1 An Overview and Brief History of Feedback Control

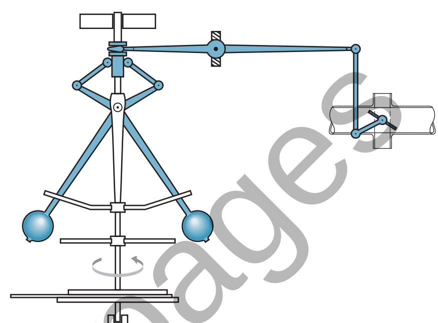

A Perspective on Feedback Control

Feedback control of dynamic systems is a very old concept with many characteristics that have evolved over time. The central idea is that a dynamic system's output can be measured and fed back to a controller of some kind then used to affect the system. There are several variations on this theme.

A system that involves a person controlling a machine, as in driving an automobile, is called **manual** control. A system that involves machines only, as when room temperature can be set by a thermostat, is called **automatic** control. Systems designed to hold an output steady against unknown disturbances are called **regulators**, while systems designed to track a reference signal are called **tracking** or **servo** systems. Control systems are also classified according to the information used to compute the controlling action. If the controller does *not* use a measure of the system output being controlled in computing the control action to take, the system is called **open-loop** control. If the controlled output signal *is* measured and fed back for use in the control computation, the system is called **closed-loop** or **feedback** control. There are many other important properties of control systems in addition to these most basic characteristics. For example, we will mainly consider feedback of current measurements as opposed to predictions of the future; however, a very familiar example illustrates the limitation imposed by that assumption. When driving a car, the use of simple feedback corresponds to driving in a thick fog where one can *only see the road immediately at the front of the car* and is unable to see the future required position! Looking at the road ahead is a form of predictive control and this information, which has obvious advantages, would always be used where it is available. In most automatic control situations studied in this book, observation of the future track or disturbance is not possible. In any case, the control designer should study the process to see if any information could anticipate either a track to be followed or a disturbance to be rejected. If such a possibility is feasible, the control designer should use it to **feedforward** an early warning to the control system. An example of this is in the control of steam pressure in the boiler of an electric power generation plant. The electricity demand cycle over a day is well known; therefore, when it is known that there will soon be an increased need for electrical power, that information can be fed forward to the boiler controller in anticipation of a soon-to-be-demanded increase in steam flow.

The applications of feedback control have never been more exciting than they are today. Feedback control is an essential element in aircraft of all types: most manned aircraft, and all unmanned aircraft from large military aircraft to small drones. The FAA has predicted that the number of drones registered in the U.S. will reach 7 million by 2020! Automatic landing and collision avoidance systems in airliners are now being used routinely, and the use of satellite navigation in future designs promises a revolution in our ability to navigate aircraft in an ever more crowded airspace. The use of feedback control in driverless cars is an essential element to their success. They are now under extensive development, and predictions have been made that driverless cars will ultimately reduce the number of cars on the road by a very large percentage. The use of feedback control in surgical robotic systems is also emerging. Control is essential to the operation of systems from cell phones to jumbo jets and from washing machines to oil refineries as large as a small city. The list goes on and on. In fact, many engineers refer to control as a *hidden technology* because of its essential importance to so many devices and systems while being mainly out of sight. The future will no doubt see engineers create even more imaginative applications of feedback control.

Chapter Overview

In this chapter, we begin our exploration of feedback control using a simple familiar example: a household furnace controlled by a thermostat. The generic components of a control system are identified within the context of this example. In another example in Section 1.2—an automobile cruise control—we will develop the elementary static equations and assign numerical values to elements of the system model in order to compare the performance of open-loop control to that of feedback control when dynamics are ignored. Section 1.3 then introduces the key elements in control system design. In order to provide a context for our studies, and to give you a glimpse of how the field has evolved, Section 1.4 provides a brief history of control theory and design. In addition, later chapters have brief sections of additional historical notes on the topics covered there. Finally, Section 1.5 provides a brief overview of the contents and organization of the entire book.

1.1 A Simple Feedback System

In feedback systems, the variable being controlled—such as temperature or speed—is measured by a sensor and the measured information is fed back to the controller to influence the controlled variable. The principle is readily illustrated by a very common system, the household furnace controlled by a thermostat. The components of this system and their interconnections are shown in Fig. 1.1. Such an illustration identifies the major parts of the system and shows the directions of information flow from one component to another.

