

Nonlinear Control Systems

By Matthew R. James

Department of Systems Engineering, Research School of Information Sciences and Engineering, Australian National University, Canberra, ACT 0200, Australia.

1. Introduction

Control systems are prevalent in nature and in man-made systems. Natural regulation occurs in biological and chemical processes, and may serve to maintain the various constituents at their appropriate levels, for example. In the early days of the industrial revolution, governors were devised to regulate the speed of steam engines, while in modern times, computerised control systems have become commonplace in industrial plants, robot manipulators, aircraft and spacecraft, etc. Indeed, the highly maneuverable X-29 aircraft using forward swept wings is possible only because of its control systems, and moreover, control theory has been crucial in NASA's Apollo and Space Shuttle programmes. Control systems such as in these examples use in an essential way the idea of *feedback*, the central theme of this chapter.

Control theory is the branch of engineering/science concerned with the design and analysis of control systems. Linear control theory treats systems for which an underlying linear model is assumed, and is a relatively mature subject, complete with firm theoretical foundations and a wide range of powerful and applicable design methodologies; see e.g., Anderson & Moore (1990), Kailath (1980). In contrast, nonlinear control theory deals with systems for which linear models are not adequate, and is relatively immature, especially in relation to applications. In fact, linear systems techniques are frequently employed in spite of the presence of nonlinearities. Nonetheless, nonlinear control theory is exciting and vitally important, and is the subject of a huge and varied range of research worldwide.

The aim of this chapter is to convey to readers of *Complex Systems* something of the flavour of the subject, the techniques, the computational issues, and some of the applications. To place this chapter in perspective, in relation to the other chapters in this book, it is worthwhile citing Brockett's remark that control theory is a *prescriptive science*, whereas physics, biology, etc, are *descriptive sciences*, see Brockett (1976b). Computer science shares some of the prescriptive qualities of control theory, in the sense that some objective is prescribed, and means are sought to fulfill it. It is this *design* aspect that is most important here. Indeed, control systems are designed to *influence* the behaviour of the system being controlled in order to achieve a desired level of performance. Brockett categorised control theory briefly:

(i) To express models in input-output form, thereby identifying those variables which can be manipulated and those which can be observed.

(ii) To develop methods for regulating the response of systems by modifying the dynamical nature of the system—e.g. stabilisation.

(iii) To optimise the performance of the system relative to some performance index.

In addition, feedback design endeavours to compensate for disturbances and uncertainty. This chapter attempts to highlight the fundamental role played by feedback in control theory. Additional themes are stability, robustness, optimisation, information and com-

putational complexity. It is impossible in a few pages or even in a whole book to do justice to our aims, and the material presented certainly omits many important aspects.

This chapter does not discuss in any detail neural networks, fuzzy logic control, and chaos. These topics, while interesting and potentially important, are not sufficiently well developed in the context of control theory at present. In §2, some of the basic ideas of feedback are introduced and illustrated by example. The basic definition of state space models is given in §3, and this framework is used throughout this chapter. Differential geometry provides useful tools for nonlinear control, as will be seen in §4 on the important system theoretic concepts of controllability and observability. The possibility of feedback linearisation of nonlinear systems is discussed in §5, and the problem of stabilisation is the topic of §6. Optimisation-based methods are very important in control theory, and are discussed in the remaining sections (§§7–13). The basic ideas of deterministic and stochastic optimal control are reviewed, and issues such as robustness and stabilisation are discussed. Some results concerning output feedback problems are briefly touched upon, and computational methods and issues of computational complexity are covered. This chapter is written in a tutorial style, with emphasis on intuition. Some parts of this chapter are more technically demanding than other parts, and it is hoped that the references provided will be of use to interested readers.

2. The Power of Feedback

2.1. *What is Feedback?*

The conceptual framework of system theory depends in a large part on the distinction between different types of variables in a model, in particular, the inputs and outputs. It may happen that certain variables have natural interpretations as inputs since they are readily manipulated, while other measured variables serve as outputs. For instance, in a robot arm, the torque applied by the joint actuators is an input, whereas the measured joint angles are outputs. However, the distinction is not always clear and is a matter of choice in the modelling process. In econometrics, is the prime interest rate an input, or an output—i.e., is it an independent variable, or one which responds to other influences? For us, a *system* Σ is a causal map which assigns an output $y(\cdot)$ to each input $u(\cdot)$, drawn in Figure 1. Causality means that past values of the output do not depend on future values of the input.



FIGURE 1. Open loop system.

Feedback occurs when the *information* in the output is made available to the input, usually after some processing through another system. A general feedback system may involve additional inputs and outputs with specific meanings, Figure 2.

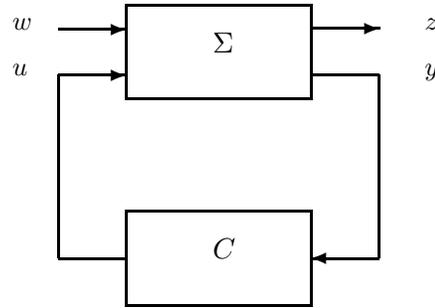


FIGURE 2. General closed loop feedback system.

In the configuration of Figure 2, the system Σ is often called the *plant* (the system being controlled), while the system C is called the *controller*. The inputs w may include reference signals to be followed, as well as unknown disturbances (such as modelling errors and noise), and the output z is a performance variable, say an error quantity with ideal value zero. The feedback system shown in Figure 2 is called a *closed loop* system, for obvious reasons, whereas, the system Σ alone as in Figure 1 is called an *open loop* system. Thus a feedback system is characterised by the information flow that occurs or is permitted among the various inputs and outputs. The modelling process consists of building a model of this form, and once this is done, the task of designing the controller to meet the given objectives can be tackled.

2.2. Feedback Design

To illustrate some of the benefits of feedback, we consider the following system:

$$m\ddot{\theta} - mg \sin \theta = \tau. \quad (2.1)$$

Equation (2.1) is a simple nonlinear model of an inverted pendulum, or a one-link robot manipulator. The angle θ is measured from the vertical ($\theta = 0$ corresponds to the vertical equilibrium), and τ is the torque applied by a motor to the revolute joint attaching the pendulum to a frame. The motor is the active component, and the pendulum is controlled by adjusting the motor torque appropriately. Thus the input is $u = \tau$. If the joint angle is measured, then the output is $y = \theta$. If the angular velocity is also measured, then $y = (\theta, \dot{\theta})$. The pendulum has length 1 m, mass m kg, and g is acceleration due to gravity.

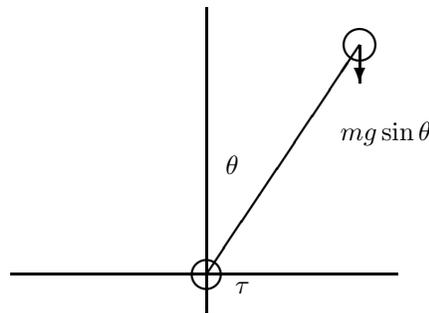


FIGURE 3. Inverted pendulum or 1-link robot arm.

The model neglects effects such as friction, motor dynamics, etc.

We begin by analysing the stability of the equilibrium $\theta = 0$ of the homogeneous system

$$m\ddot{\theta} - mg \sin \theta = 0 \quad (2.2)$$

corresponding to zero motor torque (no control action). The linearisation of (2.2) at $\theta = 0$ is

$$m\ddot{\theta} - mg\theta = 0, \quad (2.3)$$

since $\sin \theta \approx \theta$ for small θ . This equation has general solution $\alpha e^{\sqrt{g}t} + \beta e^{-\sqrt{g}t}$, and, because of the first term (exponential growth), $\theta = 0$ is not a stable equilibrium for (2.2). This means that if the pendulum is initially set-up in the vertical position, then a small disturbance can cause the pendulum to fall. We would like to design a control system to prevent this from happening, i.e. to restore the pendulum to its vertical position in case a disturbance causes it to fall.

One could design a stabilising feedback controller for the linearised system

$$m\ddot{\theta} - mg\theta = \tau \quad (2.4)$$

and use it to control the nonlinear system (2.1). This will result in a locally asymptotically stable system, which may be satisfactory for small deviations θ from 0.

Instead, we adopt an approach which is called the computed torque method in robotics, or feedback linearisation (see §5). This method has the potential to yield a globally stabilising controller. The nonlinearity in (2.1) is not *neglected* as in the linearisation method just mentioned, but rather it is *cancelled*. This is achieved by applying the *feedback control law*

$$\tau = -mg \sin \theta + \tau' \quad (2.5)$$

to (2.1), where τ' is a new torque input, and results in the *closed loop system*

$$m\ddot{\theta} = \tau'. \quad (2.6)$$

Equation (2.6) describes a linear system with input τ' , shown in Figure 4.

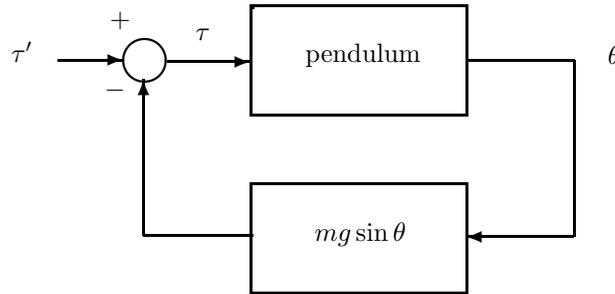


FIGURE 4. Feedback linearisation of the pendulum.

Linear systems methods can now be used to choose τ' in such a way that the pendulum is stabilised. Indeed, let's set

$$\tau' = -k_1 \dot{\theta} - k_2 \theta.$$

Then (2.6) becomes

$$m\ddot{\theta} + k_1 \dot{\theta} + k_2 \theta = 0. \quad (2.7)$$

If we select the *feedback gains*

$$k_1 = 3m, \quad k_2 = 2m,$$

then the system (2.7) has general solution $\alpha e^{-t} + \beta e^{-2t}$, and so $\theta = 0$ is now globally asymptotically stable. Thus the effect of any disturbance will decay, and the pendulum will be restored to its vertical position.

The feedback controller just designed and applied to (2.1) is

$$\tau = -mg \sin \theta - 3m\dot{\theta} - 2m\theta. \quad (2.8)$$

The torque τ is an explicit function of the angle θ and its derivative, the angular velocity $\dot{\theta}$. The controller is a *feedback* controller because it takes measurements of θ and $\dot{\theta}$ and uses this *information* to adjust the motor torque in a way which is stabilising, see Figure 5. For instance, if θ is non-zero, then the last term (proportional term) in (2.8) has the effect of forcing the motor to act in a direction opposite to the natural tendency to fall. The second term (derivative term) in (2.8) responds to the speed of the pendulum. An additional integral term may also be added. The proportional-integral-derivative (PID) combination is widespread in applications.

