Chapter 1

Introduction

1.1 What Is Mathematical Control Theory?

Mathematical control theory is the area of application-oriented mathematics that deals with the basic principles underlying the analysis and design of control systems. To *control* an object means to influence its behavior so as to achieve a desired goal. In order to implement this influence, engineers build devices that incorporate various mathematical techniques. These devices range from Watt's steam engine governor, designed during the English Industrial Revolution, to the sophisticated microprocessor controllers found in consumer items —such as CD players and automobiles— or in industrial robots and airplane autopilots.

The study of these devices and their interaction with the object being controlled is the subject of this book. While on the one hand one wants to understand the fundamental limitations that mathematics imposes on what is achievable, irrespective of the precise technology being used, it is also true that technology may well influence the type of question to be asked and the choice of mathematical model. An example of this is the use of difference rather than differential equations when one is interested in digital control.

Roughly speaking, there have been two main lines of work in control theory, which sometimes have seemed to proceed in very different directions but which are in fact complementary. One of these is based on the idea that a good model of the object to be controlled is available and that one wants to somehow *optimize* its behavior. For instance, physical principles and engineering specifications can be —and are— used in order to calculate that trajectory of a spacecraft which minimizes total travel time or fuel consumption. The techniques here are closely related to the classical calculus of variations and to other areas of optimization theory; the end result is typically a preprogrammed flight plan. The other main line of work is that based on the constraints imposed by *uncertainty* about the model or about the environment in which the object operates. The central tool here is the use of *feedback* in order to correct for deviations from the desired behavior. For instance, various feedback control systems are used during actual space flight in order to compensate for errors from the precomputed trajectory. Mathematically, stability theory, dynamical systems, and especially the theory of functions of a complex variable, have had a strong influence on this approach. It is widely recognized today that these two broad lines of work deal just with different aspects of the same problems, and we do not make an artificial distinction between them in this book.

Later on we shall give an axiomatic definition of what we mean by a "system" or "machine." Its role will be somewhat analogous to that played in mathematics by the definition of "function" as a set of ordered pairs: not itself the object of study, but a necessary foundation upon which the entire theoretical development will rest. In this Chapter, however, we dispense with precise definitions and will use a very simple physical example in order to give an intuitive presentation of some of the goals, terminology, and methodology of control theory.

The discussion here will be informal and not rigorous, but the reader is encouraged to follow it in detail, since the ideas to be given underlie everything else in the book. Without them, many problems may look artificial. Later, we often refer back to this Chapter for motivation.

1.2 Proportional-Derivative Control

One of the simplest problems in robotics is that of controlling the position of a single-link rotational joint using a motor placed at the pivot. Mathematically, this is just a pendulum to which one can apply a torque as an external force (see Figure 1.1).

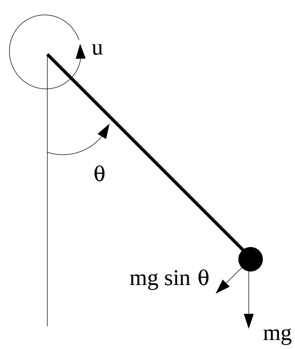

Figure 1.1: *Pendulum.*

We assume that friction is negligible, that all of the mass is concentrated at the end, and that the rod has unit length. From Newton's law for rotating objects, there results, in terms of the variable θ that describes the counterclockwise angle with respect to the vertical, the second-order nonlinear differential equation

$$m\ddot{\theta}(t) + mg\sin\theta(t) = u(t), \tag{1.1}$$

where m is the mass, g the acceleration due to gravity, and $u(t)$ the value of the external torque at time t (counterclockwise being positive). We call $u(\cdot)$ the *input* or *control* function. To avoid having to keep track of constants, let us assume that units of time and distance have been chosen so that $m = g = 1$.

The vertical stationary position ($\theta = \pi, \dot{\theta} = 0$) is an equilibrium when no control is being applied ($u \equiv 0$), but a small deviation from this will result in an unstable motion. Let us assume that our objective is to apply torques as needed to correct for such deviations. For small $\theta - \pi$,

$$\sin\theta = -(\theta - \pi) + o(\theta - \pi)\,.$$

Here we use the standard "little-o" notation: $o(x)$ stands for some function $g(x)$ for which

$$\lim_{x\to 0} \frac{g(x)}{x} = 0\,.$$

Since only small deviations are of interest, we drop the nonlinear part represented by the term $o(\theta-\pi)$. Thus, with $\varphi := \theta - \pi$ as a new variable, we replace equation (1.1) by the linear differential equation

$$\ddot{\varphi}(t) - \varphi(t) = u(t) \tag{1.2}$$

as our object of study. (See Figure 1.2.) Later we will analyze the effect of the ignored nonlinearity.

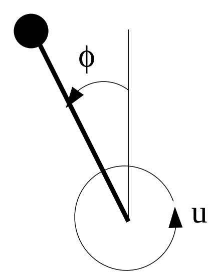

Figure 1.2: *Inverted pendulum.*

Our objective then is to bring φ and $\dot{\varphi}$ to zero, for any small nonzero initial values $\varphi(0), \dot{\varphi}(0)$ in equation (1.2), and preferably to do so as fast as possible, with few oscillations, and without ever letting the angle and velocity become too large. Although this is a highly simplified system, this kind of "servo" problem illustrates what is done in engineering practice. One typically wants to achieve a desired value for certain variables, such as the correct idling speed in an automobile's electronic ignition system or the position of the read/write head in a disk drive controller.

A naive first attempt at solving this control problem would be as follows: If we are to the left of the vertical, that is, if $\varphi = \theta - \pi > 0$, then we wish to move to the right, and therefore, we apply a negative torque. If instead we are to the right, we apply a positive, that is to say counterclockwise, torque. In other words, we apply *proportional feedback*

$$u(t) = -\alpha\varphi(t)\,, \tag{1.3}$$

where α is some positive real number, the *feedback gain*.

Let us analyze the resulting *closed-loop* equation obtained when the value of the control given by (1.3) is substituted into the *open-loop* original equation (1.2), that is

$$\ddot{\varphi}(t) - \varphi(t) + \alpha\varphi(t) = 0\,. \tag{1.4}$$

If $\alpha > 1$, the solutions of this differential equation are all oscillatory, since the roots of the associated characteristic equation

$$z^2 + \alpha - 1 = 0 \tag{1.5}$$

are purely imaginary, $z = \pm i\sqrt{\alpha - 1}$. If instead $\alpha < 1$, then all of the solutions except for those with

$$\dot{\varphi}(0) = -\varphi(0)\sqrt{1-\alpha}$$

diverge to $\pm\infty$. Finally, if $\alpha = 1$, then each set of initial values with $\dot{\varphi}(0) = 0$ is an equilibrium point of the closed-loop system. Therefore, in none of the cases is the system guaranteed to approach the desired configuration.