We can easily analyze the operation of this system qualitatively from the graph. Suppose both the temperature in the room where the thermostat is located and the outside temperature are significantly below the reference temperature (also called the setpoint) when power is applied. The thermostat will be *on* and the control logic will open the furnace gas valve and light the fire box. This will cause heat Q_{in} to be supplied to the house at a rate that will be significantly larger than the heat loss Q_{out}. As a result, the room temperature will rise until it exceeds the thermostat reference setting by a small amount. At this time, the furnace will be turned off and the room temperature will start to fall toward the outside value. When it falls a small amount below the setpoint,[1] the thermostat will come on again and the cycle will repeat. Typical plots of room temperature along with the furnace cycles of on and off are shown in Fig. 1.1. The outside temperature remains at 50°F and the thermostat is initially set at 55°F. At 6 a.m., the thermostat is stepped to 65°F and the furnace brings it to that level and cycles the temperature around that value thereafter. Notice the house is well insulated, so the fall of temperature with the furnace off is significantly slower than the rise with the furnace on. From this example, we can identify the generic components of the elementary feedback control system, as shown in Fig. 1.2.

The central component of this feedback system is the **process** whose output is to be controlled. In our example the process would be the house whose output is the room temperature and the **disturbance** to

[1] The **setpoint**, **reference**, and **desired input** are all the same thing and shown in Figs. 1.1–1.3.

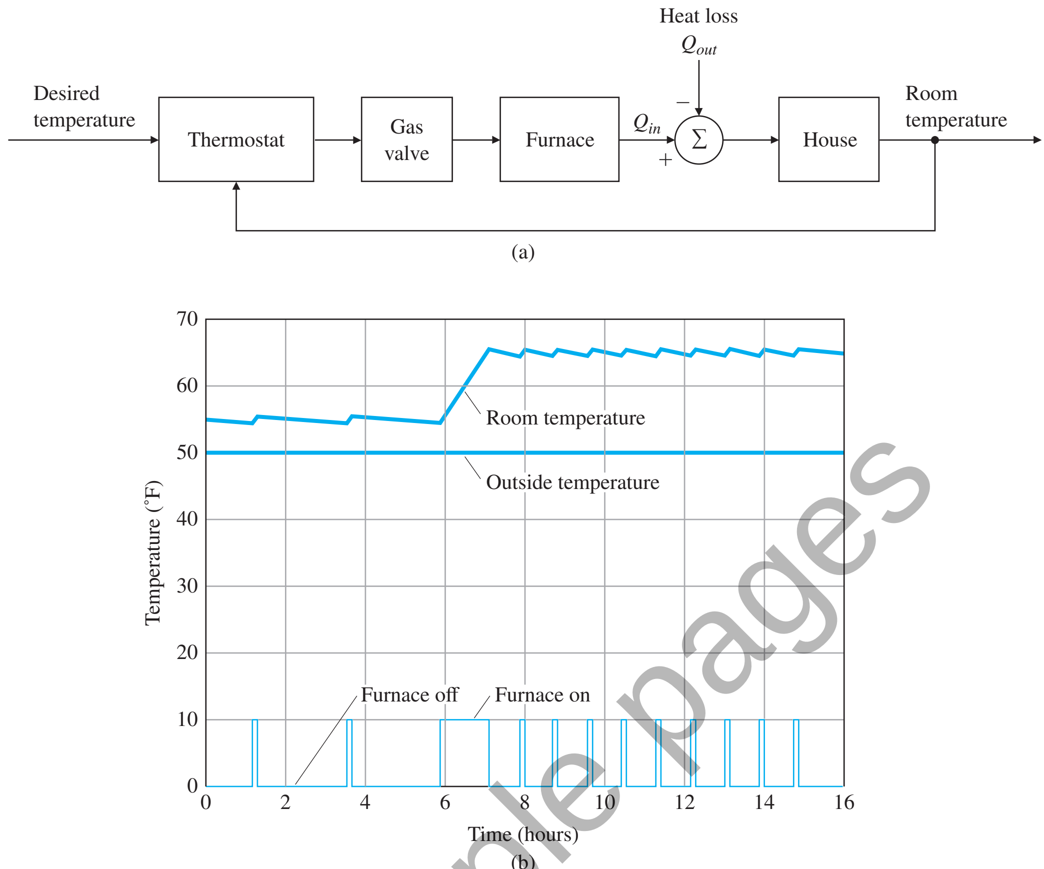

Figure 1.1

Feedback control: (a) component block diagram of a room temperature control system; (b) plot of room temperature and furnace action

the process is the flow of heat from the house, Q_{out}, due to conduction through the walls and roof to the lower outside temperature. (The outward flow of heat also depends on other factors such as wind, open doors, and so on.) The design of the process can obviously have a major impact on the effectiveness of the controls. The temperature of a well-insulated house with thermopane windows is clearly easier to control than otherwise. Similarly, the design of aircraft with control in mind makes a world of difference to the final performance. In every case, the earlier the concepts of control are introduced into the process design, the better. The **actuator** is the device that can influence the controlled variable of the process. In our case, the actuator is a gas furnace. Actually, the furnace usually has a pilot light or striking mechanism, a gas valve, and a blower fan, which turns on or off depending on the air temperature in the furnace. These details illustrate the fact that many feedback systems contain components that themselves