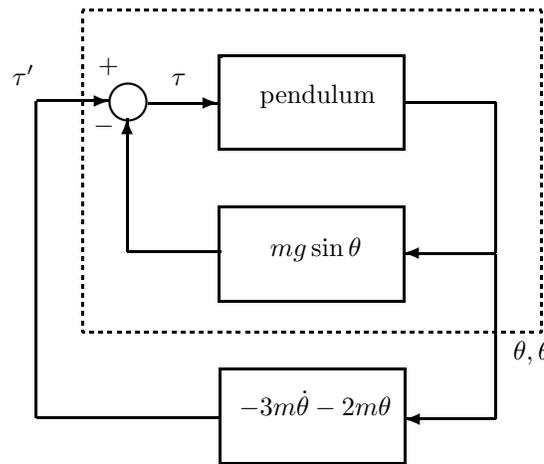


FIGURE 5. Feedback stabilisation of the pendulum using feedback linearisation.

It is worth noting that the feedback controller (2.8) has fundamentally altered the dynamics of the pendulum. With the controller in place, it is no longer an unstable nonlinear system. Indeed, in a different context, it was reported in the article Hunt & Johnson (1993) that it is possible to remove the effects of *chaos* with a suitable controller, even using a simple proportional control method.

Note that this design procedure requires explicit knowledge of the model parameters (length, mass, etc), and if they are not known exactly, performance may be degraded. Similarly, unmodelled influences (such as friction, motor dynamics) also impose significant practical limitations. Ideally, one would like a design which is *robust* and tolerates these negative effects.

The control problem is substantially more complicated if measurements of the joint angle and/or angular velocity are not available. The problem is then one of partial state information, and an output feedback design is required. These issues are discussed in §4.2 and §12.

It is also interesting to note that the natural state space (corresponding to $(\theta, \dot{\theta})$) for this example is the torus $M = S^1 \times \mathbf{R}$ —a manifold and not a vector space. Thus it is not surprising to find that differential geometric methods provide important tools in nonlinear control theory.

3. State Space Models

The systems discussed in §2.1 were represented in terms of block diagrams showing the information flows among key variables, the various inputs and outputs. No internal details were given for the blocks. In §2.2 internal details were given by the physical model for the pendulum, expressed in terms of a differential equation for the joint angle θ . In general, it is often convenient to use a differential equation model to describe the input-output behaviour of a system, and we will adopt this type of description throughout this chapter. Systems will be described by first-order differential equations in *state space* form, viz:

$$\begin{cases} \dot{x} &= f(x, u) \\ y &= h(x), \end{cases} \quad (3.1)$$

where the state x takes values in \mathbf{R}^n (or in some manifold M in which case we interpret (3.1) in local coordinates), the input $u \in U \subset \mathbf{R}^m$, the output $y \in \mathbf{R}^p$. In the case of the pendulum, the choice for a state mentioned above was $x = (\theta, \dot{\theta})$. The controlled vector field f is assumed sufficiently smooth, as is the observation function h . (In this chapter we will not concern ourselves with precise regularity assumptions.)

The state x need not necessarily have a physical interpretation, and such representations of systems are far from unique. The state space model may be derived from fundamental principles, or via an *identification* procedure. In either case, we assume such a model is given. If a control input $u(t)$, $t \geq 0$ is applied to the system, then the state trajectory $x(t)$, $t \geq 0$ is determined by solving the differential equation (3.1), and the output is given by $y(t) = h(x(t))$, $t \geq 0$.

The state summarises the internal status of the system, and its value $x(t)$ at time t is sufficient to determine future values $x(s)$, $s > t$ given the control values $u(s)$, $s > t$. The state is thus a very important quantity, and the ability of the controls to manipulate it as well as the information about it available in the observations are crucial factors in determining the quality of performance that can be achieved.

A lot can be done at the input-output level, without regard to the internal details of the system blocks. For instance, the *small gain theorem* concerns stability of a closed loop system in an input-output sense, see Vidyasagar (1993). However, in this chapter, we will restrict our attention to state space models and methods based on them.

4. Controllability and Observability

In this section we look at two fundamental issues concerning the nature of nonlinear systems. The model we use for a nonlinear system Σ is

$$\begin{cases} \dot{x} &= f(x) + g(x)u \\ y &= h(x), \end{cases} \quad (4.1)$$

where f , g and h are smooth. In equation (4.1), $g(x)u$ is short for $\sum_{i=1}^m g_i(x)u_i$, where g_i , $i = 1, \dots, m$ are smooth vector fields. It is of interest to know how well the controls

u influence the states x (controllability and reachability), and how much information concerning the states is available in the measured outputs y (observability and reconstructability.)

Controllability and observability have been the subject of a large amount of research over the last few decades; we refer the reader to the paper Hermann & Krener (1977) and the books Isidori (1985), Nijmeijer & van der Schaft (1990), Vidyasagar (1993), and the references contained in them. The differential geometric theory for controllability is “dual” to the corresponding observability theory, in that the former uses vector fields and distributions while the latter uses 1-forms and codistributions. In both cases, the Frobenius theorem is a basic tool, see Boothby (1975). The Frobenius theorem concerns the existence of (integral) submanifolds which are tangent at each point to a given field of subspaces (of the tangent spaces), and is a generalisation of the idea that solutions to differential equation define one dimensional submanifolds tangent to the one dimensional subspaces spanned by the vector field.

4.1. Controllability

Here is a basic question. Given two states x^0 and x^1 , is it possible to find a control $t \mapsto u(t)$ which steers Σ from x^0 to x^1 ? If so, we say that x^1 is *reachable* from x^0 , or that x^0 is *controllable* to x^1 . This question can be answered (at least locally) using differential geometric methods.

Consider the following special situation with two inputs and no drift f :

$$\dot{x} = u_1 g_1(x) + u_2 g_2(x), \quad (4.2)$$

where g_1 and g_2 are smooth vector fields in \mathbf{R}^3 ($n = 3$). The problem is to determine, locally, the states which can be reached from a given point x^0 .

Suppose that in a neighbourhood of x^0 the vector fields g_1, g_2 are linearly independent, so that (4.2) can move in two independent directions: setting $u_2 = 0$ and $u_1 = \pm 1$ causes Σ to move along integral curves of g_1 , and similarly movement along g_2 is effected by setting $u_1 = 0$ and $u_2 = \pm 1$. Clearly Σ can move in the direction of any linear combination of g_1 and g_2 . But can Σ move in a direction independent of g_1 and g_2 ? The answer is yes provided the following algebraic condition holds:

$$g_1(x^0), g_2(x^0), [g_1, g_2](x^0) \text{ linearly independent.} \quad (4.3)$$

Here, $[g_1, g_2]$ denotes the Lie bracket of the two vector fields g_1 and g_2 , defined by

$$[g_1, g_2] = \frac{\partial g_2}{\partial x} g_1 - \frac{\partial g_1}{\partial x} g_2, \quad (4.4)$$

see Boothby (1975). To move in the $[g_1, g_2]$ direction, the control must be switched between g_1 and g_2 appropriately. Indeed, let

$$u(t) = (u_1(t), u_2(t)) = \begin{cases} (1, 0) & \text{if } 0 \leq t \leq \tau \\ (0, 1) & \text{if } \tau \leq t \leq 2\tau \\ (-1, 0) & \text{if } 2\tau \leq t \leq 3\tau \\ (0, -1) & \text{if } 3\tau \leq t \leq 4\tau. \end{cases}$$

Then the solution of (4.2) satisfies

$$x(4\tau) = x^0 + \frac{1}{2}\tau^2 [g_1, g_2](x^0) + O(\tau^3)$$

for small τ ; readily verified using Taylor’s formula. This is illustrated in Figure 6.

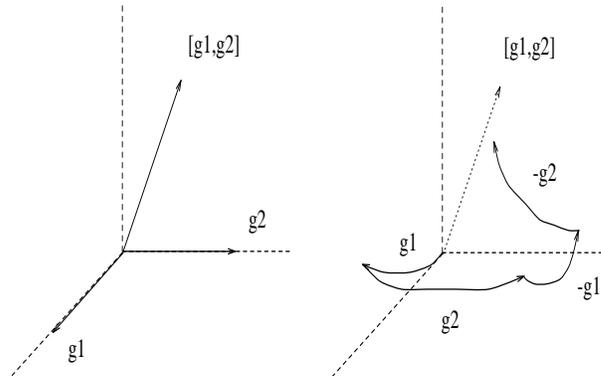


FIGURE 6. Lie bracket approximated by switching.

(By more complicated switching, it is possible to move along higher-order Lie brackets.) Using the inverse function theorem, it is possible to show that condition (4.3) implies that the set of points reachable from x^0 contains a non-empty open set.

We mention a few aspects of the general theory, concerning system (4.1). Define, for a neighbourhood V of x^0 , $R^V(x^0, T)$ to be the set of states which can be reached from x^0 in time $t \leq T$ with trajectory staying inside V . Let

$$\mathcal{R} = \text{Lie}\{f, g_i, i = 1, \dots, m\}$$

denote the Lie algebra generated by the vector fields appearing in (4.1). \mathcal{R} contains all linear combinations of these vector fields and all Lie brackets, and determines all possible directions in which Σ can move using appropriate switching. Assume that the following *reachability rank condition* holds:

$$\dim \mathcal{R}(x^0) = n. \quad (4.5)$$

Here, $\mathcal{R}(x)$ is a vector subspace of the tangent space $T_x \mathbf{R}^n$, and is spanned by $\{X(x) : X \in \mathcal{R}\}$. The mapping $x \mapsto \mathcal{R}(x)$ is known as a (smooth) distribution. Then for any neighbourhood V of x^0 and any $T > 0$, the set $R^V(x^0, T)$ contains a non-empty open subset.

EXAMPLE 4.1. Consider the system in \mathbf{R}^2 : $\dot{x}_1 = u$, $\dot{x}_2 = x_1^2$. Here $f(x_1, x_2) = (0, x_1^2)$ and $g(x_1, x_2) = (1, 0)$. Then $[f, g] = -(0, 2x_1)$ and $[[f, g], g] = (0, 2)$, and consequently $\dim \mathcal{R}(x_1, x_2) = 2$ for all $(x_1, x_2) \in \mathbf{R}^2$. Therefore the reachability rank condition is satisfied everywhere. For any open neighbourhood V of $x^0 = (0, 0)$, we have $R^V((0, 0), T) = V \cap \{x_2 > 0\}$. Note that points in the upper half plane are reachable from $(0, 0)$, but it is not possible to control such points to $(0, 0)$. Thus systems with drift are not generally reversible. \square

If the reachability rank condition fails, motion may be restricted to lower-dimensional submanifolds. Indeed, if $\dim \mathcal{R}(x) = k < n$ for all x in a neighbourhood W containing x^0 , then it is possible to show, using the Frobenius theorem, that there is a manifold R_{x^0} passing through x^0 which contains $R^V(x^0, T)$ for all neighbourhoods $V \subset W$ of x^0 , $T > 0$. Further, $R^V(x^0, T)$ contains a relatively open subset of R_{x^0} .