We have seen that proportional control does not work. We proved this for the linearized model, and an exercise below will show it directly for the original nonlinear equation (1.1). Intuitively, the problem can be understood as follows. Take first the case $\alpha < 1$. For any initial condition for which $\varphi(0)$ is small but positive and $\dot{\varphi}(0) = 0$, there results from equation (1.4) that $\ddot{\varphi}(0) > 0$. Therefore, also $\dot{\varphi}$ and hence φ increase, and the pendulum moves away, rather than toward, the vertical position. When $\alpha > 1$ the problem is more subtle: The torque is being applied in the correct direction to counteract the natural instability of the pendulum, but this feedback helps build too much inertia. In particular, when already close to $\varphi(0) = 0$ but moving at a relatively large speed, the controller (1.3) keeps pushing toward the vertical, and overshoot and eventual oscillation result.

The obvious solution is to keep $\alpha > 1$ but to modify the proportional feedback (1.3) through the addition of a term that acts as a brake, penalizing velocities. In other words, one needs to add damping to the system. We arrive then at a *PD*, or *proportional-derivative* feedback law,

$$u(t) = -\alpha\varphi(t) - \beta\dot{\varphi}(t)\,, \tag{1.6}$$

with $\alpha > 1$ and $\beta > 0$. In practice, implementing such a controller involves measurement of both the angular position and the velocity. If only the former is easily available, then one must estimate the velocity as part of the control algorithm; this will lead later to the idea of *observers*, which are techniques for reliably performing such an estimation. We assume here that $\dot{\varphi}$ can indeed be measured. Consider then the resulting closed-loop system,

$$\ddot{\varphi}(t) + \beta\dot{\varphi}(t) + (\alpha - 1)\varphi(t) = 0 \,. \tag{1.7}$$

The roots of its associated characteristic equation

$$z^2 + \beta z + \alpha - 1 = 0 \tag{1.8}$$

are

$$\frac{-\beta \pm \sqrt{\beta^2 - 4(\alpha - 1)}}{2} \,,$$

both of which have negative real parts. Thus all the solutions of (1.2) converge to zero. The system has been *stabilized* under feedback. This convergence may be oscillatory, but if we design the controller in such a way that in addition to the above conditions on α and β it is true that

$$\beta^2 > 4(\alpha - 1) \,, \tag{1.9}$$

then all of the solutions are combinations of decaying exponentials and no oscillation results.

We conclude from the above discussion that through a suitable choice of the gains α and β it is possible to attain the desired behavior, at least for the linearized model. That this same design will still work for the original nonlinear model, and, hence, assuming that this model was accurate, for a real pendulum, is due to what is perhaps the most important fact in control theory —and for that matter in much of mathematics— namely that first-order approximations are sufficient to characterize local behavior. Informally, we have the following *linearization principle*:

Designs based on linearizations work **locally** *for the original system*

The term "local" refers to the fact that satisfactory behavior only can be expected for those initial conditions that are close to the point about which the linearization was made. Of course, as with any "principle," this is not a theorem. It can only become so when precise meanings are assigned to the various terms and proper technical assumptions are made. Indeed, we will invest some effort in this text to isolate cases where this principle may be made rigorous. One of these cases will be that of stabilization, and the theorem there will imply that if we can stabilize the linearized system (1.2) for a certain choice of parameters α, β in the law (1.6), then the same control law does bring initial conditions of (1.1) that start close to $\theta = \pi, \dot{\theta} = 0$ to the vertical equilibrium.

Basically because of the linearization principle, a great deal of the literature in control theory deals exclusively with linear systems. From an engineering point of view, local solutions to control problems are often enough; when they are not, ad hoc methods sometimes may be used in order to "patch" together such local solutions, a procedure called *gain scheduling*. Sometimes, one may even be lucky and find a way to transform the problem of interest into one that is globally linear; we explain this later using again the pendulum as an example. In many other cases, however, a genuinely nonlinear approach is needed, and much research effort during the past few years has been directed toward that goal. In this text, when we develop the basic definitions and results for the linear theory we will always do so with an eye toward extensions to nonlinear, global, results.

An Exercise

As remarked earlier, proportional control (1.3) by itself is inadequate for the original nonlinear model. Using again $\varphi = \theta - \pi$, the closed-loop equation becomes

$$\ddot{\varphi}(t) - \sin\varphi(t) + \alpha\varphi(t) = 0\,. \tag{1.10}$$

The next exercise claims that solutions of this equation typically will not approach zero, no matter how the feedback gain α is picked.

Exercise 1.2.1 Assume that α is any fixed real number, and consider the ("energy") function of two real variables

$$V(x,y) := \cos x - 1 + \frac{1}{2}(\alpha x^2 + y^2)\,. \tag{1.11}$$

Show that $V(\varphi(t), \dot{\varphi}(t))$ is constant along the solutions of (1.10). Using that $V(x,0)$ is an analytic function and therefore that its zero at $x = 0$ is isolated, conclude that there are initial conditions of the type $\varphi(0) = \varepsilon, \dot{\varphi}(0) = 0$, with ε arbitrarily small, for which the corresponding solution of (1.10) does not satisfy that $\varphi(t) \to 0$ and $\dot{\varphi}(t) \to 0$ as $t \to \infty$. □

1.3 Digital Control

The actual physical implementation of (1.6) need not concern us here, but some remarks are in order. Assuming again that the values $\varphi(t)$ and $\dot{\varphi}(t)$, or equivalently $\theta(t)$ and $\dot{\theta}(t)$, can be measured, it is necessary to take a linear combination of these in order to determine the torque $u(t)$ that the motor must apply. Such combinations are readily carried out by circuits built out of devices called *operational amplifiers*. Alternatively, the damping term can be separately implemented directly through the use of an appropriate device (a "dashpot"), and the torque is then made proportional to $\varphi(t)$.

A more modern alternative, attractive especially for larger systems, is to convert position and velocity to digital form and to use a computer to calculate the necessary controls. Still using the linearized inverted pendulum as an illustration, we now describe some of the mathematical problems that this leads to.

A typical approach to computer control is based on the *sample-and-hold* technique, which can be described as follows. The values $\varphi(t)$ and $\dot\varphi(t)$ are measured only at discrete instants or *sampling times*

$$0,\ \delta,\ 2\delta,\ 3\delta, \ldots, k\delta, \ldots$$

The control law is updated by a program at each time $t = k\delta$ on the basis of the sampled values $\varphi(k\delta)$ and $\dot\varphi(k\delta)$. The output of this program, a value v_k, is then fed into the system as a control (*held constant* at that value) during the interval $[k\delta, k\delta + \delta]$.