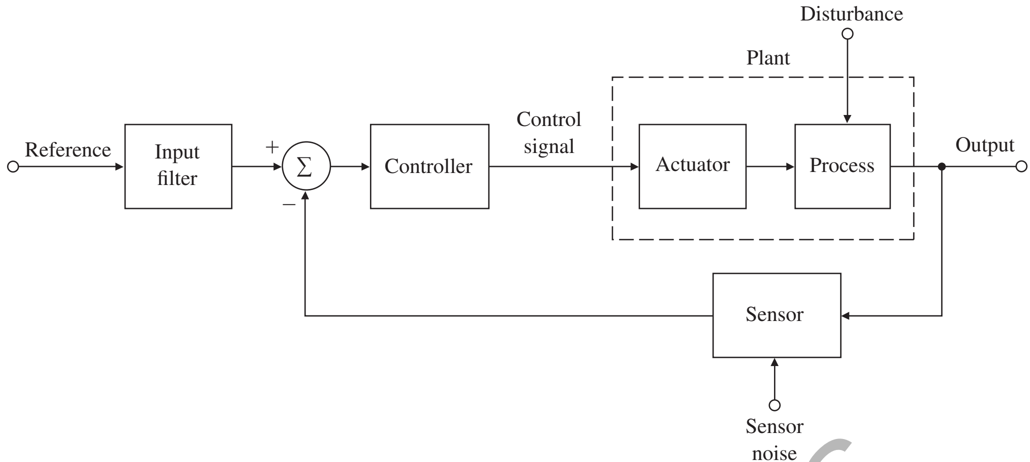

Figure 1.2

Component block diagram of an elementary feedback control

form other feedback systems.[2] The central issue with the actuator is its ability to move the process output with adequate speed and range. The furnace must produce more heat than the house loses on the worst day, and must distribute it quickly if the house temperature is to be kept in a narrow range. Power, speed, and reliability are usually more important than accuracy. Generally, the process and the actuator are intimately connected and the control design centers on finding a suitable input or control signal to send to the actuator. The combination of process and actuator is called the **plant**, and the component that actually computes the desired control signal is the **controller**. Because of the flexibility of electrical signal processing, the controller typically works on electrical signals, although the use of pneumatic controllers based on compressed air has a long and important place in process control. With the development of digital technology, cost-effectiveness and flexibility have led to the use of digital signal processors as the controller in an increasing number of cases. The component labeled **thermostat** in Fig. 1.1 measures the room temperature and is called the **sensor** in Fig. 1.2, a device whose output inevitably contains sensor noise. Sensor selection and placement are very important in control design, for it is sometimes not possible for the true controlled variable and the sensed variable to be the same. For example, although we may really wish to control the house temperature as a whole, the thermostat is in one particular room, which may or may not be at the same temperature as the rest of the house. For instance, if the thermostat is set to 68°F but is placed in the living room near a roaring fireplace, a person working in

[2]Jonathan Swift (1733) said it this way: "So, Naturalists observe, a flea Hath smaller fleas that on him prey; And these have smaller still to bite 'em; And so proceed, *ad infinitum*." Swift, J., On Poetry: A Rhapsody, 1733, J. Bartlett, ed., Familiar Quotations, 15th ed., Boston: Little Brown, 1980.

the study could still feel uncomfortably cold.[3,4] As we will see, in addition to placement, important properties of a sensor are the accuracy of the measurements as well as low noise, reliability, and linearity. The sensor will typically convert the physical variable into an electrical signal for use by the controller. Our general system also includes an **input filter** whose role is to convert the reference signal to electrical form for later manipulation by the controller. In some cases, the input filter can modify the reference command input in ways that improve the system response. Finally, there is a **controller** to compute the difference between the reference signal and the sensor output to give the controller a measure of the system error. The thermostat on the wall includes the sensor, input filter, and the controller. A few decades ago, the user simply set the thermostat manually to achieve the desired room temperature at the thermostat location. Over the last few decades, the addition of a small computer in the thermostat has enabled storing the desired temperature over the day and week and more recently, thermostats have gained the ability to learn what the desired temperature should be and to base that value, in part, on whether anybody will be home soon! A thermostat system that includes a motion detector can determine whether anybody is home and learns from the patterns observed what the desired temperature profile should be. The process of learning the desired setpoint is an example of artificial intelligence (AI) or machine learning, which is gaining acceptance in many fields as the power and affordability of computers improve. The combination of feedback control, AI, sensor fusion, and logic to tie it all together will become an essential feature in many future devices such as drones, driverless cars, and many others.

