
Feedback Systems

An Introduction for Scientists and Engineers
SECOND EDITION

Karl Johan Åström
Richard M. Murray

Version v3.0e (4 Jun 2015)

This is the electronic edition of *Feedback Systems* and is available from <http://www.cds.caltech.edu/~murray/amwiki>. Hardcover editions may be purchased from Princeton University Press, <http://press.princeton.edu/titles/8701.html>.

This manuscript is for personal use only and may not be reproduced, in whole or in part, without written consent from the publisher (see <http://press.princeton.edu/permissions.html>).

PRINCETON UNIVERSITY PRESS
PRINCETON AND OXFORD

Chapter One

Introduction

Feedback is a central feature of life. The process of feedback governs how we grow, respond to stress and challenge, and regulate factors such as body temperature, blood pressure and cholesterol level. The mechanisms operate at every level, from the interaction of proteins in cells to the interaction of organisms in complex ecologies.

M. B. Hoagland and B. Dodson, *The Way Life Works*, 1995 [HD95].

In this chapter we provide an introduction to the basic concept of *feedback* and the related engineering discipline of *control*. We focus on both historical and current examples, with the intention of providing the context for current tools in feedback and control. Much of the material in this chapter is adapted from [Mur03], and the authors gratefully acknowledge the contributions of Roger Brockett and Gunter Stein to portions of this chapter.

1.1 What Is Feedback?

A *dynamical system* is a system whose behavior changes over time, often in response to external stimulation or forcing. The term *feedback* refers to a situation in which two (or more) dynamical systems are connected together such that each system influences the other and their dynamics are thus strongly coupled. Simple causal reasoning about a feedback system is difficult because the first system influences the second and the second system influences the first, leading to a circular argument. It makes conventional causal reasoning difficult and it is necessary to analyze the system as a whole. A consequence of this is that the behavior of feedback systems is often counterintuitive, and it is therefore necessary to resort to formal methods to understand them.

Figure 1.1 illustrates in block diagram form the idea of feedback. We often use




Figure 1.1: Open and closed loop systems. (a) The output of system 1 is used as the input of system 2, and the output of system 2 becomes the input of system 1, creating a closed loop system. (b) The interconnection between system 2 and system 1 is removed, and the system is said to be open loop.




Figure 1.2: The centrifugal governor and the steam engine. The centrifugal governor on the left consists of a set of flyballs that spread apart as the speed of the engine increases. The steam engine on the right uses a centrifugal governor (above and to the left of the flywheel) to regulate its speed. (Credit: Machine a Vapeur Horizontale de Philip Taylor [1828].)

the terms *open loop* and *closed loop* when referring to such systems. A system is said to be a closed loop system if the systems are interconnected in a cycle, as shown in Figure 1.1a. If we break the interconnection, we refer to the configuration as an open loop system, as shown in Figure 1.1b.

As the quote at the beginning of this chapter illustrates, a major source of examples of feedback systems is biology. Biological systems make use of feedback in an extraordinary number of ways, on scales ranging from molecules to cells to organisms to ecosystems. One example is the regulation of glucose in the bloodstream through the production of insulin and glucagon by the pancreas. The body attempts to maintain a constant concentration of glucose, which is used by the body's cells to produce energy. When glucose levels rise (after eating a meal, for example), the hormone insulin is released and causes the body to store excess glucose in the liver. When glucose levels are low, the pancreas secretes the hormone glucagon, which has the opposite effect. Referring to Figure 1.1, we can view the liver as system 1 and the pancreas as system 2. The output from the liver is the glucose concentration in the blood, and the output from the pancreas is the amount of insulin or glucagon produced. The interplay between insulin and glucagon secretions throughout the day helps to keep the blood-glucose concentration constant, at about 90 mg per 100 mL of blood.

An early engineering example of a feedback system is a centrifugal governor, in which the shaft of a steam engine is connected to a flyball mechanism that is itself connected to the throttle of the steam engine, as illustrated in Figure 1.2. The system is designed so that as the speed of the engine increases (perhaps because of a lessening of the load on the engine), the flyballs spread apart and a linkage causes the throttle on the steam engine to be closed. This in turn slows down the engine, which causes the flyballs to come back together. We can model this system as a closed loop system by taking system 1 as the steam engine and system 2 as

the governor. When properly designed, the flyball governor maintains a constant speed of the engine, roughly independent of the loading conditions. The centrifugal governor was an enabler of the successful Watt steam engine, which fueled the industrial revolution.