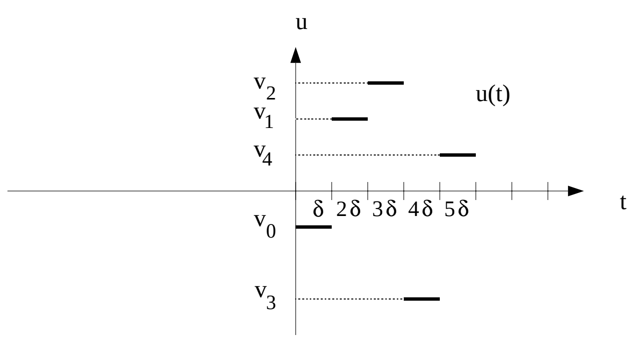

Figure 1.3: *Sampled control.*

For simplicity we assume here that the computation of v_k can be done quickly relative to the length δ of the sampling intervals; otherwise, the model must be modified to account for the extra delay. To calculate the effect of applying the constant control

$$u(t) \equiv v_k \quad \text{if} \quad t \in [k\delta, k\delta + \delta] \tag{1.12}$$

we solve the differential equation (1.2) with this function u. By differentiation one can verify that the general solution is, for $t \in [k\delta, k\delta + \delta]$,

$$\varphi(t) \;=\; \frac{\varphi(k\delta) + \dot\varphi(k\delta) + v_k}{2} e^{t-k\delta} + \frac{\varphi(k\delta) - \dot\varphi(k\delta) + v_k}{2} e^{-t+k\delta} - v_k\,, \tag{1.13}$$

so

$$\dot\varphi(t) \;=\; \frac{\varphi(k\delta) + \dot\varphi(k\delta) + v_k}{2} e^{t-k\delta} - \frac{\varphi(k\delta) - \dot\varphi(k\delta) + v_k}{2} e^{-t+k\delta}\,. \tag{1.14}$$

Thus, applying the constant control u gives rise to new values for $\varphi(k\delta + \delta)$ and $\dot\varphi(k\delta + \delta)$ at the end of the interval via the formula

$$\begin{pmatrix} \varphi(k\delta + \delta) \\ \dot\varphi(k\delta + \delta) \end{pmatrix} = A \begin{pmatrix} \varphi(k\delta) \\ \dot\varphi(k\delta) \end{pmatrix} + Bv_k\,, \tag{1.15}$$

where

$$A = \begin{pmatrix} \cosh\delta & \sinh\delta \\ \sinh\delta & \cosh\delta \end{pmatrix} \tag{1.16}$$

and

$$B = \begin{pmatrix} \cosh\delta - 1 \\ \sinh\delta \end{pmatrix}. \tag{1.17}$$

In other words, if we let $x_0, x_1, \ldots$ denote the sequence of two dimensional vectors

$$x_k := \begin{pmatrix} \varphi(k\delta) \\ \dot{\varphi}(k\delta) \end{pmatrix},$$

then $\{x_k\}$ satisfies the recursion

$$x_{k+1} = Ax_k + Bv_k\,. \tag{1.18}$$

Assume now that we wish to program our computer to calculate the constant control values v_k to be applied during any interval via a linear transformation

$$v_k := Fx_k \tag{1.19}$$

of the measured values of position and velocity at the start of the interval. Here F is just a row vector (f_1, f_2) that gives the coefficients of a linear combination of these measured values. Formally we are in a situation analogous to the PD control (1.6), except that we now assume that the measurements are being made only at discrete times and that a constant control will be applied on each interval. Substituting (1.19) into the difference equation (1.18), there results the new difference equation

$$x_{k+1} = (A + BF)x_k\,. \tag{1.20}$$

Since for any k

$$x_{k+2} = (A + BF)^2 x_k\,, \tag{1.21}$$

it follows that, if one finds gains f_1 and f_2 with the property that the matrix $A + BF$ is nilpotent, that is,

$$(A + BF)^2 = 0\,, \tag{1.22}$$

then we would have a controller with the property that after two sampling steps necessarily $x_{k+2} = 0$. That is, both φ and $\dot{\varphi}$ vanish after these two steps, and the system remains at rest after that. This is the objective that we wanted to achieve all along. We now show that this choice of gains is always possible. Consider the characteristic polynomial

$$\begin{aligned} \det(zI - A - BF) \;=\; z^2 \;&+\; (-2\cosh\delta - f_2\sinh\delta - f_1\cosh\delta + f_1)z \\ &-\; f_1\cosh\delta + 1 + f_1 + f_2\sinh\delta\,. \end{aligned} \tag{1.23}$$

It follows from the Cayley-Hamilton Theorem that condition (1.22) will hold provided that this polynomial reduces to just z^2. So we need to solve for the

f_i the system of equations resulting from setting the coefficient of z and the constant term to zero. This gives

$$f_1 = -1/2 \frac{2\cosh\delta - 1}{\cosh\delta - 1} \quad \text{and} \quad f_2 = -1/2 \frac{2\cosh\delta + 1}{\sinh\delta}\,. \tag{1.24}$$

We conclude that it is always possible to find a matrix F as desired. In other words, using sampled control we have been able to achieve stabilization of the system. Moreover, this stability is of a very strong type, in that, at least theoretically, it is possible to bring the position and velocity *exactly* to zero in finite time, rather than only asymptotically as with a continuous-time controller. This strong type of stabilization is often called *deadbeat control*; its possibility (together with ease of implementation and maintenance, and reliability) constitutes a major advantage of digital techniques.

1.4 Feedback Versus Precomputed Control

Note that the first solution that we provided to the pendulum control problem was in the form of a *feedback* law (1.6), where $u(t)$ could be calculated in terms of the current position and velocity, which are "fed back" after suitable weighings. This is in contrast to "open-loop" design, where the expression of the entire control function $u(\cdot)$ is given in terms of the initial conditions $\varphi(0), \dot\varphi(0)$, and one applies this function $u(\cdot)$ blindly thereafter, with no further observation of positions and velocities. In real systems there will be random perturbations that are not accounted for in the mathematical model. While a feedback law will tend to correct automatically for these, a precomputed control takes no account of them. This can be illustrated by the following simple examples.