This text will present methods for analyzing feedback control systems and will describe the most important design techniques engineers can use in applying feedback to solve control problems. We will also study the specific advantages of feedback that compensate for the additional complexity it demands.

1.2 A First Analysis of Feedback

The value of feedback can be readily demonstrated by quantitative analysis of a simplified model of a familiar system, the cruise control of an automobile (see Fig. 1.3). To study this situation analytically, we

[3] In the renovations of the kitchen in the house of one of the authors, the new ovens were placed against the wall where the thermostat was mounted on the other side. Now when dinner is baked in the kitchen on a cold day, the author freezes in his study unless the thermostat is reset.

[4] The story is told of the new employee at the nitroglycerin factory who was to control the temperature of a critical part of the process manually. He was told to "keep that reading below 300°." On a routine inspection tour, the supervisor realized that the batch was dangerously hot and found the worker holding the thermometer under cold water tap to bring it down to 300°. They got out just before the explosion. Moral: sometimes automatic control is better than manual.

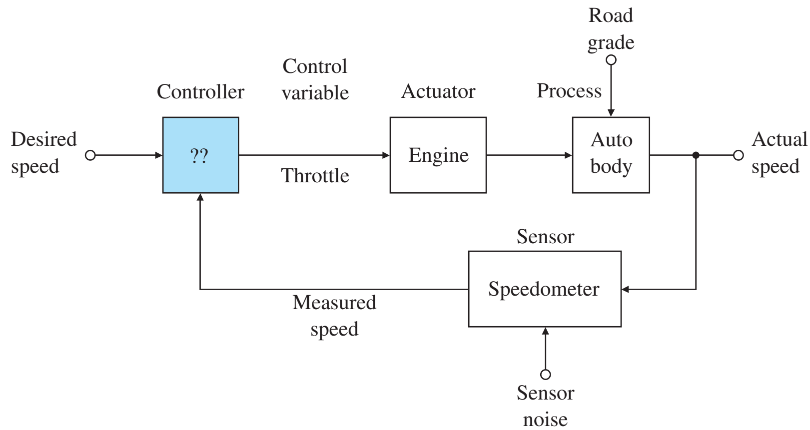

Figure 1.3
Component block diagram of automobile cruise control

need a mathematical **model** of our system in the form of a set of quantitative relationships among the variables. For this example, we ignore the dynamic response of the car and consider only the steady behavior. (Dynamics will, of course, play a major role in later chapters.) Furthermore, we assume that for the range of speeds to be used by the system, we can approximate the relations as linear. After measuring the speed of the vehicle on a level road at 65 mph, we find that a 1° change in the throttle angle (our control variable, u) causes a 10 mph change in speed (the output variable, y), hence the value 10 in the box between u and y in Fig. 1.4, which is a **block diagram** of the plant. Generally, the block diagram shows the mathematical relationships of a system in graphical form. From observations while driving up and down hills, it is found that when the grade changes by 1%, we measure a speed change of 5 mph, hence the value 0.5 in the upper box in Fig. 1.4, which reflects that a 1% grade change has half the effect of a 1° change in the throttle angle. The speedometer is found to be accurate to a fraction of 1 mph and will be considered exact. In the block diagram, the connecting lines carry signals and a block is like an ideal amplifier which multiplies the signal at its input by the value marked in the block to give the output signal. To sum two or more signals, we show lines for the signals coming into a **summer**, a circle with the summation sign Σ inside. An algebraic sign (plus or minus) beside each arrow head indicates whether the input

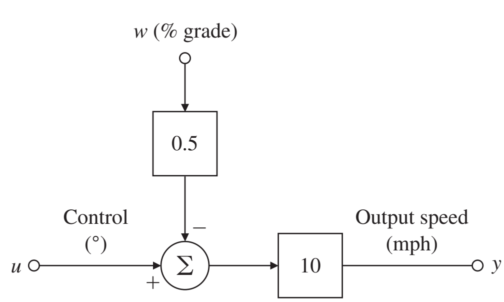

Figure 1.4
Block diagram of the cruise control plant

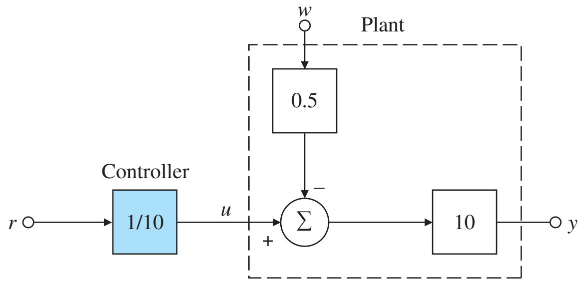

Figure 1.5
Open-loop cruise control

adds to or subtracts from the total output of the summer. For this analysis, we wish to compare the effects of a 1% grade on the output speed when the reference speed is set for 65 with and without feedback to the controller.

