \frac{\partial^2 V^\circ}{\partial x \partial t}(x^\circ(t),t).\end{aligned}$$

The two mixed partial terms in this equation are also included in (11.2). Combining these two equations, we thereby eliminate these terms to obtain

$$\begin{aligned}0 = \frac{d}{dt}p(t) &+ \frac{\partial \ell}{\partial x}(x^\circ(t),u^\circ(t),t) \\ &+ \underbrace{\frac{\partial V^\circ}{\partial x}(x^\circ(t),t)}_{p^T(t)}\frac{\partial f}{\partial x}(x^\circ(t),u^\circ(t),t).\end{aligned} \tag{11.3}$$

From the form of the Hamiltonian

$$H(x,p,u,t) = p^T f(x,u,t) + \ell(x,u,t),$$

the differential equation (11.3) may be written

$$\dot{p}(t) = -\nabla_x H(x^\circ(t),p(t),u^\circ(t),t)$$

with the boundary condition

$$p(t_1) = \nabla_x V^\circ(x^\circ(t_1),t_1) = \nabla_x m(x^\circ(t_1),t_1).$$

This is not a proof, since for example we do not know if the value function V° will be smooth. However, it does make the following result seem plausible. For a proof see [8].

Theorem 11.1.1 (Minimum Principle)**.** *Suppose that t_1 is fixed, $x(t_1)$ is free, and suppose that u° is a solution to the optimal control problem (11.1). Then*

(a) There exists a costate vector $p(t)$ such that for $t_0 \le t \le t_1$,

$$u^\circ(t) = \arg\min_u H(x^\circ(t),p(t),u,t)$$

(b) The pair (p,x°) satisfy the 2-point boundary value problem:

$$\begin{aligned}\dot{x}^\circ(t) &= \nabla_p H(x^\circ(t),p(t),u^\circ(t),t) \qquad \Big(= f(x^\circ(t),u^\circ(t),t)\Big) \\ \dot{p}(t) &= -\nabla_x H(x^\circ(t),p(t),u^\circ(t),t)\end{aligned} \tag{11.4}$$

with the two boundary conditions

$$x(t_0) = x_0; \qquad p(t_1) = \nabla_x m(x(t_1),t_1).$$

□

11.2 Minimum Principle and Lagrange multipliers⋆

Optimal control involves a functional minimization which is similar in form to ordinary optimization in $\mathbb{R}^m$ as described in a second year calculus course. For *unconstrained* optimization problems in $\mathbb{R}^m$, the main idea is to look at the derivative of the function V to be minimized, and find points in $\mathbb{R}^m$ at which the derivative is zero, so that $\nabla V(x) = \vartheta$. Such an x is called a *stationary point* of the optimization problem. By examining all stationary points of the function to be minimized one can frequently find among these the optimal solution. The *calculus of variations* is the infinite dimensional generalization of ordinary optimization in $\mathbb{R}^m$. Conceptually, it is no more complex than its finite dimensional counterpart.

In more advanced calculus courses, solutions to *constrained* optimization problems in $\mathbb{R}^m$ are addressed using *Lagrange multipliers*. Suppose for example that one desires to minimize the function $V(x)$ subject to the constraint $g(x) = \vartheta$, $x \in \mathbb{R}^m$, where $g\colon \mathbb{R}^m \to \mathbb{R}^d$. Consider the new cost function $\widehat{V}(x,p) = V(x) + p^T g(x)$, where $p \in \mathbb{R}^d$. The vector p is known as a Lagrange multiplier. The point of extending the state space in this way is that we can now solve an unconstrained optimization problem, and very often this gives a solution to the constrained problem. Indeed, if x^0, p^0 is a stationary point for $\widehat{V}$ then we must have

$$\begin{aligned} \vartheta &= \nabla_x \widehat{V}\,(x^0,p^0) = \nabla V(x^0) + \nabla g(x^0)p^0 && (11.5)\\ \vartheta &= \nabla_p \widehat{V}\,(x^0,p^0) = g(x^0) && (11.6) \end{aligned}$$

Equation (11.5) implies that the gradient of V is a linear combination of the gradients of the components of g at x_0, which is a necessary condition for optimality under very general conditions. Figure 11.1 illustrates this with $m = 2$ and $d = 1$. Equation (11.6) is simply a restatement of the constraint $g = \vartheta$. This Lagrange multiplier approach can also be generalized to infinite dimensional problems of the form (11.1), and this then gives a direct derivation of the Minimum Principle which does not rely on the HJB equations.