Assume that we are only interested in the problem of controlling (1.2) when starting from the initial position $\varphi(0) = 1$ and velocity $\dot\varphi(0) = -2$. Some trial and error gives us that the control function

$$u(t) = 3e^{-2t} \tag{1.25}$$

is adequate for this purpose, since the solution when applying this forcing term is

$$\varphi(t) = e^{-2t}\,.$$

It is certainly true that $\varphi(t)$ and its derivative approach zero, actually rather quickly. So (1.25) solves the original problem. If we made any mistakes in estimating the initial velocity, however, the control (1.25) is no longer very useful:

Exercise 1.4.1 Show that if the differential equation (1.2) is again solved with the right-hand side equal to (1.25) but now using instead the initial conditions

$$\varphi(0) = 1,\ \dot\varphi(0) = -2 + \varepsilon\,,$$

where ε is any positive number (no matter how small), then the solution satisfies

$$\lim_{t\to+\infty} \varphi(t) = +\infty\,.$$

If $\varepsilon < 0$, show that then the limit is $-\infty$. □

A variation of this is as follows. Suppose that we measured correctly the initial conditions but that a momentary power surge affects the motor controlling the pendulum. To model this, we assume that the differential equation is now

$$\ddot{\varphi}(t) - \varphi(t) = u(t) + d(t), \tag{1.26}$$

and that the disturbance $d(\cdot)$ is the function

$$d(t) = \begin{cases} \varepsilon & \text{if } t \in [1,2] \\ 0 & \text{otherwise.} \end{cases}$$

Here ε is some positive real number.

Exercise 1.4.2 Show that the solution of (1.26) with initial conditions $\varphi(0) = 1, \dot{\varphi}(0) = -2$, and u chosen according to (1.25) diverges, but that the solution of

$$\ddot{\varphi}(t) - \varphi(t) = -\alpha\varphi(t) - \beta\dot{\varphi}(t) + d(t) \tag{1.27}$$

still satisfies the condition that φ and $\dot{\varphi}$ approach zero. □

One can prove in fact that the solution of equation (1.27) approaches zero even if $d(\cdot)$ is an arbitrary decaying function; this is an easy consequence of results on the *input/output stability* of linear systems.

Not only is it more robust to errors, but the feedback solution is also in this case simpler than the open-loop one, in that the explicit form of the control as a function of time need not be calculated. Of course, the cost of implementing the feedback controller is that the position and velocity must be continuously monitored.

There are various manners in which to make the advantages of feedback mathematically precise. One may include in the model some characterization of the uncertainty, for instance, by means of specifying a probability law for a disturbance input such as the above $d(\cdot)$. In any case, one can always pose a control problem directly as one of finding feedback solutions, and we shall often do so.

The second solution (1.12)-(1.19) that we gave to the pendulum problem, via digital control, is in a sense a combination of feedback and precomputed control. But in terms of the sampled model (1.18), which ignores the behavior of the system in between sampling times, digital control can be thought of as a purely feedback law. For the times of interest, (1.19) expresses the control in terms of the "current" state variables.

1.5 State-Space and Spectrum Assignment

We now have seen two fundamentally different ways in which to control the linearized inverted pendulum (1.2). They both involve calculating certain gains: α and β in the case of the continuously acting PD controller (1.6), or the entries f_1 and f_2 of F for sampled feedback. In both cases we found that appropriate choices can be made of these coefficients that result, for instance, in decaying exponential behavior if (1.9) holds, or deadbeat control if (1.24) is used.

That such choices are possible is no coincidence. *For very general linear systems, it is possible to obtain essentially arbitrary asymptotic behavior.* One of the basic results in control theory is the *Pole-Shifting Theorem*, also called the "Pole-Assignment" or the "Spectrum Assignment" Theorem, which makes this fact precise.

In order to discuss this Theorem, it is convenient to use the *state-space* formalism. In essence, this means that, instead of studying a high-order differential equation such as (1.1) or (1.2), we replace it by a system of first-order differential equations. For instance, instead of (1.2) we introduce the first-order vector equation

$$\dot{x}(t) = Ax(t) + Bu(t)\,, \tag{1.28}$$

where $x(t)$ is the column vector $(\varphi(t), \dot{\varphi}(t))'$ (prime indicates transpose), and where A and B are the matrices

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\,,\ B = \begin{pmatrix} 0 \\ 1 \end{pmatrix}\,. \tag{1.29}$$

An equation such as (1.28) is an example of what is called a *linear, continuous-time, time-invariant, finite dimensional control system*.

The term "continuous-time" refers to the fact that the time variable t is continuous, in contrast to a system defined by a difference equation such as (1.18). The terminology "time-invariant" is used to indicate that the time t does not appear as an independent variable in the right-hand side of the equation. If the mass of the pendulum were to change in time, for instance, one might model this instead through a time-*varying* equation. "Finite dimensional" refers to the fact that the state $x(t)$ can be characterized completely at each instant by a finite number of parameters (in the above case, two). For an example of a system that is not finite dimensional, consider the problem of controlling the temperature of an object by heating its boundary; the description in that case incorporates a partial differential equation, and the characterization of the state of the system at any given time requires an infinite amount of data (the temperatures at all points).

If instead of the linearized model we had started with the nonlinear differential equation (1.1), then in terms of the same vector x we would obtain a set of first-order equations

$$\dot{x}(t) = f(x(t), u(t))\,, \tag{1.30}$$

where now the right-hand side f is not necessarily a linear function of x and u. This is an example of what we will later call a nonlinear continuous-time system.

Of course, we are interested in systems described by more than just two state variables; in general $x(t)$ will be an n-dimensional vector. Moreover, often there is more than one independent control acting on a system. For instance, in a three-link robotic arm there may be one motor acting at each of its joints (the "shoulder," "elbow," and "wrist"). In such cases the control $u(t)$ at each instant is not a scalar quantity but an m-dimensional vector that indicates the magnitude of the external forces being applied at each input channel. Mathematically, one just lets f in (1.30) be a function of $n+m$ variables. For linear systems (1.28) or (1.18), one allows A to be an $n \times n$ matrix and B to be a rectangular matrix of size $n \times m$.

The transformation from a high-order single equation to a system of first-order equations is exactly the same as that often carried out when establishing existence results for, or studying numerical methods for solving, ordinary differential equations. For any σ, the vector $x(\sigma)$ contains precisely the amount of data needed in order to solve the equation for $\tau > \sigma$, assuming that the control $u(\cdot)$ is specified for $\tau > \sigma$. The *state* $x(\sigma)$ thus summarizes all the information about the past of the system needed in order to understand its future behavior, except for the purely external effects due to the input.

Later we reverse things and *define* systems through the concept of state. We will think of (1.28) not as a differential equation but rather, for each pair of times $\tau > \sigma$, as a rule for obtaining any value $x(\tau)$ from the specifications of $x(\sigma)$ and of the restriction ω of the control $u(\cdot)$ to the interval between σ and τ. We will use the notation

$$x(\tau) = \phi(\tau, \sigma, x, \omega) \tag{1.31}$$

for this rule. We may read the right-hand side as "the state at time τ resulting from starting at time σ in state x and applying the input function ω." Because solutions of differential equations do not exist in general for all pairs of initial and final times $\sigma < \tau$, typically ϕ may be defined only on a subset of the set of possible quadruples $(\tau, \sigma, x, \omega)$.