Open-loop control

In the first case, shown in Fig. 1.5, the controller does not use the speedometer reading but sets $u = r/10$, where r is the reference speed, which is, 65 mph. This is an example of an **open-loop control system.** The term *open-loop* refers to the fact that there is no closed path or loop around which the signals go in the block diagram; that is, the control variable u is independent of the output variable, y. In our simple example, the open-loop output speed, y_{ol}, is given by the equations

$$\begin{aligned} y_{ol} &= 10(u - 0.5w) \\ &= 10\left(\frac{r}{10} - 0.5w\right) \\ &= r - 5w. \end{aligned}$$

The error in output speed is

$$e_{ol} = r - y_{ol} \tag{1.1}$$

$$= 5w, \tag{1.2}$$

and the percent error is

$$\% \, error = 500\frac{w}{r}. \tag{1.3}$$

If $r = 65$ and the road is level, then $w = 0$ and the speed will be 65 with no error. However, if $w = 1$ corresponding to a 1% grade, then the speed will be 60 and we have a 5-mph error, which is a 7.69% error in the speed. For a grade of 2%, the speed error would be 10 mph, which is an error of 15.38%, and so on. The example shows that there would be no error when $w = 0$, but this result depends on the controller gain being the exact inverse of the plant gain of 10. In practice, the plant gain is subject to change and if it does, errors are introduced by this means also. If there is an error in the plant gain in open-loop control, the percent speed error would be the same as the percent plant-gain error.

The block diagram of a feedback scheme is shown in Fig. 1.6, where the controller gain has been set to 10. In this simple example, we have assumed that we have an ideal sensor providing a measurement of y_{cl}. In this case, the equations are

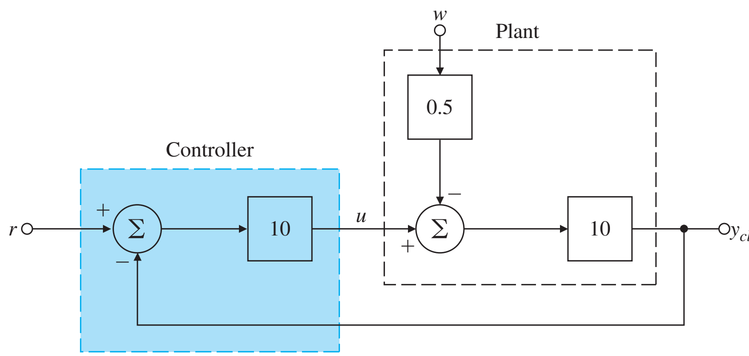

Figure 1.6
Closed-loop cruise control

$$y_{cl} = 10u - 5w,$$
$$u = 10(r - y_{cl}).$$

Combining them yields

$$y_{cl} = 100r - 100y_{cl} - 5w,$$
$$101y_{cl} = 100r - 5w,$$
$$y_{cl} = \frac{100}{101}r - \frac{5}{101}w,$$
$$e_{cl} = \frac{r}{101} + \frac{5w}{101}.$$

Thus, the feedback has reduced the sensitivity of the speed error to the grade by a factor of 101 when compared with the open-loop system. Note, however, that there is now a small speed error on level ground because even when $w = 0$,

$$y_{cl} = \frac{100}{101}r = 0.99r \text{ mph}.$$

This error will be small as long as the loop gain (product of plant and controller gains) is large.[5] If we again consider a reference speed of 65 mph and compare speeds with a 1% grade, the percent error in the output speed is

$$\%\ error = 100\frac{\frac{65\times 100}{101} - \left(\frac{65\times 100}{101} - \frac{5}{101}\right)}{\frac{65\times 100}{101}} \tag{1.4}$$

$$= 100\frac{5\times 101}{101\times 65\times 100} \tag{1.5}$$

$$= 0.0769\%. \tag{1.6}$$

[5] In case the error is too large, it is common practice to *reset* the reference, in this case to $\frac{101}{100}r$, so the output reaches the true desired value.