To generalize the Lagrange multiplier approach we must first generalize the concept of a stationary point. Suppose that F is a functional on $D^r[t_0,t_1]$. That is, for any function $z \in D^r[t_0,t_1]$, $F(z)$ is a real number. For any $\eta \in D^r[t_0,t_1]$ we can define a directional derivative as follows:

$$D_\eta F\,(z) = \lim_{\epsilon \to 0} \frac{F(z+\epsilon\eta) - F(z)}{\epsilon},$$

whenever the limit exists. The function $z(t,\epsilon) = z(t) + \epsilon\eta(t)$ may be viewed as a perturbation of z, as shown in Figure 11.2

We call z_0 a *stationary point* of F if $D_\eta F\,(z_0) = 0$ for *any* $\eta \in D^r[t_0,t_1]$. If z_0 is a minimum of F, then a perturbation cannot decrease the cost, and hence we must have for any η,

$$F(z_0 + \epsilon\eta) \geq F(z_0), \qquad \epsilon \in \mathbb{R}.$$

From the definition of the derivative it then follows that an optimal solution must be a stationary point, just as in the finite dimensional case!

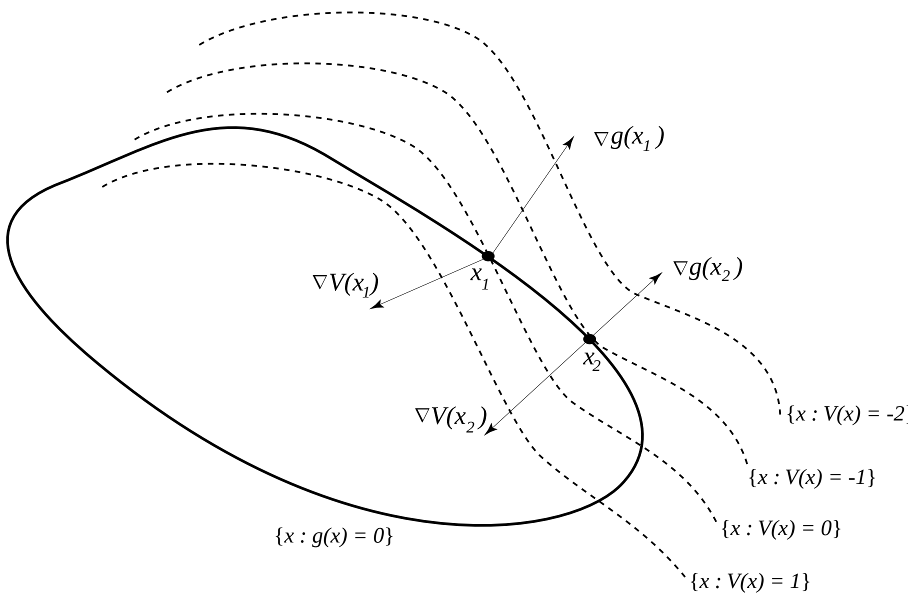

Fig. 11.1. An optimization problem on $\mathbb{R}^2$ with a single constraint $g = 0$. The point x_1 is not optimal: By moving to the right along the constraint set, the function $V(x)$ will decrease. The point x_2 is a minimum of V subject to the constraint $g(x) = 0$ since at this point we have $V(x_2) = -1$. V can get no lower - for example, $V(x)$ is never equal to -2 since the level set $\{x : V(x) = -2\}$ does not intersect the constraint set $\{x : g(x) = 0\}$. At x_2, the gradient of V and the gradient of g are parallel.

The problem at hand can be cast as a constrained functional minimization problem,

$$\begin{array}{ll} \text{Minimize} & V(x,u) = \int \ell \, d\tau + m \\ \text{Subject to} & \dot{x} - f = \vartheta, \qquad x \in D^n[t_0,t_1],\ u \in D^m[t_0,t_1]. \end{array}$$

To obtain necessary conditions for optimality for constrained problems of this form we can extend the Lagrange multiplier method. We outline this approach in five steps:

Step 1. Append the state equations to obtain the new cost functional

$$\widehat{V}(x,u) = \int_{t_0}^{t_1} \ell \, dt + m(x(t_1)) + \int_{t_0}^{t_1} p^T (f - \dot{x}) \, dt$$

The *Lagrange multiplier vector* p lies in $D^n[t_0,t_1]$. The purpose of this is to gain an expression for the cost in which x and u can be varied independently.

Step 2. Use integration by parts to eliminate the derivative of x in the definition of $\widehat{V}$:

$$\int_{t_0}^{t_1} p^T \dot{x} \, dt = \int_{t_0}^{t_1} p^T dx = \left(p^T x\right)\Big|_{t_0}^{t_1} - \int_{t_0}^{t_1} \dot{p}^T x \, dt$$

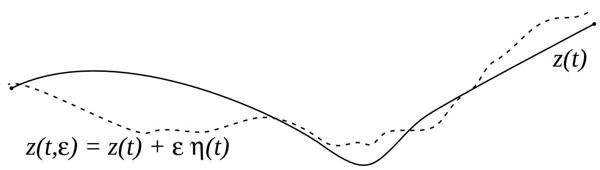

Fig. 11.2. A perturbation of the function $z \in D[t_0, t_1]$.