An advantage of doing things at this level of abstraction is that many objects other than those described by ordinary differential equations also are representable by a rule such as (1.31). For example, discrete-time equations (1.18) also give rise to systems in this sense, since x_{k+i} depends only on

$$x_k, v_k, v_{k+1}, \ldots, v_{k+i-1}$$

for all k and i. The idea of using an abstract transition mapping ϕ originates with the mathematical theory of dynamical systems. The difference with that theory is that here we are interested centrally in understanding the effect of controls.

Returning to the issue of feedback, note that using the matrix formalism (1.28) and denoting by F the vector of gains

$$F = (\,-\alpha \quad -\beta\,)$$

we may write the PD feedback law (1.6) as

$$u(t) = Fx(t)\,. \tag{1.32}$$

When we substitute this into (1.28), we obtain precisely the closed-loop equations (1.7). So the characteristic equation (1.8) of the closed-loop system is the determinant of the matrix

$$\begin{pmatrix} z & -1 \\ \alpha - 1 & z + \beta \end{pmatrix},$$

that is, the characteristic polynomial of $A + BF$. The zeros of the characteristic equation are the eigenvalues of $A + BF$.

We conclude from the above discussion that the problems of finding gains α, β for the PD controller with the properties that all solutions of (1.7) converge to zero, or that all solutions converge with no oscillation, are particular cases of the problem of finding a row vector F so that $A + BF$ has eigenvalues with certain desired properties. This is precisely the same problem that we solved for the discrete-time system (1.18) in order to obtain deadbeat control, where we needed $A + BF$ to have a double eigenvalue at zero. Note that in the second case the matrices A and B are different from the corresponding ones for the first problem. Note also that when dealing with vector-input systems, for which $u(t)$ is not a scalar, the F in equation (1.32) must be a matrix instead of a row vector.

Thus one arrives at the following purely algebraic question suggested by both the PD and the deadbeat digital control problems:

> *Given a pair of real matrices (A, B), where A is square of size $n \times n$ and B is of size $n \times m$, characterize the set of all possible eigenvalues of $A + BF$ as F ranges over the set of all possible real matrices of size $m \times n$.*

The Spectrum Assignment Theorem says in essence that for almost any pair of matrices A and B it is possible to obtain *arbitrary* eigenvalues for $A + BF$ using suitable feedback laws F, subject only to the obvious constraint (since $A + BF$ is a real matrix) that complex eigenvalues must appear in conjugate pairs. "Almost any" means that this will be true for all pairs that describe what we will call *controllable* systems, those systems for which it is possible to steer any state to any other state. We will see later that controllability corresponds to a simple nondegeneracy condition relating A and B; when there is just one control ($m = 1$) the condition is simply that B, seen as a column vector, must be cyclic for A.

This Theorem is most often referred to as the *Pole-Shifting Theorem*, a terminology that is due to the fact that the eigenvalues of $A+BF$ are also the poles of the resolvent operator

$$(zI - A - BF)^{-1},$$

or equivalently, of the rational function $1/\det(zI - A - BF)$; this rational function appears often in classical control design.

As a consequence of the Pole-Shifting Theorem we know that, save for very exceptional cases, given any system

$$\dot{x}(t) = Ax(t) + Bu(t)$$

it will be possible to find a feedback law $u(t) = Fx(t)$ such that all solutions of the closed-loop equation

$$\dot{x}(t) = (A + BF)x(t)$$

decay exponentially. To conclude this, one uses the Pole-Shifting Theorem to obtain a matrix F so that all eigenvalues of $A+BF$ are negative real numbers. Another illustration is the general deadbeat control problem: Given any discrete-time system (1.18), one can in general find an F so that $(A+BF)^n = 0$, simply by assigning zero eigenvalues. Again, a nondegeneracy restriction will be needed, as illustrated by the counterexample $A = I, B = 0$ for which no such F exists.

The Pole-Shifting Theorem is central to linear systems theory and is itself the starting point for more interesting analysis. Once we know that arbitrary sets of eigenvalues can be assigned, it becomes of interest to compare the performance of different such sets. Among those that provide stability, some may be more useful than others because they have desirable transient characteristics such as lack of "overshoot." Also, one may ask what happens when certain entries of F are restricted to vanish, which corresponds to constraints on what can be implemented.

Yet another possibility is to consider a cost criterion to be optimized, as we do when we discuss optimal control theory. For example, it is possible to make trajectories approach the origin arbitrarily fast under the PD controller (1.6), but this requires large gains α, β, which means that the controls $u(t)$ to be applied will have large magnitudes. Thus, there is a trade-off between the cost of controlling and the speed at which this may be achieved.

In a different direction, one may investigate to what extent the dynamics of nonlinear and/or infinite dimensional systems are modifiable by feedback, and this constitutes one of the main areas of current research in control theory.

Classical Design

The next exercise on elementary ordinary differential equations is intended to convey the flavor of some of the simpler performance questions that might be asked when assigning eigenvalues and to mention some standard terminology. Techniques from the theory of functions of a complex variable are often used in this type of study, which is typically the subject of undergraduate engineering courses in the "classical design" or "frequency design" of control systems. We include this discussion as an illustration but will not treat frequency design in this text.

Exercise 1.5.1 Consider again the pendulum linearized about its unstable upper position, given by the equation $\ddot{\varphi}(t) - \varphi(t) = u(t)$, and assume that we use the PD control law $u(t) = -\alpha\varphi(t) - \beta\dot{\varphi}(t)$ to obtain an asymptotically stable closed-loop system $\ddot{\varphi}(t) + b\dot{\varphi}(t) + a\varphi(t) = 0$ (with $a = \alpha - 1 > 0$ and $b = \beta > 0$). Introduce the *natural frequency* $\omega := \sqrt{a}$ and the *damping factor* $\zeta := b/(2\sqrt{a})$, so that the equation now reads

$$\ddot{\varphi}(t) + 2\zeta\omega\dot{\varphi}(t) + \omega^2\varphi(t) = 0\,. \tag{1.33}$$

(A) Prove the following facts:

1. If $\zeta < 1$ (the "underdamped" case), all solutions are decaying oscillations.
2. If $\zeta = 1$ (the "critically damped" case) or if $\zeta > 1$ ("overdamped"), then all solutions for which $\varphi(0) \neq 0$ are such that $\varphi(t) = 0$ for at most one $t > 0$.
3. If $\zeta \geq 1$, then every solution that starts from rest at a displaced position, that is, $\varphi(0) \neq 0$ and $\dot{\varphi}(0) = 0$, approaches zero monotonically. (In this case, we say that there is no "overshoot.")
4. Show rough plots of typical solutions under the three cases $\zeta < 1$, $\zeta = 1$, and $\zeta > 1$.