The reduction of the speed sensitivity to grade disturbances and plant gain in our example is due to the loop gain of 100 in the feedback case. Unfortunately, there are limits to how high this gain can be made; when dynamics are introduced, the feedback can make the response worse than before, or even cause the system to become unstable. The dilemma is illustrated by another familiar situation where it is easy to change a feedback gain. If one tries to raise the gain of a public-address amplifier too much, the sound system will squeal in a most unpleasant way. This is a situation where the gain in the feedback loop—from the speakers to the microphone through the amplifier back to the speakers—is too much. The issue of how to get the gain as large as possible to reduce the errors without making the system become unstable is called the design trade-off and is what much of feedback control design is all about.

The design trade-off

1.3 Feedback System Fundamentals

To achieve good control there are typical goals:

- **Stability**. The system must be stable at all times. This is an absolute requirement.
- **Tracking**. The system output must track the command reference signal as closely as possible.
- **Disturbance rejection**. The system output must be as insensitive as possible to disturbance inputs.
- **Robustness**. The aforementioned goals must be met even if the model used in the design is not completely accurate or if the dynamics of the physical system change over time.

The requirement of **stability** is basic and instability may have two causes. In the first place, the system being controlled may be unstable. This is illustrated by the Segway vehicle, which will simply fall over if the control is turned off. A second cause of instability may be the addition of feedback! Such an instability is called a “vicious circle,” where the feedback signal that is circled back makes the situation worse rather than better. Stability will be discussed in much more detail in Chapters 3 and 4.

There are many examples of the requirement of having the system’s output track a command signal. For example, driving a car so the vehicle stays in its lane is **command tracking**. Today, this is done by the driver; however, there are schemes now under development where the car’s “autodriver” will carry out this task using feedback control while the driver does other things, for example, surfing the Internet. Similarly, flying an airplane on the approach to landing requires that a glide path be accurately tracked by the pilot or an autopilot. It is routine for today’s aircraft autopilots to carry out this including the flare to the actual touchdown. The autopilot accepts inputs from the Instrument Landing System (ILS) that provides an electronic signal showing the desired landing trajectory, then commands the aircraft control surfaces so it follows the desired trajectory as closely as possible.

Disturbance rejection is one of the very oldest applications of feedback control. In this case, the "command" is simply a constant setpoint to which the output is to be held as the environment changes. A very common example of this is the room thermostat whose job it is to hold the room temperature close to the setpoint as outside temperature and wind change, and as doors and windows are opened and closed.

Finally, to design a controller for a dynamic system, it is necessary to have a **mathematical model** of the dynamic response of the system being controlled in all but the simplest cases. Unfortunately, almost all physical systems are very complex and often nonlinear. As a result, the design will usually be based on a simplified model and must be **robust** enough that the system meets its performance requirements when applied to the real device. Furthermore, as time and the environment change, even the best of models will be in error because the system dynamics have changed. Again, the design must not be too **sensitive** to these inevitable changes and it must work well enough regardless.

The **tools** available to control engineers to design and build feedback control systems have evolved over time. The development of digital computers has been especially important both as computation aids and as embedded control devices. As computation devices, computers have permitted identification of increasingly complex models and the application of very sophisticated control design methods. Also, as embedded devices, digital controllers have permitted the implementation of very complex control laws. Control engineers must not only be skilled in using these design tools, but also need to understand the concepts behind these tools to be able to make the best use of them. Also important is that the control engineer understands both the capabilities and the limitations of the controller devices available.

1.4 A Brief History

Interesting histories of early work on feedback control have been written by Mayr (1970) and Åström (2014), who trace the control of mechanisms to antiquity. Two of the earliest examples are the control of flow rate to regulate a water clock and the control of liquid level in a wine vessel, which is thereby kept full regardless of how many cups are dipped from it. The control of fluid flow rate is reduced to the control of fluid level, since a small orifice will produce constant flow if the pressure is constant, which is the case if the level of the liquid above the orifice is constant. The mechanism of the liquid-level control invented in antiquity and still used today (for example, in the water tank of the ordinary flush toilet) is the **float valve**. As the liquid level falls, so does the float, allowing the flow into the tank to increase; as the level rises, the flow is reduced and if necessary cut off. Figure 1.7 shows how a float valve operates. Notice here the sensor and actuator are not separate

Liquid-level control

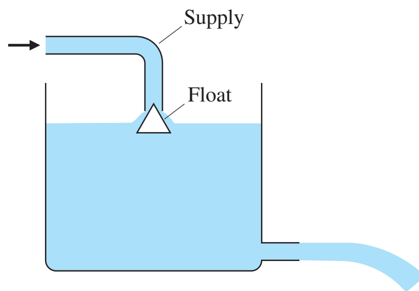

Figure 1.7
Early historical control of liquid level and flow

devices but are contained in the carefully shaped float-and-supply-tube combination.