4.2. Observability

For many reasons it is useful to have knowledge of the state of the system being controlled. However, it is not always possible to measure all components of the state, and so knowledge of them must be inferred from the outputs or measurements that are available. In control system design, filters or observers are used to compute an estimate $\hat{x}(t)$

of the state $x(t)$ by processing the outputs $y(s)$, $0 \leq s \leq t$. See §12.1. We now ask what information about the states can be obtained from such an output record.

Two states x^0 and x^1 are *indistinguishable* if no matter what control is applied, the corresponding state trajectories always produce the same output record—there is no way of telling the two states apart by watching the system.

For simplicity, we begin with an uncontrolled single-output system in \mathbf{R}^2 ,

$$\begin{cases} \dot{x} &= f(x) \\ y &= h(x), \end{cases} \quad (4.6)$$

and consider the problem of distinguishing a state x^0 from its neighbours. This is possible if the following algebraic condition holds:

$$dh(x^0), dL_f h(x^0) \text{ linearly independent.} \quad (4.7)$$

Here, for a real valued function φ , $d\varphi = (\frac{\partial\varphi}{\partial x_1}, \dots, \frac{\partial\varphi}{\partial x_n})$ (a 1-form), and $L_f\varphi$ denotes the Lie derivative

$$L_f\varphi = d\varphi \cdot f = \sum_{i=1}^n \frac{\partial\varphi}{\partial x_i} f_i.$$

Let $x^i(t, x^i)$ and $y^i(t, x^i)$, $i = 1, 2$, denote the state and output trajectories corresponding to the initial states x^0, x^1 . If these two states are indistinguishable, then

$$y^0(t, x^0) = y^1(t, x^1), \text{ for all } t \geq 0. \quad (4.8)$$

Evaluating (4.8) at $t = 0$ gives

$$h(x^0) = h(x^1), \quad (4.9)$$

and differentiating (4.8) with respect to t and setting $t = 0$ gives

$$L_f h(x^0) = L_f h(x^1). \quad (4.10)$$

Define $\Phi : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ by $\Phi(x) = (h(x) - h(x^0), L_f h(x) - L_f h(x^0))$. Then $\frac{\partial\Phi}{\partial x}(x^0) = (dh(x^0), dL_f h(x^0))$, and so the condition (4.7) implies, by the inverse function theorem, the existence of an open set V of x^0 such that states in V are distinguishable from x^0 . This follows from (4.9) and (4.10) because if x^1 is indistinguishable from x^0 , then $\Phi(x^1) = \Phi(x^0) = (0, 0)$. The inverse function theorem implies that Φ is a local diffeomorphism, forcing $x^1 = x^0$ for any indistinguishable x^1 near x^0 .

Returning to the general situation, given a neighbourhood V of x^0 , define $I^V(x^0)$ to be the set of states in V which are indistinguishable from x^0 with trajectories staying inside V . Let \mathcal{O} denote the smallest vector subspace of $C^\infty(\mathbf{R}^n)$ containing h_1, \dots, h_p and closed with respect to Lie differentiation by f, g_1, \dots, g_m . \mathcal{O} is called the *observation space*, and it defines the *observability codistribution*

$$d\mathcal{O}(x) = \text{span} \{d\varphi : \varphi \in \mathcal{O}\}.$$

For each $x \in \mathbf{R}^n$, $d\mathcal{O}(x)$ is a subspace of the cotangent space $T_x^*\mathbf{R}^n$. If the *observability rank condition*

$$\dim \mathcal{O}(x^0) = n \quad (4.11)$$

holds, then there exists a neighbourhood W of x^0 such that for all neighbourhoods $V \subset W$ of x^0 ,

$$I^V(x^0) \cap V = \{x^0\}.$$

EXAMPLE 4.2. Consider the 1-dimensional system: $\dot{x} = u$, $y = \sin x$. Here $f(x) = 0$, $g(x) = 1$, and $h(x) = \sin x$. The observation space \mathcal{O} is the span (over \mathbf{R}) of $\{\sin x, \cos x\}$,

and the observability codistribution is $d\mathcal{O}(x) = \text{span}\{\cos x, \sin x\} = T_x^*\mathbf{R} \equiv \mathbf{R}$ for all $x \in \mathbf{R}$. Thus the observability rank condition is everywhere satisfied. Thus $I^{(-\pi, \pi)}(0) \cap (-\pi, \pi) = \{0\}$. In this sense the system is locally observable, but not globally, since $I^{\mathbf{R}}(0) = \{2\pi k, k \in \mathbf{Z}\}$. (It is possible to construct a quotient system on \mathbf{R}/\mathbf{Z} which is observable.) \square

In cases where the observability rank condition fails, the indistinguishable states may be contained in integral submanifolds of the observability codistribution. Precise results can be obtained using the Frobenius theorem.

5. Feedback Linearisation

In §2 we saw that it was possible to linearise the inverted pendulum (or one-link manipulator) using nonlinear state feedback. The most general results involve state space transformations as well (Hunt *et al* (1983), Nijmeijer & van der Schaft (1990), Isidori (1985), Vidyasagar (1993)). Feedback linearisation methods have been successfully applied in the design of a helicopter control system.

For clarity, we consider the following single-input system Σ ($m = 1$):

$$\dot{x} = f(x) + g(x)u, \quad (5.1)$$

where f and g are smooth vector fields with $f(0) = 0$. The (local) feedback linearisation problem is to find, if possible, a (local) diffeomorphism S on \mathbf{R}^n with $S(0) = 0$, and a state feedback law

$$u = \alpha(x) + \beta(x)v \quad (5.2)$$

with $\alpha(0) = 0$ and $\beta(x)$ invertible for all x such that the resulting system

$$\dot{x} = f(x) + g(x)\alpha(x) + g(x)\beta(x)v \quad (5.3)$$

transforms under $z = S(x)$ to a controllable linear system

$$\dot{z} = Az + Bv. \quad (5.4)$$

Thus the closed loop system relating v and z is linear, see Figure 7.

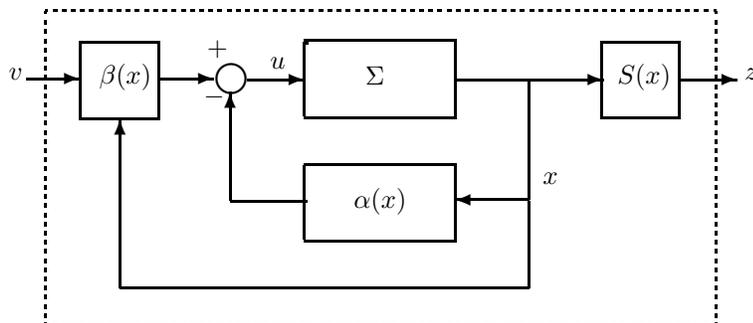


FIGURE 7. Feedback linearisation.

The main results concerning this single-input case assert that the nonlinear system (5.1) is feedback linearisable to the controllable linear system (5.4) if and only if

- (i) $\{L_f^k g, 0 \leq k \leq n-1\}$ is linearly independent, and
- (ii) $\{L_f^k g, 0 \leq k \leq n-2\}$ is involutive.

(Vector fields X_1, \dots, X_r are involutive if each Lie bracket $[X_i, X_j]$ can be expressed as a linear combination of X_1, \dots, X_r .)

Once a system has been feedback linearised, further control design can take place using linear systems methods to obtain a suitable control law $v = Kz = KS(x)$ achieving a desired goal (e.g. stabilisation as in §2).

EXAMPLE 5.1. In robotics, this procedure is readily implemented, and is known as the *computed torque method*, see Craig (1989). The dynamic equations for an n -link manipulator are

$$M(\theta)\ddot{\theta} + V(\theta, \dot{\theta}) + G(\theta) = \tau, \quad (5.5)$$

where θ is an n -vector of joint angles, $M(\theta)$ is the (invertible) mass matrix, $V(\theta, \dot{\theta})$ consists of centrifugal and coriolis terms, $G(\theta)$ accounts for gravity, and τ is the n -vector of joint torques applied at each joint. Feedback linearisation is achieved by the choice

$$\tau = M(\theta)\tau' + V(\theta, \dot{\theta}) + G(\theta). \quad (5.6)$$

The resulting linear system is

$$\ddot{\theta} = \tau'. \quad (5.7)$$

The state $x = (\theta, \dot{\theta})$ is $2n$ -dimensional, and the reader can easily express the dynamics and feedback law in state space form. \square

6. Feedback Stabilisation

Stabilisation is one of the most important problems in control theory. After all, one cannot do much with a system that is wildly out of control! In this section we take a quick look at some of the issues that arise.

The (local) feedback stabilisation problem for the system

$$\dot{x} = f(x) + g(x)u, \quad (6.1)$$

where f and $g = g_1, \dots, g_m$ are smooth vector fields with $f(0) = 0$, is to find, if possible, a function $\alpha : \mathbf{R}^n \rightarrow \mathbf{R}^m$ such that the state feedback law

$$u = \alpha(x) \quad (6.2)$$

results in a (locally) asymptotically stable closed loop system:

$$\dot{x} = f(x) + g(x)\alpha(x). \quad (6.3)$$

That is, for any initial condition x^0 (in a neighbourhood of 0), the resulting closed-loop trajectory of (6.3) should satisfy $x(t) \rightarrow 0$ as $t \rightarrow \infty$. As is apparent from our discussion so far, systems which are feedback linearisable can be stabilised by smooth feedback simply by using a stabilising linear feedback for the linearised system in conjunction with the linearising transformation and feedback.

According to classical ordinary differential equation (ODE) theory, in order that solutions to (6.3) exist and are unique, the function α must satisfy some smoothness requirements, say Lipschitz continuity. Thus it is not unreasonable to restrict one's search to include only sufficiently smooth α . Also, considering the definition of controllability, one may guess that some form of this property is needed. However, a necessary condition was given in Brockett (1983) which implies that controllability is not enough for there to exist a continuous stabilising feedback of the form (6.2). Indeed, there does not exist a continuous stabilising feedback for the drift-free system

$$\dot{x} = g(x)u, \quad (6.4)$$

when $m < n$, even if controllable. Such systems cannot even be stabilised by continuous dynamic feedback. Thus one is forced to relax the smoothness requirement, or look for different types of feedback functions.

Recently, some remarkable results have been obtained which assert that satisfaction of the reachability rank condition is enough to imply the existence of a stabilising *time-varying periodic* feedback

$$u = \alpha(x, t) \quad (6.5)$$

for drift-free systems. It is very interesting that time-varying feedback laws can succeed when time-invariant ones fail. Generalisations of these results to the case of systems with drift have been obtained.