Step 3. Recall the definition of the Hamiltonian,

$$H(x, p, u, t) := \ell(x, u, t) + p^T f(x, u, t).$$

Combining this with the formulas given in the previous steps gives

$$\begin{aligned} \widehat{V}(x, u) &= \int_{t_0}^{t_1} H(x, p, u, t)\, dt + \int_{t_0}^{t_1} \dot{p}^T x\, dt \\ &\quad + p^T(t_0)x(t_0) - p^T(t_1)x(t_1) + m(x(t_1)) \end{aligned} \tag{11.7}$$

Step 4. Suppose that u° is an optimal control, and that x° is the corresponding optimal state trajectory. The Lagrange multiplier theorem asserts that the pair (x°, u°) is a stationary point of $\widehat{V}$ for some p°. Perform perturbations of the optimal control and state trajectories to form

$$u(t, \delta) = u^\circ(t) + \delta\psi(t) \qquad x(t, \epsilon) = x^\circ(t) + \epsilon\eta(t) \tag{11.8}$$

Since we are insisting that $x(t_0) = x_0$, we may assume that $\eta(t_0) = \vartheta$. Consider first variations in ϵ, with δ set equal to zero. Letting $\widehat{V}(\epsilon) = \widehat{V}(x^\circ + \epsilon\eta, u^\circ)$, we must have

$$\frac{d}{d\epsilon}\widehat{V}(\epsilon) = 0.$$

Using (11.7) to compute the derivative gives

$$\begin{aligned} 0 = \int_{t_0}^{t_1} \frac{\partial}{\partial x} H(x^\circ, p, u^\circ, t)\eta(t)\, dt + \int_{t_0}^{t_1} \dot{p}^T \eta\, dt \\ -p^T(t_1)\eta(t_1) + p^T(t_0)\eta(t_0) + \frac{\partial}{\partial x} m(x^\circ(t_1))\eta(t_1). \end{aligned} \tag{11.9}$$

Similarly, by considering perturbations in u° we obtain for any $\psi \in D^m[t_0, t_1]$,

$$\vartheta = \int_{t_0}^{t_1} \frac{\partial}{\partial u} H(x^\circ, p, u^\circ, t)\psi(t)\, dt \tag{11.10}$$

This simpler expression is obtained because only the first term in (11.7) depends upon u.

Step 5. We can choose $\eta(t)$ *freely* in (11.9). From this it follows that

$$\frac{\partial H}{\partial x} + \dot{p}^T = \vartheta^T \Rightarrow \dot{p} = -\nabla_x H$$

and since $\eta(t_0) = \vartheta$, and $\eta(t_1)$ is free,

$$-p^T(t_1) + \frac{\partial m}{\partial x}(t_1) = \vartheta^T \Rightarrow p(t_1) = \nabla_x m(x(t_1))$$

Similarly, by (11.10) we have

$$\frac{\partial H}{\partial u} = \vartheta^T \Rightarrow \nabla_u H = \vartheta$$

In fact, if u is to be a minimum of $\hat{V}$, then in fact it must minimize H pointwise. These final equations then give the Minimum Principle Theorem 11.1.1.

From this proof it is clear that many generalizations of the Minimum Principle are possible. Suppose for instance that the final state $x(t_1) = x_1$ is specified. Then the perturbation η will satisfy $\eta(t_1) = \vartheta$, and hence using (11.9), it is impossible to find a boundary condition for p. None is needed in this case, since to solve the $2n$-dimensional coupled state and costate equations, it is enough to know the initial and final conditions of x°.

11.3 The penalty approach*

A third heuristic approach to the Minimum Principle involves relaxing the hard constraint $\dot{x} - f = \vartheta$, but instead impose a large, yet "soft" constraint by defining the cost

$$\widehat{V}(x,u) = \int_{t_0}^{t_1} \ell(x(t),u(t),t)\,dt + \frac{k}{2}\int_{t_0}^{t_1} |\dot{x}(t) - f(x(t),u(t),t)|^2\,dt + m(x(t_1))$$

The constant k in this equation is assumed large, so that $\dot{x}(t) - f(x(t),u(t),t) \approx \vartheta$.

We assume that a pair (x_k, u_k) exists which minimizes $\widehat{V}_k$. Letting (x°, u°) denote a solution to the original optimization problem, we have by optimality,

$$\widehat{V}(x_k,u_k) \le \widehat{V}(x^\circ,u^\circ) = V^\circ.$$

Assuming ℓ and m are positive, this gives the following uniform bound

$$\int_{t_0}^{t_1} |\dot{x}(t) - f(x(t),u(t),t)|^2\,dt \le \frac{2}{k}V^\circ$$

Hence, for large k, the pair (x_k, u_k) will indeed approximately satisfy the differential equation $\dot{x} = f$.

If we perturb x_k to form $x_k + \epsilon\eta$ and define $\widehat{V}(\epsilon) = \widehat{V}(x_k + \epsilon\eta, u_k)$ then we must have $d/d\epsilon\,\widehat{V}(\epsilon) = 0$ when $\epsilon = 0$. Using the definition of $\hat{V}$ gives