Feedback has many interesting properties that can be exploited in designing systems. As in the case of glucose regulation or the flyball governor, feedback can make a system resilient toward external influences. It can also be used to create linear behavior out of nonlinear components, a common approach in electronics. More generally, feedback allows a system to be insensitive both to external disturbances and to variations in its individual elements.

Feedback has potential disadvantages as well. It can create dynamic instabilities in a system, causing oscillations or even runaway behavior. Another drawback, especially in engineering systems, is that feedback can introduce unwanted sensor noise into the system, requiring careful filtering of signals. It is for these reasons that a substantial portion of the study of feedback systems is devoted to developing an understanding of dynamics and a mastery of techniques in dynamical systems.

Feedback systems are ubiquitous in both natural and engineered systems. Control systems maintain the environment, lighting and power in our buildings and factories; they regulate the operation of our cars, consumer electronics and manufacturing processes; they enable our transportation and communications systems; and they are critical elements in our military and space systems. For the most part they are hidden from view, buried within the code of embedded microprocessors, executing their functions accurately and reliably. Feedback has also made it possible to increase dramatically the precision of instruments such as atomic force microscopes (AFMs) and telescopes.

In nature, homeostasis in biological systems maintains thermal, chemical and biological conditions through feedback. At the other end of the size scale, global climate dynamics depend on the feedback interactions between the atmosphere, the oceans, the land and the sun. Ecosystems are filled with examples of feedback due to the complex interactions between animal and plant life. Even the dynamics of economies are based on the feedback between individuals and corporations through markets and the exchange of goods and services.

1.2 What Is Control?

The term *control* has many meanings and often varies between communities. In this book, we define control to be the use of algorithms and feedback in engineered systems. Thus, control includes such examples as feedback loops in electronic amplifiers, setpoint controllers in chemical and materials processing, “fly-by-wire” systems on aircraft and even router protocols that control traffic flow on the Internet. Emerging applications include high-confidence software systems, autonomous vehicles and robots, real-time resource management systems and biologically engineered systems. At its core, control is an *information science* and includes the use of information in both analog and digital representations.




Figure 1.3: Components of a computer-controlled system. The upper dashed box represents the process dynamics, which include the sensors and actuators in addition to the dynamical system being controlled. Noise and external disturbances can perturb the dynamics of the process. The controller is shown in the lower dashed box. It consists of a filter and analog-to-digital (A/D) and digital-to-analog (D/A) converters, as well as a computer that implements the control algorithm. A system clock controls the operation of the controller, synchronizing the A/D, D/A and computing processes. The operator input is also fed to the computer as an external input.

A modern controller senses the operation of a system, compares it against the desired behavior, computes corrective actions based on a model of the system's response to external inputs and actuates the system to effect the desired change. This basic *feedback loop* of sensing, computation and actuation is the central concept in control. The key issues in designing control logic are ensuring that the dynamics of the closed loop system are stable (bounded disturbances give bounded errors) and that they have additional desired behavior (good disturbance attenuation, fast responsiveness to changes in operating point, etc). These properties are established using a variety of modeling and analysis techniques that capture the essential dynamics of the system and permit the exploration of possible behaviors in the presence of uncertainty, noise and component failure.

A typical example of a control system is shown in Figure 1.3. The basic elements of sensing, computation and actuation are clearly seen. In modern control systems, computation is typically implemented on a digital computer, requiring the use of analog-to-digital (A/D) and digital-to-analog (D/A) converters. Uncertainty enters the system through noise in sensing and actuation subsystems, external disturbances that affect the underlying system operation and uncertain dynamics in the system (parameter errors, unmodeled effects, etc). The algorithm that computes the control action as a function of the sensor values is often called a *control law*. The system can be influenced externally by an operator who introduces *command signals* to the system.

Control engineering relies on and shares tools from physics (dynamics and modeling), computer science (information and software) and operations research (optimization, probability theory and game theory), but it is also different from these subjects in both insights and approach.

Perhaps the strongest area of overlap between control and other disciplines is in the modeling of physical systems, which is common across all areas of engineering and science. One of the fundamental differences between control-oriented modeling and modeling in other disciplines is the way in which interactions between subsystems are represented. Control relies on a type of input/output modeling that allows many new insights into the behavior of systems, such as disturbance attenuation and stable interconnection. Model reduction, where a simpler (lower-fidelity) description of the dynamics is derived from a high-fidelity model, is also naturally described in an input/output framework. Perhaps most importantly, modeling in a control context allows the design of *robust* interconnections between subsystems, a feature that is crucial in the operation of all large engineered systems.

Control is also closely associated with computer science since virtually all modern control algorithms for engineering systems are implemented in software. However, control algorithms and software can be very different from traditional computer software because of the central role of the dynamics of the system and the real-time nature of the implementation.