EXAMPLE 6.1. Consider the following drift-free system in \mathbf{R}^3 (Pomet (1992)):

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} u_1 + \begin{pmatrix} 0 \\ x_1 \\ 1 \end{pmatrix} u_2. \quad (6.6)$$

Since $m = 2 < n = 3$, this system cannot be stabilised by a continuous feedback. This system satisfies the reachability rank condition, since, as the reader can verify, the vector fields g_1 , g_2 and $[g_1, g_2]$ are everywhere linearly independent. It is stabilised by the periodic feedback

$$u_1 = x_2 \sin t - (x_1 + x_2 \cos t) \quad (6.7)$$

$$u_2 = -(x_1 + x_2 \cos t)x_1 \cos t - (x_1 x_2 + x_3).$$

The reader can check that this feedback is stabilising by showing that the function

$$V(t, x) = \frac{1}{2}(x_1 + x_2 \cos t)^2 + \frac{1}{2}x_2^2 + \frac{1}{2}x_3^2 \quad (6.8)$$

is a *Lyapunov function* for the time-varying closed-loop system:

$$\begin{aligned} \frac{d}{dt}V(t, x(t)) &= -(x_1(t) + x_2(t) \cos t)^2 \\ &\quad - (x_1(t)x_2(t) + x_3(t) + x_1(t)(x_1(t) + x_2(t) \cos t) \cos t)^2 < 0. \end{aligned}$$

(For a detailed discussion of Lyapunov methods, see, e.g., Vidyasagar (1993).) \square

A number of other interesting approaches to the stabilisation problem are available in the literature. To give one more example, artificial neural networks have been proposed as a means for constructing stabilising feedback functions α , see, e.g., Sontag (1992), Sontag (1993). Because, in general, continuous stabilisers do not exist, one has to go beyond single-hidden-layer nets with continuous activation. In fact, (asymptotic) controllability is enough to ensure the existence of stabilising controllers using two-hidden-layer nets, Sontag (1992).

In §9 we shall see that stabilising controllers can sometimes be obtained using optimal control methods.

7. Optimisation-Based Design

Dynamic optimisation techniques constitute a very powerful and general design methodology. Optimisation techniques are frequently used in engineering and other application areas. The basic idea is to formulate the design problem in terms of an optimisation problem whose solution is regarded as the ‘best’ design. The engineering component is the formulation of the optimisation problem, whereas the procedure for solving the optimisation problem relies on mathematics. The practical utility of dynamic optimisation

methods for nonlinear systems is limited by the computational complexity involved, see §13.

Optimal control and game theory are subjects with long and rich histories, and there are many excellent books discussing various aspects; see, e.g., Basar & Olsder (1982), Bellman (1957), Fleming & Rishel (1975), Fleming & Soner (1993), Pontryagin *et al* (1962) and the references contained in them. The remainder of the chapter is concerned with optimisation-based design.

8. Deterministic Optimal Control

In deterministic optimal control theory, there is no uncertainty, and the evolution of the system is determined precisely by the initial state and the control applied. We associate with the nonlinear system

$$\dot{x} = f(x, u) \quad (8.1)$$

a *performance index* or *cost function*

$$J = \int_{t_0}^{t_1} L(x(s), u(s)) ds + \psi(x(t_1)), \quad (8.2)$$

where $x(\cdot)$ is the trajectory resulting from the initial condition $x(t_0) = x^0$ and control $u(\cdot)$. The optimal control problem is to find a control $u(\cdot)$ (open loop) which minimises the cost J .

The function $L(x, u)$ assigns a cost along the state and control trajectories, while ψ is used to provide a penalty for the final state $x(t_1)$. For example, if one wishes to find a controller which steers the system from x^0 to near the origin with low control energy, then simply choose

$$L(x, u) = \frac{1}{2}|x|^2 + \frac{1}{2}|u|^2 \quad (8.3)$$

and $\psi(x) = \frac{1}{2}|x|^2$. Some weighting factors can be included if desired. The cost function penalises trajectories which are away from 0 and use large values of u . In general, these functions are chosen to match the problem at hand.

There are two fundamental methods for solving and analysing such optimal control problems. One is concerned with necessary conditions, the *Pontryagin minimum principle* (PMP) (Pontryagin *et al* (1962)), while the other, *dynamic programming* (DP) (Bellman (1957)), is concerned with sufficient conditions and verification.

8.1. Pontryagin Minimum Principle

Let's formally derive the necessary conditions to be satisfied by an optimal control $u^*(\cdot)$. Introduce Lagrange multipliers $\lambda(t)$, $t_0 \leq t \leq t_1$ to include the dynamics (8.1) (a constraint) in the cost function J :

$$\bar{J} = \int_{t_0}^{t_1} [L(x(s), u(s)) + \lambda(s)'(f(x(s), u(s)) - \dot{x}(s))] ds + \psi(x(t_1)).$$

Integrating by parts, we get

$$\bar{J} = \int_{t_0}^{t_1} [H(x(s), u(s), \lambda(s)) + \dot{\lambda}(s)'x(s)] ds - \lambda(t_1)'x(t_1) + \lambda(t_0)'x(t_0) + \psi(x(t_1)),$$

where $H(x, u, \lambda) = L(x, u) + \lambda'f(x, u)$. Consider a small control variation δu , and apply the control $u + \delta u$. This leads to variations δx and $\delta \bar{J}$ in the state trajectory and

augmented cost function. Indeed, we have

$$\delta\bar{J} = [\psi_x - \lambda']\delta x|_{t=t_1} + \lambda'\delta x|_{t=t_0} + \int_{t_0}^{t_1} [(H_x + \dot{\lambda}')\delta x + H_u\delta u] ds.$$

If u^* is optimal, then it is necessary that $\delta\bar{J} = 0$. These considerations lead to the following equations which must be satisfied by an optimal control u^* and resulting trajectory x^* :

$$\dot{x}^* = f(x^*, u^*), \quad t_0 \leq t \leq t_1, \quad (8.4)$$

$$\dot{\lambda}^{*'} = -H_x(x^*, u^*, \lambda^*), \quad t_0 \leq t \leq t_1, \quad (8.5)$$

$$H(x^*, u^*, \lambda^*) = \min_u H(x^*, u, \lambda^*), \quad t_0 \leq t \leq t_1, \quad (8.6)$$

$$\lambda^*(t_1) = \psi_x(x^*(t_1))', \quad x^*(t_0) = x^0. \quad (8.7)$$

The PMP gets its name from (8.6), which asserts that the Hamiltonian H must be minimised along the optimal trajectory. The optimal control is open loop (it is not a feedback control).

To solve an optimal control problem using the PMP, one first solves the two-point boundary value problem (8.4)–(8.7) for a candidate optimal control, state and adjoint trajectories. One must then check the optimality of the candidate. Note that there is no guarantee that an optimal control will exist in general, and this should be checked. In practice, solving (8.4)–(8.7) must be done numerically, and there are many methods available for doing this, see, e.g., Jennings *et al* (1991), Teo *et al* (1991), Hindmarsh (1983).

An interesting application of this method is to the Moon Landing Problem; see the example discussed in the book Fleming & Rishel (1975).

8.2. Dynamic Programming

The DP approach considers a family of optimisation problems parameterised by the initial time t and state $x(t) = x$:

$$J(x, t; u) = \int_t^{t_1} L(x(s), u(s)) ds + \psi(x(t_1)). \quad (8.8)$$

The *value function* describes the optimal value of this cost:

$$V(x, t) = \inf_{u(\cdot)} J(x, t; u), \quad (8.9)$$

where the minimum is over all (admissible) open loop controls. The fundamental principle of dynamic programming asserts that for any $r \in [t, t_1]$,

$$V(x, t) = \inf_{u(\cdot)} \left[\int_t^r L(x(s), u(s)) ds + V(x(r), r) \right]. \quad (8.10)$$

Thus the optimal cost starting at time t equals the optimal value of the accumulated cost on the interval $[t, r]$, plus the optimal cost obtained if we started again at time r , but from the state $x(r)$. This principle leads to a partial differential equation (PDE) to be satisfied by V , called the dynamic programming equation (DPE) or Hamilton–Jacobi–Bellman (HJB) equation:

$$\begin{cases} \frac{\partial V}{\partial t} + \min_u [DV \cdot f(x, u) + L(x, u)] = 0 & \text{in } \mathbf{R}^n \times (t_0, t_1), \\ V(x, t_1) = \psi(x) & \text{in } \mathbf{R}^n. \end{cases} \quad (8.11)$$

The DPE here is a nonlinear first-order PDE, the significance of which is the *verification theorem* (Fleming & Rishel (1975), Fleming & Soner (1993)). If there exists a continuously differentiable solution V of the DPE (8.11) (satisfying some additional technical conditions) and if $\mathbf{u}^*(x, t)$ is the control value attaining the minimum in (8.11), then

$$u^*(t) = \mathbf{u}^*(x^*(t), t), \quad t_0 \leq t \leq t_1, \quad (8.12)$$

is optimal, where $x^*(\cdot)$ is the resulting optimal state trajectory. It is very important to note that *the optimal control is a state feedback controller!* This turns out to be a general principle in DP, and so *feedback arises naturally in dynamic optimisation.*

To see why (8.12) is optimal, let $u(\cdot)$ be any control. If V is continuously differentiable, then

$$V(x(t_1), t_1) = V(x, t) + \int_t^{t_1} \left[\frac{\partial V}{\partial t}(x(s), s) + DV(x(s), s) \cdot f(x(s), u(s)) \right] ds,$$

which together with (8.11) implies

$$\begin{aligned} V(x, t) &= - \int_t^{t_1} \left[\frac{\partial V}{\partial t}(x(s), s) + DV(x(s), s) \cdot f(x(s), u(s)) \right] ds + \psi(x(t_1)) \\ &\leq \int_t^{t_1} L(x(s), u(s)) ds + \psi(x(t_1)) \end{aligned}$$

with equality if $u = u^*$. Setting $t = t_0$ and $x(t_0) = x^0$ we get

$$\begin{aligned} J &= V(x^0, 0) \\ &= \int_{t_0}^{t_1} L(x^*(s), u^*(s)) ds + \psi(x^*(t_1)) \\ &\leq \int_{t_0}^{t_1} L(x(s), u(s)) ds + \psi(x(t_1)) \end{aligned}$$

for any other control u .

In general, the DPE (8.11) does not have smooth solutions, and so the verification theorem no longer applies in this form. In fact, the DPE must be interpreted in a generalised sense, viz., the viscosity sense (Crandall & Lions (1984), Fleming & Soner (1993)). However, more sophisticated verification theorems are available using the methods of non-smooth analysis, see Clarke (1983). The controllers obtained are not, in general, continuous.