(B) Consider again the underdamped case. Using rough plots of the solutions of the equation with $\varphi(0) = -1, \dot{\varphi}(0) = 0$ both in the case when ζ is near zero and in the case when ζ is near 1, illustrate the fact that there is a trade-off between the speed of response (how fast one gets to, and stays near, the desired value 0 for φ) and the proportion of overshoot (the maximum of $\varphi(t)$ given the initial magnitude 1). Show that for $\zeta = \sqrt{2}/2$ this overshoot is $e^{-\pi} < 0.05$.

(C) Assume now that the objective is not to stabilize the system (1.2) about $\varphi = \dot{\varphi} = 0$, but to make it assume some other desired value $\varphi = \varphi_d$, still with velocity $\dot{\varphi}$ approaching zero. If we were already at rest at this position $\varphi = \varphi_d$, then the constant control $u \equiv -\varphi_d$ would keep us there. Otherwise, we must add to this a correcting term. We then modify the PD controller to measure instead the deviation from the desired value φ_d:

$$u(t) = -\alpha(\varphi(t) - \varphi_d) - \beta\dot{\varphi}(t) - \varphi_d\,.$$

Show that with this control, and still assuming $\alpha > 1, \beta > 0$, the solutions of the closed-loop system indeed approach $\varphi = \varphi_d, \dot{\varphi} = 0$. One says that the solution has been made to *track* the desired *reference* value φ_d. □

The value $\zeta = \sqrt{2}/2$ appears often in engineering practice, because this results in an overshoot of less than 5%, a figure considered acceptable. If we write θ for the angle between the negative real axis and the vector z, where z is either root of the characteristic equation

$$z^2 + 2\zeta\omega z + \omega^2 = 0\,,$$

then $\zeta = \cos\theta$. Thus, a value $\theta = 45°$ or less is often sought. If in addition one wants the magnitude of solutions to decay at a given exponential rate, say, as $e^{-\lambda t}$, then one looks at placing eigenvalues in a region that forces the desired stability degree, such as that illustrated in Figure 1.4, which exhibits a "damping margin" of 45° and a "stability margin" of λ.

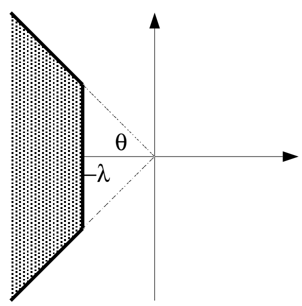

Figure 1.4: *Typical desired eigenvalue locations.*

Because solutions of more general linear equations can be written as sums of those corresponding to complex pairs and real roots, designers often apply these same rules to systems of order higher than two.

1.6 Outputs and Dynamic Feedback

We introduce now another major component of control theory, the idea of measurements and outputs.

In some situations all of the variables needed in order to implement a feedback control are readily available, but other times this is not the case. For instance, the speed of a car can be obtained easily via a speedometer, but the ground speed of an aircraft may be difficult to estimate. To model mathematically such a constraint, one adds to the description of a system the specification of a *measurement* function. This is a function of the states, with scalar or vector values, which summarizes all of the information that is available to the control algorithm.

Returning once more to the inverted pendulum example, let us assume that in designing the PD controller (1.6) we are restricted to measuring directly only the angular position $\varphi(t)$. We model this by adding to the state space equation

(1.28), where A and B are given by (1.29), the mapping

$$C : \mathbb{R}^2 \to \mathbb{R} : \ x \mapsto x_1 \,, \tag{1.34}$$

or equivalently its matrix in the canonical basis, $C = (1, 0)$, that picks the variable φ. We write $y(t)$ for the allowed measurement at time t, that is,

$$y(t) = Cx(t) \,. \tag{1.35}$$

More generally, we define a *linear system with output* by giving triples of matrices (A, B, C), where for some integers n, m, p, A is $n \times n$, B is $n \times m$, and C is $p \times n$. The integer p can be thought of as the number of independent measurements that can be taken at each instant. One also defines nonlinear and discrete-time systems with output in an analogous way, or more abstractly by adding an output map to the axiomatic definition of system via transitions (1.31).

A linear feedback law that depends only on the allowed measurements is then, in terms of this notation, of the form

$$u(t) = KCx(t) \,. \tag{1.36}$$

In other words, F is restricted to factor as $F = KC$, where C is given. As discussed when dealing with proportional-only control, such *static output feedback* is in general inadequate for control purposes. The set of eigenvalues achievable for matrices $A + BKC$ by varying K is severely restricted.

Of course, there is no theoretical reason for restricting attention to static feedback laws. A controller that would incorporate a device for differentiating the observed position $x_1(t) = \varphi(t)$ could then use this derivative.

However, differentiation tends to be an undesirable operation, because it is very sensitive to noise. To illustrate this sensitivity, consider the problem of finding the derivative c of the function

$$\xi(t) = ct \,,$$

where c is an unknown constant. Assume that the data available consists of $\xi(t)$ measured under additive noise,

$$y(t) = \xi(t) + n(t) \,,$$

where $n(t)$ is of the form

$$n(t) = d \sin \omega t \,,$$

and d and ω are unknown constants, ω being large. (This is realistic in that noise effects tend to be of "high frequency" compared to signals, but of course in real applications the noise cannot be modeled by a deterministic signal of constant frequency, and a probabilistic description of $n(t)$ is needed for an accurate analysis.) If we simply differentiate, the result

$$\dot{y}(t) = c + d\omega \cos \omega t$$

can be very far from the desired value c since ω is large. An alternative is to use the fact that the noise $n(t)$ averages to zero over long intervals; thus, we can cancel its effect by integration. That is, if we compute

$$\int_0^t y(\tau)\, d\tau$$

and we then multiply this by $2/t^2$, we obtain

$$c - \frac{2d}{\omega t^2}(1 - \cos \omega t)\,, \tag{1.37}$$

which converges to c as $t \to \infty$ and hence provides an accurate estimate, asymptotically, of the desired derivative.

The topic of *observers* deals with obtaining estimates of all the state variables from the measured ones, using only integration —or, in discrete-time, summations. It is an easy consequence of general theorems on the existence of observers that one can achieve also stabilization of the inverted pendulum, as well as of most other continuous-time linear systems, using a controller that is itself a continuous-time linear system. Such a controller is said to incorporate *integral* or *dynamic* feedback, and it includes a differential equation, driven by the observations, that calculates the necessary estimate. Analogously, for discrete-time systems with partial observations, the theorems result in the existence of observers that are themselves discrete-time linear systems.