Drebbel's incubator

A more recent invention described by Mayr (1970) is a system, designed by Cornelis Drebbel in about 1620, to control the temperature of a furnace used to heat an incubator[6] (see Fig. 1.8). The furnace consists of a box to contain the fire, with a flue at the top fitted with a damper. Inside the fire box is the double-walled incubator box, the hollow walls of which are filled with water to transfer the heat evenly to the incubator. The temperature sensor is a glass vessel filled with alcohol and mercury and placed in the water jacket around the incubator box. As the fire heats the box and water, the alcohol expands and the riser floats up, lowering the damper on the flue. If the box is too cold, the alcohol contracts, the damper is opened, and the fire burns hotter.

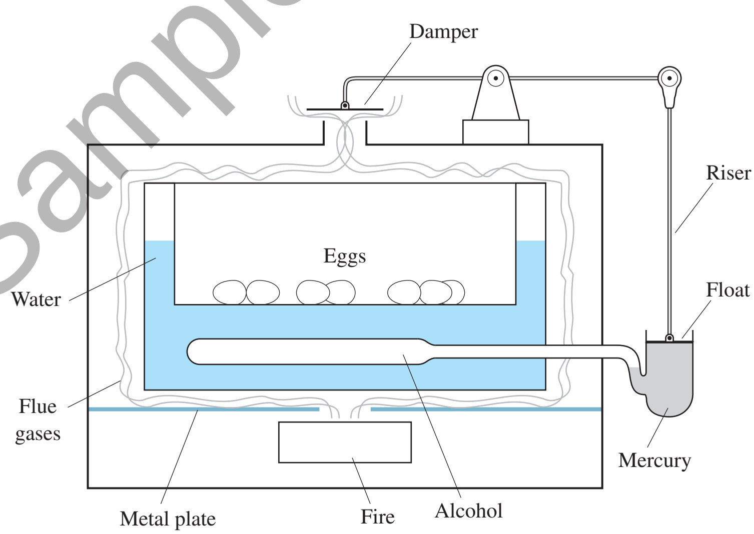

Figure 1.8
Drebbel's incubator for hatching chicken eggs

[6]French doctors introduced incubators into the care of premature babies over 100 years ago.

The desired temperature is set by the length of the riser, which sets the opening of the damper for a given expansion of the alcohol.

Fly-ball governor

A famous problem in the chronicles of control systems was the search for a means to control the rotation speed of a shaft. Much early work (Fuller, 1976) seems to have been motivated by the desire to automatically control the speed of the grinding stone in a wind-driven flour mill. Of various methods attempted, the one with the most promise used a conical pendulum, or **fly-ball governor**, to measure the speed of the mill. The sails of the driving windmill were rolled up or let out with ropes and pulleys, much like a window shade, to maintain fixed speed. However, it was adaptation of these principles to the steam engine in the laboratories of James Watt around 1788 that made the fly-ball governor famous. An early version is shown in Fig. 1.9, while Figs. 1.10 and 1.11 show a close-up of a fly-ball governor and a sketch of its components.

The action of the fly-ball governor (also called a centrifugal governor) is simple to describe. Suppose the engine is operating in equilibrium. Two weighted balls spinning around a central shaft can be seen to describe a cone of a given angle with the shaft. When a load is suddenly applied to the engine, its speed will slow, and the balls of the governor will drop to a smaller cone. Thus the ball angle is used to sense the output speed. This action, through the levers, will open the main valve to the steam chest (which is the actuator) and admit more steam to the engine, restoring most of the lost speed. To hold the steam valve at a new position, it is necessary for the fly balls to rotate at a different angle, implying that the speed under load is not exactly the same as before. We saw this effect earlier with cruise control, where feedback control gave a very small error. To recover the exact same speed in the system, it would require resetting the desired speed setting by changing the length of the rod from the lever to the valve. Subsequent inventors

Figure 1.9
Photograph of an early Watt steam engine
Source: Chronicle/Alamy Stock Photo

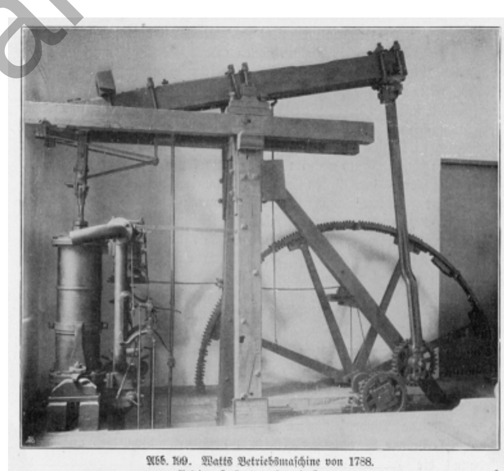

Figure 1.10
Close-up of the fly-ball governor
Source: Washington Imaging/Alamy Stock Photo