EXAMPLE 8.1. To see that optimal controllers can easily be discontinuous, consider the classical minimum time problem. The dynamics are $\dot{x}_1 = x_2$, $\dot{x}_2 = u$, and the controls take values in the compact set $U = [-1, 1]$. The problem is to find the controller which steers the system from an initial state $x = (x_1, x_2)$ to the origin in minimum time. The value function is simply the minimum time possible:

$$T(x) = \inf_{u(\cdot)} \left\{ \int_0^{t_f} 1 dt : x(0) = x, x(t_f) = 0 \right\}. \quad (8.13)$$

For a number of reasons, it is convenient to use the discounted minimum time function

$$S(x) = 1 - e^{-T(x)},$$

with DPE

$$\begin{cases} S(x) &= \min_{u \in [-1, 1]} \{ DS(x) \cdot f(x, u) + 1 \} \\ S(0) &= 0. \end{cases} \quad (8.14)$$

The value function $S(x)$ is not differentiable at $x = 0$, rather, it is only Holder continuous

there (with exponent $\frac{1}{2}$), see Figure 8. The optimal state feedback $\mathbf{u}^*(x)$ equals either -1 or $+1$, and has a single discontinuity across the so-called switching curve, see Figure 9.

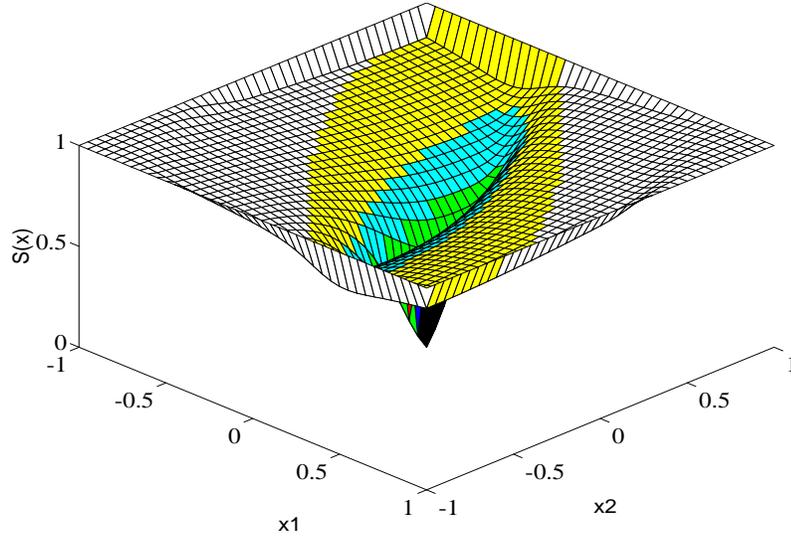


FIGURE 8. A non-smooth value function.

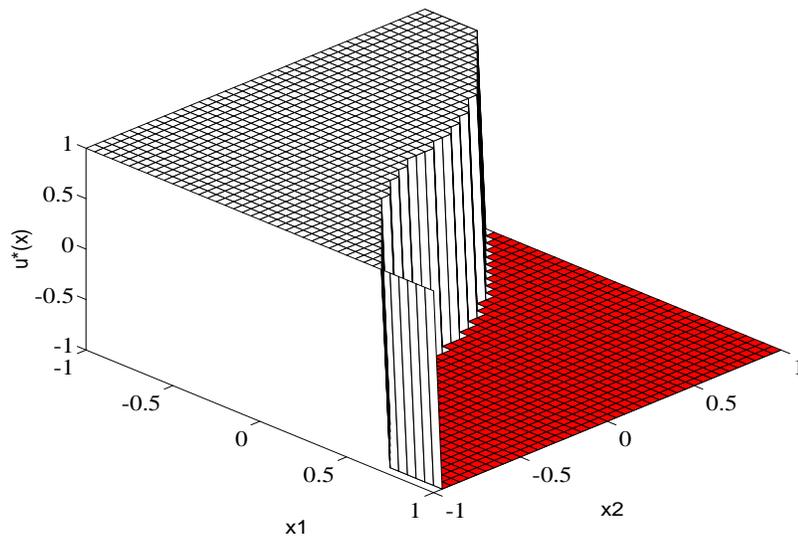


FIGURE 9. A discontinuous optimal feedback controller.

□

9. Optimal Stabilisation

The basic optimal control problem described in §8 is a *finite horizon* problem, because it is defined on a finite time interval. *Infinite horizon* problems are defined on $[0, \infty)$, and so the asymptotic behaviour of the trajectories and controls is an issue. Indeed, infinite horizon problems can be used to obtain stabilising state feedback controllers. (The minimum time problem is also an infinite horizon problem, but of a different type.)

To this end, we consider the following infinite horizon value function:

$$V(x) = \inf_{u(\cdot)} \int_0^{\infty} L(x(s), u(s)) ds. \quad (9.1)$$

The DPE is now of the following stationary type

$$\min_u [DV \cdot f(x, u) + L(x, u)] = 0 \text{ in } \mathbf{R}^n, \quad (9.2)$$

and, assuming validity of the verification theorem, the optimal controller $\mathbf{u}^*(x)$ is the control value achieving the minimum in (9.2):

$$u^*(t) = \mathbf{u}^*(x^*(t)), \quad t \geq 0. \quad (9.3)$$

Under suitable hypotheses, this controller is a stabilising state feedback controller. For instance, a sufficiently strong form of controllability would be needed, the cost term $L(x, u)$ must adequately reflect the behaviour of the states (observability), and $\mathbf{u}^*(x)$ must be sufficiently regular. However, it can easily happen that no controller is capable of making the cost (9.1) finite. Often, a discount factor is used, especially in the stochastic case.

EXAMPLE 9.1. Consider the inverted pendulum discussed in §2. A state space form of the dynamic equation (2.1) is

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} x_2 \\ g \sin x_1 \end{pmatrix} + \begin{pmatrix} 0 \\ g/m \end{pmatrix} \tau. \quad (9.4)$$

Choose the cost term L as in (8.3), in which case the value function (9.1) takes the form

$$V(x) = \inf_{\tau(\cdot)} \int_0^{\infty} \frac{1}{2} (|x_1(s)|^2 + |x_2(s)|^2 + |\tau(s)|^2) ds. \quad (9.5)$$

Since the inverted pendulum is (exponentially) feedback stabilisable (with smooth feedback), the value function will be finite, but not necessarily smooth. The optimal state feedback controller, if it exists, need not necessarily be smooth. However, it is possible to solve this problem numerically by solving the DPE (9.2) using the methods described in §13. Figures 10–12 show the value function, optimal state feedback control, and a simulation showing that it is stabilising.

From Figure 11, it is evident that the control is linear near the origin, and so is close to the optimal controller obtained by using the linearisation $\sin x_1 \approx x_1$ there. Away from the origin, the control is definitely nonlinear. The control appears to be continuous in the region shown in Figure 11. Note that the initial angle in the simulation, Figure 12, is well away from the linear region. \square

10. Stochastic Optimal Control

The simple deterministic optimal control problems discussed above do not take into account disturbances which may affect performance. In stochastic control theory dis-

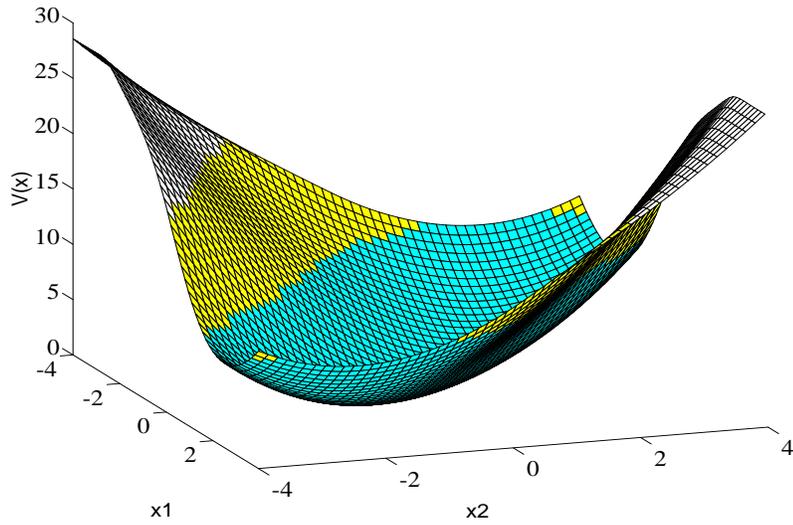


FIGURE 10. Value function for the pendulum.

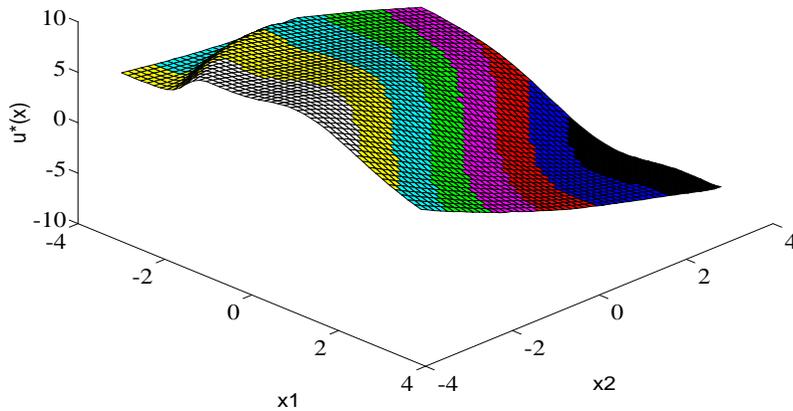


FIGURE 11. Optimal stabilising feedback for the pendulum.

turbances are modelled as random or stochastic processes, and the performance index averages over them.

The nonlinear system (8.1) is replaced by the *stochastic differential equation* model

$$dx = f(x, u) dt + \varepsilon dw, \quad (10.1)$$

where w is a standard Brownian motion. Formally, “ $\frac{dw}{dt}$ ” is white noise, and models disturbances as random fluctuations. If the noise variance ε^2 is small, then the solution of (10.1) is close to that of (8.1) (c.f. semiclassical limits in physics). The cost function

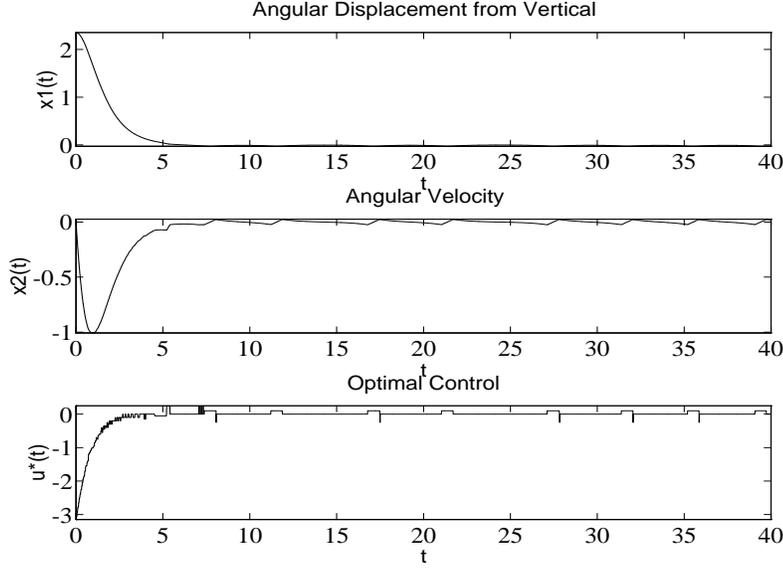


FIGURE 12. Optimal trajectories for the stabilised pendulum.

is

$$J = \mathbf{E} \left[\int_{t_0}^{t_1} L(x(s), u(s)) ds + \psi(x(t_1)) \right], \quad (10.2)$$

where the expectation denotes integration with respect to the probability measure associated with w .