In particular, there exist for the pendulum coefficients α, β, μ, ν, so that

$$\dot{z}(t) = \nu z(t) + \mu x_1(t) \tag{1.38}$$

together with a feedback law

$$u(t) = -\alpha x_1(t) - \beta z(t) \tag{1.39}$$

stabilizes the system. The controller solves the differential equation (1.38) with an arbitrary initial condition $z(0)$ and feeds the linear combination (1.39) of the measured variable x_1 and the estimated variable z back into the system. Later we develop in detail a systematic approach to the construction of such dynamic controllers, but in this simple example it is easy to find the necessary parameters by analyzing the equations.

Exercise 1.6.1 Consider the system of three differential equations obtained from (1.28) together with (1.38), where $u(t)$ is given by (1.39). Find numbers α, β, μ, ν, such that all solutions of this system approach zero. □

When observations are subject to noise, the speed at which one can estimate the unmeasured variables —and hence how fast one can control— is limited by the magnitude of this noise. For instance, the convergence of the expression in equation (1.37) is slower when d is large (noise has high intensity) or ω is small (noise has large bandwidth). The resulting trade-offs give rise to problems of *stochastic filtering* or *stochastic state estimation*, and the *Kalman Filter* is a technique for the analysis of such problems.

PID Control

Dynamic control is also useful in order to cancel unwanted disturbances even if all state variables are available for feedback. We illustrate this using again the inverted pendulum model. Assume that an unknown but constant disturbance $d(t) \equiv e$ can act on the system, which is therefore described by

$$\ddot{\varphi}(t) - \varphi(t) = u(t) + e. \tag{1.40}$$

One might still try to use the PD controller (1.6) to stabilize this system independently of e. Since stability should hold in particular when $e = 0$, one should have $\alpha > 1$ and $\beta > 0$ as before. But the PD controller is not sufficient:

Exercise 1.6.2 Show that if $e \neq 0$ and $\alpha > 1$, $\beta > 0$, then no solution of

$$\ddot{\varphi}(t) + \beta\dot{\varphi}(t) + (\alpha - 1)\varphi(t) = e$$

is such that $\varphi(t) \to 0$ as $t \to \infty$. □

Since e is a constant, it can be thought of as the solution of the differential equation

$$\dot{x}_0(t) = 0\,. \tag{1.41}$$

If we think of x_0 as a quantity that cannot be measured, then the previous discussion about observers suggests that there may exist a dynamic controller:

Exercise 1.6.3 Find numbers α, β, μ with the following property: For each $e \in \mathbb{R}$, all of the solutions of the system of equations

$$\begin{aligned}
\dot{x}_0(t) &= x_1(t)\\
\dot{x}_1(t) &= x_2(t)\\
\dot{x}_2(t) &= -\mu x_0(t) + (1-\alpha)x_1(t) - \beta x_2(t) + e
\end{aligned}$$

converge to $(e/\mu, 0, 0)$. (*Hint:* First solve this problem for the homogeneous system that results if $e = 0$ by finding parameters that make the matrix associated to the equations have all eigenvalues with negative real parts. Then show that the same parameters work in the general case.) □

Thus, using a controller that integrates $x_1(t) = \varphi(t)$ and feeds back the combination

$$u(t) = -\alpha\varphi(t) - \beta\dot{\varphi}(t) - \mu\int_0^t \varphi(\tau)d\tau \tag{1.42}$$

one can ensure that both φ and $\dot{\varphi}$ converge to zero. Another way of thinking of this controller is as follows: If the integral term is not used, the value $\varphi(t)$ approaches a constant *steady-state error*; the effect of the integral is to offset a nonzero error. The controller (1.42) is a *PID*, or *proportional-integral-derivative* feedback, the control mechanism most widely used in linear systems applications.

More generally, the problem of canceling the effect of disturbances that are not necessarily constant gives rise to the study of *disturbance rejection* or *disturbance decoupling* problems. The resulting dynamic controllers typically incorporate an “internal model” of the possible mechanism generating such disturbances, as with the simple case (1.41).

Observability, Duality, Realization, and Identification

A basic system theoretic property that arises naturally when studying observers and dynamic feedback is that of *observability*. Roughly speaking, this means that all information about the state x should in principle be recoverable from knowledge of the measurements y that result; in a sense it is a counterpart to the concept of controllability —mentioned earlier— which has to do with finding controls u that allow one to attain a desired state x. In the case of linear systems, this analogy can be made precise through the idea of *duality*, and it permits obtaining most results for observers as immediate corollaries of those for controllability. This duality extends to a precise correspondence between optimal control problems —in which one studies the trade-off between cost of control and speed of control— and filtering problems —which are based on the trade-off between magnitude of noise and speed of estimation.

Once outputs or measurements are incorporated into the basic definition of a system, one can pose questions that involve the relationship between controls u and observations y. It is then of interest to characterize the class of *input/output (i/o) behaviors* that can be obtained and conversely to study the *realization* problem: “Given an observed i/o behavior, what possible systems could give rise to it?” In other words, if we start with a “black box” model that provides only information about how u affects y, how does one deduce the differential — or, in discrete-time, difference— equation that is responsible for this behavior?

Besides obvious philosophical interest, such questions are closely related to *identification* problems, where one wishes to estimate the i/o behavior itself from partial and possibly noisy data, because state-space descriptions serve to parametrize such possible i/o behaviors. Conversely, from a synthesis viewpoint, realization techniques allow us to compute a state representation, and hence also construct a physical system, that satisfies given i/o specifications.

The main results on the realization problem show that realizations essentially are unique provided that they satisfy certain minimality or irredundancy requirements. We will provide the main theorems of realization theory for linear systems. The underlying properties turn out to be closely related to other system theoretic notions such as controllability and observability.

1.7 Dealing with Nonlinearity

Ultimately, linear techniques are limited by the fact that real systems are more often than not nonlinear. As discussed above, local analysis and control design in general can be carried out satisfactorily using just first-order information, but global analysis requires more powerful methods.

One case in which global control *can* be achieved with ideas from linear system theory is that in which the system can be reduced, after some transformation, to a linear one. For instance, a change of coordinates in the space of the x variable might result in such a simplification. A situation that has appeared many times in applications is that in which a particular *feedback* control linearizes the system. Consider again the pendulum, but now the original nonlinear model.