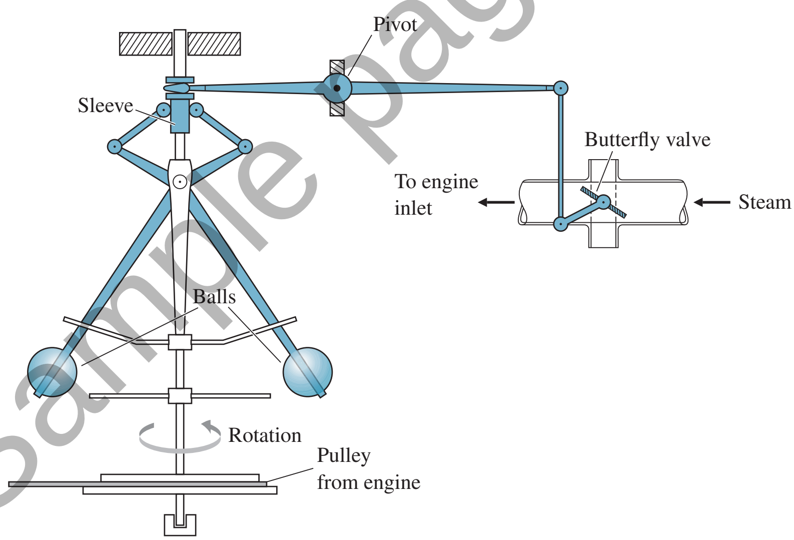

Figure 1.11
Operating parts of a fly-ball governor

introduced mechanisms that integrated the speed error to provide automatic reset. In Chapter 4, we will analyze these systems to show that such integration can result in feedback systems with zero steady-state error to constant disturbances.

Beginnings of control theory

Because Watt was a practical man, he did not engage in theoretical analysis of the governor, similar to the millwrights earlier. Fuller (1976) has traced the early development of control theory to a period of studies from Christiaan Huygens in 1673 to James Clerk Maxwell in 1868. Fuller gives particular credit to the contributions of G. B. Airy, professor of mathematics and astronomy at Cambridge University from 1826 to 1835 and Astronomer Royal at Greenwich Observatory from 1835 to 1881. Airy was concerned with speed control; if his telescopes could be rotated counter to the rotation of the earth, a fixed star could be observed for extended periods. Using the centrifugal-pendulum governor he discovered that it was capable of unstable motion—"and the machine (if I may so express myself) became perfectly wild" (Airy, 1840; quoted in Fuller, 1976). According to Fuller, Airy was the first worker to discuss instability in a feedback control system and the first to analyze such a system using differential equations. These attributes signal the beginnings of the study of feedback control dynamics.

Stability analysis

The first systematic study of the stability of feedback control was apparently given in the paper "On Governors" by Maxwell (1868).[7] In this paper, Maxwell developed the differential equations of the governor, linearized them about equilibrium, and stated that stability depends on the roots of a certain (characteristic) equation having negative real parts. Maxwell attempted to derive conditions on the coefficients of a polynomial that would hold if all the roots had negative real parts. He was successful only for second- and third-order cases. Determining criteria for stability was the problem for the Adams Prize of 1877, which was won by E. J. Routh.[8] His criterion, developed in his essay, remains of sufficient interest that control engineers are still learning how to apply his simple technique. Analysis of the characteristic equation remained the foundation of control theory until the invention of the electronic feedback amplifier by H. S. Black in 1927 at Bell Telephone Laboratories.

Shortly after publication of Routh's work, the Russian mathematician Lyapunov (1892) began studying the question of stability of motion. His studies were based on the nonlinear differential equations of motion, and also included results for linear equations that are equivalent to Routh's criterion. His work was fundamental to what is now called the state-variable approach to control theory, but was not introduced into the control literature until about 1958.

Frequency response

The development of the feedback amplifier is briefly described in an interesting article based on a talk by Bode (1960) reproduced in Bellman and Kalaba (1964). With the introduction of electronic amplifiers, long-distance telephoning became possible in the decades following World War I. However, as distances increased, so did the loss of electrical energy; in spite of using larger-diameter wires, increasing numbers of amplifiers were needed to replace the lost energy. Unfortunately, large numbers of amplifiers resulted in much distortion since the small nonlinearity of the vacuum tubes then used in electronic amplifiers were

[7]An exposition of Maxwell's contribution is given in Fuller (1976).

[8]E. J. Routh was first academically in his class at Cambridge University in 1854, while J. C. Maxwell was second. In 1877, Maxwell was on the Adams Prize Committee that chose the problem of stability as the topic for the year.