It turns out that DP methods are very convenient for stochastic problems, although there are PMP-type results available. The value function is defined by

$$V(x, t) = \inf_{u(\cdot)} \mathbf{E}_{x,t} \left[\int_t^{t_1} L(x(s), u(s)) ds + \psi(x(t_1)) \right], \quad (10.3)$$

and the DPE is

$$\begin{cases} \frac{\partial V}{\partial t} + \frac{\varepsilon^2}{2} \Delta V + \min_u [DV \cdot f(x, u) + L(x, u)] = 0 & \text{in } \mathbf{R}^n \times (t_0, t_1), \\ V(x, t_1) = \psi(x) & \text{in } \mathbf{R}^n. \end{cases} \quad (10.4)$$

This PDE is a nonlinear second-order parabolic equation, and in this particular case smooth solutions can exist due to the (non-degenerate) Laplacian term. Again, the optimal control is a state feedback controller $\mathbf{u}^{\varepsilon^*}(x, t)$ obtained by minimising the appropriate expression in the DPE (10.4):

$$u^{\varepsilon^*}(t) = \mathbf{u}^{\varepsilon^*}(x^*(t), t), \quad t_0 \leq t \leq t_1. \quad (10.5)$$

Under certain conditions, the controller $\mathbf{u}^{\varepsilon^*}(x, t)$ is related to the deterministic controller $\mathbf{u}^*(x, t)$ obtained in §8.2 by

$$\mathbf{u}^{\varepsilon^*}(x, t) = \mathbf{u}^*(x, t) + \varepsilon^2 \mathbf{u}_{stoch}(x, t) + \dots, \quad (10.6)$$

for small $\varepsilon > 0$ and for some term \mathbf{u}_{stoch} which, to first-order, compensates for the disturbances described by the stochastic model.

11. Robust Control

An alternative approach to stochastic control is to model the disturbance as an unknown but deterministic L_2 signal. This is known as robust H_∞ control, and is discussed, e.g., in the book Basar & Bernhard (1991).

The system model is

$$\dot{x} = f(x, u) + w, \quad (11.1)$$

where $w \in L_2$ models disturbances and uncertainties such as plant model errors and noise. Associated with (11.1) is a performance variable

$$z = \ell(x, u). \quad (11.2)$$

For instance, one could take $|z|^2$ to be the cost term $L(x, u)$.

The problem is to find, if possible, a stabilising controller which ensures that the effect of the disturbance w on the performance z is bounded in the following sense:

$$\int_{t_0}^{t_1} |z(s)|^2 ds \leq \gamma^2 \int_{t_0}^{t_1} |w(s)|^2 ds + \beta_0(x(t_0)) \quad (11.3)$$

for all $w \in L_2([t_0, t_1], \mathbf{R}^n)$, and all $t_1 \geq t_0$, for some finite β_0 . Here, $\gamma > 0$ is given, and is interpreted as essentially a bound on an induced norm of the input-output map $w \mapsto z$. When $w = 0$, the ideal situation (as in deterministic control), the size of z is limited by the initial state (via β_0), but if $w \neq 0$, the resulting effect on z is limited by (11.3).

The robust H_∞ control problem can be solved using game theory methods (see, e.g. Basar & Bernhard (1991) and James & Baras (1993)), which are closely related to the above-mentioned optimal control techniques. Henceforth we shall consider the following special game problem, for fixed $t_0 < t_1$.

We seek to minimise the worst-case performance index

$$J = \sup_{w(\cdot)} \left\{ \int_{t_0}^{t_1} [L(x(s), u(s)) - \frac{1}{2}\gamma^2 |w(s)|^2] ds + \psi(x(t_1)) \right\}. \quad (11.4)$$

If this cost is finite, then the resulting optimal controller achieves the goal (11.3). Applying DP, the value function is given by

$$V(x, t) = \inf_{u(\cdot)} \sup_{w(\cdot)} \left\{ \int_t^{t_1} [L(x(s), u(s)) - \frac{1}{2}\gamma^2 |w(s)|^2] ds + \psi(x(t_1)) \right\}, \quad (11.5)$$

where $x(t) = x$, as in §8.2. The DPE is the first-order nonlinear PDE, and is called the Hamilton-Jacobi-Isaacs (HJI) equation:

$$\begin{cases} \frac{\partial V}{\partial t} + \min_u \max_w [DV \cdot (f(x, u) + w) + L(x, u) - \frac{1}{2}\gamma^2 |w|^2] = 0 & \text{in } \mathbf{R}^n \times (t_0, t_1), \\ V(x, t_1) = \psi(x) & \text{in } \mathbf{R}^n. \end{cases} \quad (11.6)$$

If the DPE (11.6) has a smooth solution, then the optimal controller is given by a value $\mathbf{u}^{\gamma^*}(x, t)$ which attains the minimum in (11.6). This optimal state feedback controller is related to the optimal state feedback controller for the simple deterministic optimal control problem in §8.2 by

$$\mathbf{u}^{\gamma^*}(x, t) = \mathbf{u}^*(x, t) + \frac{1}{\gamma^2} \mathbf{u}_{robust}(x, t) + \dots, \quad (11.7)$$

for large γ , and for some \mathbf{u}_{robust} . The term \mathbf{u}_{robust} depends on the worst-case disturbance energy, and contributes to the controller appropriately.

The infinite horizon problem can also be discussed, yielding stabilising robust feedback controllers, assuming the validity of a suitable verification theorem, etc.

12. Output Feedback

The optimal feedback controllers discussed above all feed back the state x . This makes sense since no restriction was placed on the information available to the controller, and the state is by definition a quantity summarising the internal status of the system. As mentioned earlier, in many cases the state is not completely available, and the controller must use only the information gained from observing the outputs y —the controller can only (causally) feed back the output:

$$u(t) = \mathbf{u}(y(s), 0 \leq s \leq t). \quad (12.1)$$

This *information constraint* has a dramatic affect on the complexity of the optimal solution, as we shall soon see.

12.1. Deterministic Optimal Control

There is no theory of output feedback optimal control for deterministic systems without disturbances. For such systems, a state feedback controller is equivalent to an open loop controller, since in the absence of uncertainty the trajectory is completely determined by the initial state and the control applied. In the presence of disturbances (either deterministic or stochastic) this is not the case.

Because of the error correcting nature of feedback, it is still desirable to implement a feedback controller even if the design does not specifically account for disturbances. In the output feedback case, the following (suboptimal) procedure can be adopted. We suppose the output is

$$y = h(x). \quad (12.2)$$

First, solve a deterministic optimal control problem for an optimal stabilising state feedback controller $\mathbf{u}_{state}^*(x)$, as in §9.1. Then design an *observer* or *filter*, say of the form

$$\dot{\hat{x}} = G(\hat{x}, u, y) \quad (12.3)$$

to (asymptotically) estimate the state:

$$|\hat{x}(t) - x(t)| \rightarrow 0 \text{ as } t \rightarrow \infty. \quad (12.4)$$

Note that this filter uses only the information contained in the outputs and in the inputs applied. Finally implement the output feedback controller obtained by substituting the estimate into the optimal state feedback controller, see Figure 13: $u(t) = \mathbf{u}_{state}^*(\hat{x}(t))$. Under certain conditions, the closed-loop system will be asymptotically stable, see Vidyasagar (1980).

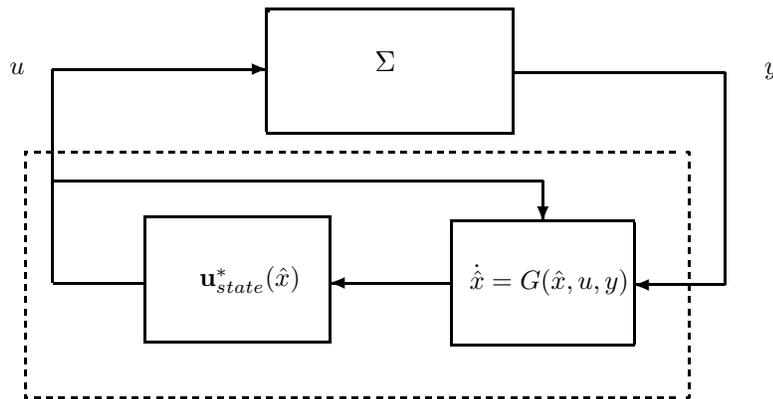


FIGURE 13. Suboptimal controller/observer structure.

12.2. *Stochastic Optimal Control*

In the stochastic framework, suppose the output is given by the equation

$$dy = h(x) dt + dv, \quad (12.5)$$

where v is a standard Brownian motion (independent of w in the state equation (10.1)). The controller must satisfy the constraint (12.1).

As above, one can combine the optimal state feedback controller with a state estimate, but this output feedback controller will not generally minimise the cost J (it will in the case of linear systems with quadratic cost). Rather, the optimal controller feeds back an *information state* $\sigma_t = \sigma(x, t)$ —an infinite dimensional quantity. In the case at hand (but not necessarily always), σ can be taken to be an unnormalised conditional density of $x(t)$ given $y(s)$, $0 \leq s \leq t$:

$$\mathbf{E}[\phi(x(t))|y(s), 0 \leq s \leq t] = \frac{\int_{\mathbf{R}^n} \phi(x) \sigma(x, t) dx}{\int_{\mathbf{R}^n} \sigma(x, t) dx}.$$

The information state is the solution of the stochastic PDE

$$d\sigma = A(u)^* \sigma dt + h\sigma dy, \quad (12.6)$$

where $A(u)^*$ is the formal adjoint of the linear operator $A(u) = \frac{\varepsilon^2}{2} \Delta + f(\cdot, u) \cdot D$. The equation (12.6) is called the Duncan-Mortensen-Zakai equation of nonlinear filtering, Zakai (1969). It is a *filter* which processes the output and produces a quantity relevant to state estimation.