If we first subtract the effect of the term $\sin\theta$ by using u, we are left with a simple linear system

$$\ddot{\theta}(t) = u(t)\,, \tag{1.43}$$

which can be controlled easily. For instance, the PD feedback law

$$u(t) = -\theta(t) - \dot{\theta}(t) \tag{1.44}$$

stabilizes (1.43). In order to stabilize the original system, we now add the subtracted term. That is, we use

$$u(t) = \sin\theta(t) - \theta(t) - \dot{\theta}(t)\,. \tag{1.45}$$

Passing from the original model to (1.43) can be thought of as the effect of applying the feedback law

$$u(t) = f(x(t)) + u'(t)\,,$$

where $f(x) = \sin x_1$ and u' is a new control; the study of such feedback transformations and the characterization of conditions under which they simplify the system is an active area of research.

The above example illustrates another issue that is important when dealing with the global analysis of nonlinear control problems. This is the fact that the mathematical structure of the state space might impose severe restrictions on what is achievable. The proposed feedback law (1.45) will stabilize

$$\ddot{\theta}(t) + \sin\theta(t) = u(t)$$

provided we can think of θ as taking arbitrary real values. However, in the physical system $\theta + 2\pi$ describes the same state as θ. Thus, the natural space to describe θ is a unit circle. When applying the control u as defined by the formula (1.45), which of the infinitely many possible values of θ should one use? It turns out that everything will behave as desired as long as one chooses u continuously in time —as the pendulum finishes one complete counterclockwise revolution, start using an angle measured in the next 2π-interval. But this choice is not unambiguous in terms of the physical coordinate θ. In other words, (1.45) is not a feedback law when the "correct" state space is used, since it is not a function of states.

The correct state space for this example is in fact the Cartesian product of a circle and the real line, since pairs $(\theta(t), \dot{\theta}(t))$ belong naturally to such a space. It can be proved that in such a space —mathematically the tangent bundle to the unit circle— there cannot exist *any* continuous feedback law that stabilizes the system, because the latter would imply that the space is diffeomorphic (smoothly equivalent) to a Euclidean space. In general, techniques from differential geometry —or, for systems described by polynomial nonlinearities, algebraic geometry— appear when studying global behavior.

When dealing with the problems posed by infinite dimensionality as opposed to nonlinearity, similar considerations force the use of techniques from functional analysis.

1.8 A Brief Historical Background

Control mechanisms are widespread in nature and are used by living organisms in order to maintain essential variables such as body temperature and blood sugar levels at desired setpoints. In engineering too, feedback control has a long history: As far back as the early Romans, one finds water levels in aqueducts being kept constant through the use of various combinations of valves.

Modern developments started during the seventeenth century. The Dutch mathematician and astronomer Christiaan Huygens designed pendulum clocks and in doing so analyzed the problem of speed control; in this work he competed with his contemporary Robert Hooke. The needs of navigation had a strong influence on scientific and technological research at that time, and accurate clocks —to allow determinations of solar time— were in great demand. The attention turned to windmills during the eighteenth century, and speed controls based on Hooke's and Huygens' ideas were built. A central idea here is the use of flyballs: Two balls attached to an axis turn together with the windmill, in such a manner that centrifugal force due to angular velocity causes them to rise; in turn this upward movement is made to affect the positions of the mill's sails. Thus, feedback was implemented by the linkages from the flyballs to the sails.

But it was the Industrial Revolution, and the adaptation by James Watt in 1769 of flyball governors to steam engines, that made control mechanisms very popular; the problem there was to regulate engine speed despite a variable load. Steady-state error could appear, and various inventors introduced variations of the integral feedback idea in order to deal with this problem.

The mathematician and astronomer George Airy was the first to attempt, around 1840, an analysis of the governor and similar devices. By 1868, there were about 75,000 Watt governors in use; that year, the Scottish physicist James Clerk Maxwell published the first complete mathematical treatment of their properties and was able to explain the sometimes erratic behavior that had been observed as well as the effect of integral control. His work gave rise to the first wave of theoretical research in control, and characterizations of stability were independently obtained for linear systems by the mathematicians A. Hurwitz and E.J. Routh. This theory was applied in a variety of different areas, such as the study of ship steering systems.

During the 1930s, researchers at Bell Telephone Laboratories developed the theory of feedback amplifiers, based on assuring stability and appropriate response for electrical circuits. This work, by H. Nyquist, H.W. Bode, and others, constitutes even today the foundation of much of frequency design. Analog computers appeared also around that time and formed the basis for implementing controllers in the chemical and petroleum industries. During the Second World War, these techniques were used in the design of anti-aircraft batteries and fire-control systems, applications that continue to be developed today. The mathematician Norbert Wiener, who developed a theory of estimation for random processes used in these applications, coined at the end of the war the term "cybernetics" to refer to control theory and related areas.

These so-called classical approaches were for the most part limited by their restriction to linear time-invariant systems with scalar inputs and outputs. Only during the 1950s did control theory begin to develop powerful general techniques that allowed treating multivariable, time-varying systems, as well as many nonlinear problems. Contributions by Richard Bellman (dynamic programming) and Rudolf Kalman (filtering, linear/quadratic optimal control, and algebraic analysis) in the United States, and by L. Pontryagin (nonlinear optimal control) in the Soviet Union, formed the basis for a very large research effort during the 1960s, which continues to this day. Present day theoretical research in control theory involves a variety of areas of pure mathematics, as illustrated for instance by the remarks in Section 1.7. Concepts and results from these areas find applications in control theory; conversely, questions about control systems give rise to new open problems in mathematics.

Excellent references for the early historical background are the papers [149], which contain a large number of literature citations, and the book [299]. See also the introductory article [226]. Other references, which in addition contain overviews of current research and open problems, are the reports [275] and [135].

1.9 Some Topics Not Covered

In an area as wide as control theory, it is impossible to cover all topics in a single text, even briefly. For instance, we concentrate exclusively on *deterministic* systems. Incorporating models for uncertainty leads to *stochastic* models, which are the subject of much research activity, when this uncertainty is expressed in probabilistic or statistical terms; references to the literature on stochastic aspects of many of the problems considered are given at various points in the text. Mathematically different but closely related is the area of *robust control*, which deals with the design of control laws that are guaranteed to perform even if the assumed model of the system to be controlled is incorrect —with the allowed deviations quantified in appropriate norms— or under the possibility of imperfect controller design. See the collection of papers [123] for pointers to a large literature. The area of *adaptive control* deals also with the control of partially unknown systems, but differs from robust control in the mathematics employed. Adaptive controllers typically make use of *identification techniques*, which produce estimates of the system parameters for use by controllers; see e.g. [282].

When using computers one should consider the effect of quantization errors on the implementation of control laws, which arise due to limited precision when real-valued signals are translated into fixed-point representations (*A/D* or *analog to digital* conversion); see e.g. [304]. Other questions relate to the interface between higher-level controllers implemented in software and lower-level servos of an analog and physical character; this gives rise to the area of *hybrid systems*; see the volume [15].

The notes at the end of the various chapters give further bibliographical references to many of the other areas that have been omitted.