Using properties of conditional expectation, the cost function J can be expressed purely in terms of σ , and the optimal control problem becomes one of full state information, with σ as the new state. Now that a suitable state has been defined for the output feedback optimal control problem, dynamic programming methods can be applied. The DPE for the value function $V(\sigma, t)$ is

$$\begin{cases} \frac{\partial V}{\partial t} + \frac{1}{2} D^2 V(h\sigma, h\sigma) + \min_{u \in U} [DV \cdot A(u)^* \sigma + \langle \sigma, L(\cdot, u) \rangle] = 0 \\ V(\sigma, T) = \langle \sigma, \psi \rangle, \end{cases} \quad (12.7)$$

where $\langle \sigma, \psi \rangle = \int_{\mathbf{R}^n} \sigma(x) \psi(x) dx$. This PDE is defined on the infinite dimensional space $L_2(\mathbf{R}^n)$, and can be interpreted in the viscosity sense, see Lions (1988). This equation is considerably more complicated than the corresponding state feedback DPE (12.7). In the spirit of dynamic programming, the optimal controller is obtained from the value $\mathbf{u}^*(\sigma, t)$ which attains the minimum in (12.7):

$$u^*(t) = \mathbf{u}^*(\sigma_t^*, t). \quad (12.8)$$

This controller is an output feedback controller, since it feeds back the information state, a quantity depending on the outputs, see Figure 14. Dynamic programming leads naturally to this feedback structure.

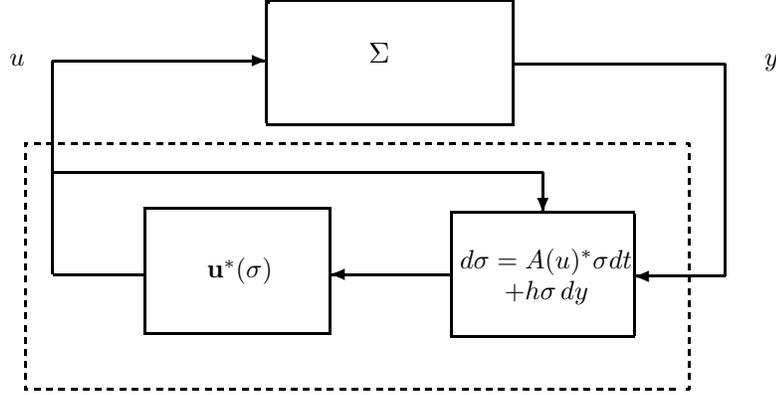


FIGURE 14. Optimal information state feedback: stochastic control.

In general, it is difficult to prove rigorously results of this type; for related results, see, e.g., Fleming & Pardoux (1982), Hijab (1989).

The choice of information state in general depends on the cost function as well as on the system model. Indeed, for the class of *risk-sensitive* cost functions, the information state includes terms in the cost, see James *et al* (1994), James *et al* (1993).

12.3. Robust Control

Consider the robust control problem with dynamics (11.1) and outputs

$$y = h(x) + v, \quad (12.9)$$

where v is another unknown deterministic L_2 signal. The resulting differential game can be transformed to an equivalent game with full state information by defining an appropriate information state, see James *et al* (1994), James *et al* (1993), James & Baras (1993).

The information state $p_t = p(x, t)$ is the solution of the nonlinear first-order PDE

$$\dot{p} = F(p, u, y), \quad (12.10)$$

where F is the nonlinear operator

$$F(p, u, y) = \max_w [Dp \cdot (f(x, u) + w) + L(x, u) - \frac{1}{2}\gamma^2(|w|^2 + |y - h(x)|^2)]. \quad (12.11)$$

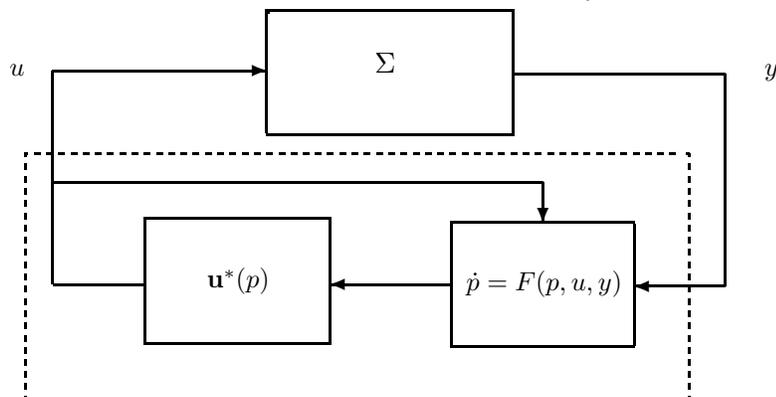
Again, equation (12.10) is a filter which processes the output to give a quantity relevant to the control problem (not necessarily state estimation), and is infinite dimensional. The value function $W(p, t)$ satisfies the DPE

$$\begin{cases} \frac{\partial W}{\partial t} + \min_{u \in U} \max_{y \in \mathbf{R}^p} \{DW \cdot F(p, u, y)\} = 0 \\ W(p, T) = (p, \psi), \end{cases} \quad (12.12)$$

where $(p, \psi) = \sup_{x \in \mathbf{R}^n} (p(x) + \psi(x))$ is the “sup-pairing” James *et al* (1994). This PDE is defined on a suitable (infinite dimensional) function space.

If the DPE (12.12) has a sufficiently smooth solution and if $\mathbf{u}^*(p, t)$ is the control value achieving the minimum, then the optimal output feedback controller is:

$$u(t) = \mathbf{u}^*(p_t^*, t). \quad (12.13)$$

FIGURE 15. Optimal information state feedback: robust H_∞ control.

Precise results concerning this problem are currently being obtained. Note that the information state for the stochastic problem is quite different to the one used in the robust control problem. In general, an information state provides the correct notion of “state” for an optimal control or game problem.

13. Computation and Complexity

The practical utility of optimisation procedures depends on how effective the methods for solving the optimisation problem are. Linear programming is popular in part because of the availability of efficient numerical algorithms capable of handling real problems (simplex algorithm). The same is true for optimal design of linear control systems (matrix Riccati equations). For nonlinear systems, the computational issues are much more involved. To obtain candidate optimal open loop controllers using the Pontryagin minimum principle, a nonlinear two-point boundary value problem must be solved. There are a number of numerical schemes available, see Jennings *et al* (1991), Teo *et al* (1991), Hindmarsh (1983). The dynamic programming method, as we have seen, requires the solution of a PDE to obtain the optimal feedback controller. In the next section, we describe a finite difference numerical method, and in the section following, make some remarks about its computational complexity. In summary, effective computational methods are crucial for optimisation-based design. It will be interesting to see nonlinear dynamic optimisation flourish as computational power continues to increase and as innovative optimal and suboptimal algorithms and approximation methods are developed.

13.1. Finite-Difference Approximation

Finite difference and finite element schemes are commonly employed to solve both linear and nonlinear PDEs. We present a finite difference scheme for the stabilisation problem discussed in §9. This entails the solution of the nonlinear PDE (9.2). The method follows the general principles in the book Kushner & Dupuis (1993).

For $h > 0$ define the grid $h\mathbf{Z}^n = \{hz : z_i \in \mathbf{Z}, i = 1, \dots, n\}$. We wish to construct an approximation $V^h(x)$ to $V(x)$ on this grid. The basic idea is to approximate the derivative $DV(x) \cdot f(x, u)$ by the finite difference

$$\sum_{i=1}^n f_i^\pm(x, u) (V(x \pm he_i) - V(x)) / h, \quad (13.1)$$

where e_1, \dots, e_n are the standard unit vectors in \mathbf{R}^n , $f_i^+ = \max(f_i, 0)$ and $f_i^- = -\min(f_i, 0)$. A finite difference replacement of (9.2) is

$$V^h(x) = \min_u \left\{ \sum_z p^h(x, z|u) V^h(z) + \Delta t(x) L(x, u) \right\}. \quad (13.2)$$

Here,

$$p^h(x, z|u) = \begin{cases} \frac{\|f(x, u)\|_1}{\max_u \|f(x, u)\|_1} & \text{if } z = x \\ \frac{f_i^\pm(x, u)}{\max_u \|f(x, u)\|_1} & \text{if } z = x \pm h e_i \\ 0 & \text{otherwise,} \end{cases} \quad (13.3)$$

$\|v\|_1 = \sum_{i=1}^n |v_i|$, and $\Delta t(x) = h / \max_u \|f(x, u)\|_1$. Equation (13.2) is obtained by substituting (13.1) for $DV \cdot f$ in (9.2) and some manipulation. In Kushner & Dupuis (1993), the quantities $p^h(x, z|u)$ are interpreted as transition probabilities for a controlled Markov chain, with (13.2) as the corresponding DPE. The optimal state feedback policy $\mathbf{u}^{*h}(x)$ for this discrete optimal control problem is obtained by finding the control value attaining the minimum in (13.2). Using this interpretation, convergence results can be proven. In our case, it is interesting that a deterministic system is approximated by a (discrete) stochastic system. See Figure 16. One can imagine a particle moving randomly on the grid trying to follow the deterministic motion or flow, on average. When the particle is at a grid point, it jumps to one of its neighbouring grid points according to the transition probability.

In practice, one must use only a finite set of grid points, say a set forming a box D^h centred at the origin. On the boundary ∂D^h , the vector field $f(x, u)$ is modified by projection, so if $f(x, u)$ points out of the box, then the relevant components are set to zero. This has the effect of constraining the dynamics to a bounded region. From now on we take (13.2) to be defined on D^h .

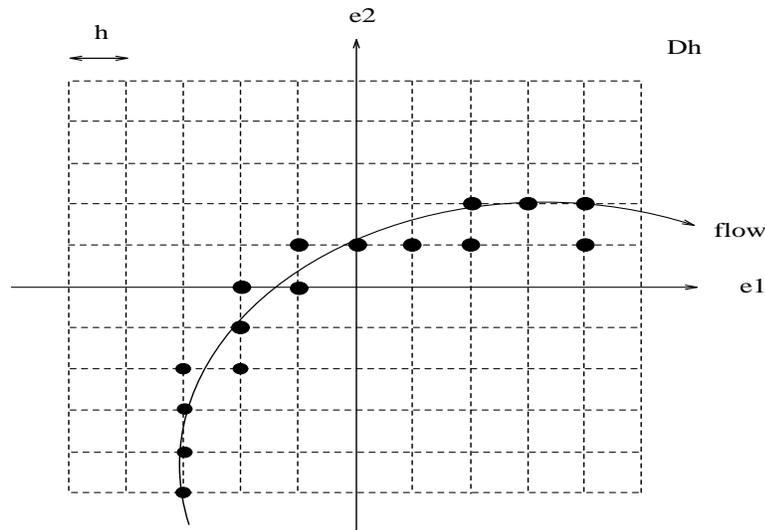


FIGURE 16. The finite difference grid showing a randomly moving particle following the deterministic flow.

Equation (13.2) is a nonlinear implicit relation for V^h , and iterative methods are used to solve it (approximately). There are two main types of methods, viz., *value space iteration*, which generates a sequence of approximations

$$V_k^h(x) = \min_u \left\{ \sum_z p^h(x, z|u) V_{k-1}^h(z) + \Delta t(x) L(x, u) \right\} \quad (13.4)$$

to V^h , and *policy space iteration*, which involves a sequence of approximations $\mathbf{u}_k^{*h}(x)$ to the optimal feedback policy $\mathbf{u}^{*h}(x)$. There are many variations and combinations of these methods, acceleration procedures, and multigrid algorithms. Note that because of the local structure of equation (13.2), *parallel computation* is natural.

Numerical results for the stabilisation of the inverted pendulum are shown in §9. The nonlinear state feedback controller $\mathbf{u}^{*h}(x)$ (Figure 11) was effective in stabilising the pendulum (Figure 12).

13.2. Computational Complexity

13.2.1. State Feedback

One can see easily from the previous subsection that the computational effort demanded by dynamic programming can quickly become prohibitive. If N is the number of grid points in each direction e_i , then the storage requirements for $V^h(x)$, $x \in D^h$ are of order N^n , a number growing exponentially with the number of state dimensions. If $N \sim 1/h$, this measure of computational complexity is roughly

$$O(1/h^n).$$

The complexity of a discounted stochastic optimal control problem was analysed by Chow & Tsitsiklis (1989). If $\varepsilon > 0$ is the desired accuracy and if the state and control spaces are of dimension n and m , then the computational complexity is

$$O(1/\varepsilon^{2n+m}),$$

measured in terms of the number of times functions of x and u need be called in an algorithm to achieve the given accuracy.

Thus the computational complexity of state feedback dynamic programming is of exponential order in the state space dimension, n . This effectively limits its use to low dimensional problems. For $n = 1$ or 2 , it is possible to solve problems on a desktop workstation. For $n = 3$ or 4 , a supercomputer such as a Connection Machine CM-2 or CM-5 is usually needed. If n is larger, then in general other methods of approximation or computation are required, see, e.g., Campillo & Pardoux (1992).

13.2.2. Output Feedback

In the state feedback case discussed above, the DPE is defined in a finite dimensional space \mathbf{R}^n . However, in the case of output feedback, the DPE is defined on an *infinite dimensional* (function) space (in which the information state takes values). Consequently, the computational complexity is vastly greater. In fact, if n is the dimension of the state x , and a grid of size $h > 0$ is used, then the computational complexity is roughly

$$O(1/h^{1/h^n}),$$

a number increasing doubly exponentially with $n!$ In the output feedback case, it is particularly imperative that effective approximations and numerical methods be developed.

14. Concluding Remarks

In this chapter we have surveyed a selection of topics in nonlinear control theory, with an emphasis on *feedback design*. As we have seen, feedback is central to control systems, and techniques from differential geometry and dynamic optimisation play leading roles. Feedback is used to stabilise and regulate a system in the presence of disturbances and uncertainty, and the main problem of control engineers is to design feedback controllers. In practice, issues of robustness and computation are crucial, and in situations with partial information (output feedback), the complexity can be very high.

The author wishes to thank A.J. Cahill for his assistance in preparing the simulations. This work was supported in part by Boeing BCAC and by the Cooperative Research Centre for Robust and Adaptive Systems.

REFERENCES

- ANDERSON, B.D.O., AND MOORE, J.B. (1990), *Optimal Control: Linear Quadratic Methods*, Prentice-Hall, Englewood Cliffs NJ.
- BASAR, T., AND BERNHARD, P. (1991), *H_∞ -Optimal Control and Related Minimax Design Problems*, Birkhauser, Boston.
- BASAR, T., AND OLSDER, G.J. (1982), *Dynamic Noncooperative Game Theory*, Academic Press, New York.
- BELLMAN, R. (1957), *Dynamic Programming*, Princeton Univ. Press, Princeton, NJ.
- BOOTHBY, W.M. (1975), *An Introduction to Differentiable Manifolds and Riemannian Geometry*, Academic Press, New York.
- BROCKETT, R.W. (1976a), Nonlinear Systems and Differential Geometry, *Proc. IEEE*, **64**(1) 61—72.
- BROCKETT, R.W. (1976b), Control Theory and Analytical Mechanics, in: C. Martin and R. Hermann, Eds., *The 1976 Ames Research Center (NASA) Conference on Geometric Control Theory*, Math Sci Press, Brookline, MA.
- BROCKETT, R.W. (1983), Asymptotic Stability and Feedback Stabilization, in: R.W. Brockett, R.S. Millman and H.J. Sussmann, Eds., *Differential Geometric Control Theory*, Birkhauser, Basel-Boston.
- CAMPILLO, F., AND PARDOUX, E. (1992), Numerical Methods in Ergodic Optimal Stochastic Control and Application, in: I. Karatzas and D. Ocone (Eds.), *Applied Stochastic Analysis*, LNCIS **177**, Springer Verlag, New York, 59—73.
- CRANDALL, M.G. AND LIONS, P.L. (1984), Viscosity Solutions of Hamilton-Jacobi Equations, *Trans. AMS*, **277** 183—186.
- CHOW, C.S., AND TSITSIKLIS, J.N. (1989), The Computational Complexity of Dynamic Programming, *J. Complexity*, **5** 466—488.
- CLARKE, F.H. (1983), *Optimization and Non-Smooth Analysis*, Wiley-Interscience, New York.
- CORON, J.M. (1992), Global Asymptotic Stabilization for Controllable Systems without Drift, *Math. Control, Signals and Systems*, **5** 295—312.
- CRAIG, J.J. (1989), *Introduction to Robotics: Mechanics and Control* (2nd Ed.), Addison Wesley, New York.
- FLEMING, W.H. (1988), Chair, *Future Directions in Control Theory: A Mathematical Perspective*, Panel Report, SIAM, Philadelphia.
- FLEMING, W.H. AND PARDOUX, E. (1982), Optimal Control for Partially Observed Diffusions, *SIAM J. Control Optim.*, **20**(2) 261—285.
- FLEMING, W.H., AND RISHEL, R.W. (1975), *Deterministic and Stochastic Optimal Control*, Springer Verlag, New York.
- FLEMING, W.H., AND SONER, H.M. (1993), *Controlled Markov Processes and Viscosity Solutions*, Springer Verlag, New York.

- HERMANN, R., AND KRENER, A.J. (1977), Nonlinear Controllability and Observability, *IEEE Trans. Automatic Control*, **22-5** 728—740.
- HIJAB, O. (1990), Partially Observed Control of Markov Processes, *Stochastics*, **28** 123—144.
- HINDMARSH, A.C. (1983), ODEPACK: A systemized Collection of ODE Solvers, in: *Scientific Computing*, R.S. Stepleman et al (Eds.), Noth Holland, Amsterdam, 55—64.
- HUNT, E.R., AND JOHNSON, G. (1993), Controlling Chaos, *IEEE Spectrum*, **30-11** 32—36.
- HUNT, L.R., SU, R., AND MEYER, G. (1983), Global Transformations of Nonlinear Systems, *IEEE Trans. Automatic Control*, **28** 24—31.
- ISIDORI, A. (1985), *Nonlinear Control Systems: An Introduction*, Springer Verlag, New York.
- JAMES, M.R., AND BARAS, J.S. (1993), Robust Output Feedback Control for Discrete-Time Nonlinear Systems, submitted to *IEEE Trans. Automatic Control*.
- JAMES, M.R., BARAS, J.S., AND ELLIOTT, R.J. (1993), Output Feedback Risk-Sensitive Control and Differential Games for Continuous-Time Nonlinear Systems, *32nd IEEE Control and Decision Conference*, San Antonio.
- JAMES, M.R., BARAS, J.S., AND ELLIOTT, R.J. (1994), Risk-Sensitive Control and Dynamic Games for Partially Observed Discrete-Time Nonlinear Systems, to appear, *IEEE Trans. Automatic Control*, April 1994.
- JENNINGS, L.S., FISHER, M.E., TEO, K.L., AND GOH, C.J. (1991), MISER3: Solving Optimal Control Problems—an Update, *Advances in Engineering Software and Workstations*, **13** 190—196.
- KAILATH, T. (1980), *Linear Systems*, Prentice-Hall, Englewood Cliffs NJ.
- KOKOTOVIC, P.V. (1985), Recent Trends in Feedback Design: An Overview, *Automatica*, **21-3** 225—236.
- KUMAR, P.R., AND VARAIYA, P. *Stochastic Systems: Estimation, Identification, and Adaptive Control*, Prentice-Hall, Englewood Cliffs NJ.
- KUSHNER, H.J., AND DUPUIS, P.G. (1993), *Numerical Methods for Stochastic Control Problems in Continuous Time*, Springer Verlag, New York.
- LIONS, P.L. (1988), Viscosity Solutions of Fully Nonlinear Second Order Equations and Optimal Stochastic Control in Infinite Dimensions, Part II: Optimal Control of Zakai's Equation, *Lecture Notes Math* **1390** 147—170.
- MEES, A.I. (1981), *Dynamics of Feedback Systems*, Wiley, Chichester.
- NIJMEIJER, H., AND VAN DER SCHAFT, A.J. (1990), *Nonlinear Dynamical Control Systems*, Springer Verlag, New York.
- POMET, J.B. (1992), Explicit Design of Time-Varying Stabilizing Control Laws for a Class of Controllable Systems without Drift, *Systems and Control Letters*, **18** 147—158.
- PONTRYAGIN, L.S., BOLTYANSKII, V.G., GAMKRILEDZE, R.V., AND MISCHENKO, E.F. (1962), *The Mathematical Theory of Optimal Processes*, Interscience, New York.
- SAMSON, C. (1990), Velocity and Torque Feedback Control of a Wheeled Mobile Robot: Stability Analysis, preprint, INRIA, Sophia Antipolis.
- SONTAG, E.D. (1992), Feedback Stabilization Using Two-Hidden-Layer Nets, *IEEE Tans. Neural Networks*, **3** 981—990.
- SONTAG, E.D. (1993), Some Topics in Neural Networks, Technical Report, Siemens Corporate Research, NJ.
- TEO, K.L., GOH, C.J., AND WONG, K.H. (1991), *A Unified Computational Approach to Optimal Control Problems*, Longman Scientific and Technical, London.
- VIDYASAGAR, M. (1980), On the Stabilization of Nonlinear Systems Using State Detection, *IEEE Trans. Automatic Control*, **25(3)** 504—509.
- VIDYASAGAR, M. (1993), *Nonlinear Systems Analysis*, Second Edition, Prentice Hall, Englewood Cliffs NJ.
- WILLEMS, J.C. (1971), The Generation of Lyapunov Functions for Input-Output Stable Systems, *SIAM J. Control*, **9** (1) 105—134.
- WILLEMS, J.C. (1972), Dissipative Dynamical Systems Part I: General Theory, *Arch. Rational*

Mech. Anal., **45** 321—351.

ZAKAI, M. (1969), On the Optimal Filtering of Diffusion Processes, *Z. Wahr. Verw. Geb.*, **11**
230—243.