Feedback Examples

Feedback has many interesting and useful properties. It makes it possible to design precise systems from imprecise components and to make relevant quantities in a system change in a prescribed fashion. An unstable system can be stabilized using feedback, and the effects of external disturbances can be reduced. Feedback also offers new degrees of freedom to a designer by exploiting sensing, actuation and computation. In this section we survey some of the important applications and trends for feedback in the world around us.

1.3 Feedback Properties

Feedback is a powerful idea which, as we have seen, is used extensively in natural and technological systems. The principle of feedback is simple: base correcting actions on the difference between desired and actual performance. In engineering, feedback has been rediscovered and patented many times in many different contexts. The use of feedback has often resulted in vast improvements in system capability, and these improvements have sometimes been revolutionary, as discussed above. The reason for this is that feedback has some truly remarkable properties. In this section we will discuss some of the properties of feedback that can be understood intuitively. This intuition will be formalized in subsequent chapters.




Figure 1.4: A feedback system for controlling the speed of a vehicle. In the block diagram on the left, the speed of the vehicle is measured and compared to the desired speed within the “Compute” block. Based on the difference in the actual and desired speeds, the throttle (or brake) is used to modify the force applied to the vehicle by the engine, drivetrain and wheels. The figure on the right shows the response of the control system to a commanded change in speed from 25 m/s to 30 m/s. The three different curves correspond to differing masses of the vehicle, between 1000 and 3000 kg, demonstrating the robustness of the closed loop system to a very large change in the vehicle characteristics.

Robustness to Uncertainty

One of the key uses of feedback is to provide robustness to uncertainty. By measuring the difference between the sensed value of a regulated signal and its desired value, we can supply a corrective action. If the system undergoes some change that affects the regulated signal, then we sense this change and try to force the system back to the desired operating point. This is precisely the effect that Watt exploited in his use of the centrifugal governor on steam engines.

As an example of this principle, consider the simple feedback system shown in Figure 1.4. In this system, the speed of a vehicle is controlled by adjusting the amount of gas flowing to the engine. Simple *proportional-integral* (PI) feedback is used to make the amount of gas depend on both the error between the current and the desired speed and the integral of that error. The plot on the right shows the results of this feedback for a step change in the desired speed and a variety of different masses for the car, which might result from having a different number of passengers or towing a trailer. Notice that independent of the mass (which varies by a factor of 3!), the steady-state speed of the vehicle always approaches the desired speed and achieves that speed within approximately 5 s. Thus the performance of the system is robust with respect to this uncertainty.

Another early example of the use of feedback to provide robustness is the negative feedback amplifier. When telephone communications were developed, amplifiers were used to compensate for signal attenuation in long lines. A vacuum tube was a component that could be used to build amplifiers. Distortion caused by the nonlinear characteristics of the tube amplifier together with amplifier drift were obstacles that prevented the development of line amplifiers for a long time. A major breakthrough was the invention of the feedback amplifier in 1927 by Harold S. Black, an electrical engineer at Bell Telephone Laboratories. Black used *negative feedback*, which reduces the gain but makes the amplifier insensitive to variations in tube characteristics. This invention made it possible to build stable amplifiers

with linear characteristics despite the nonlinearities of the vacuum tube amplifier.

Design of Dynamics

Another use of feedback is to change the dynamics of a system. Through feedback, we can alter the behavior of a system to meet the needs of an application: systems that are unstable can be stabilized, systems that are sluggish can be made responsive and systems that have drifting operating points can be held constant. Control theory provides a rich collection of techniques to analyze the stability and dynamic response of complex systems and to place bounds on the behavior of such systems by analyzing the gains of linear and nonlinear operators that describe their components.

An example of the use of control in the design of dynamics comes from the area of flight control. The following quote, from a lecture presented by Wilbur Wright to the Western Society of Engineers in 1901 [McF53], illustrates the role of control in the development of the airplane:

Men already know how to construct wings or airplanes, which when driven through the air at sufficient speed, will not only sustain the weight of the wings themselves, but also that of the engine, and of the engineer as well. Men also know how to build engines and screws of sufficient lightness and power to drive these planes at sustaining speed ... Inability to balance and steer still confronts students of the flying problem ... When this one feature has been worked out, the age of flying will have arrived, for all other difficulties are of minor importance.

The Wright brothers thus realized that control was a key issue to enable flight. They resolved the compromise between stability and maneuverability by building an airplane, the Wright Flyer, that was unstable but maneuverable. The Flyer had a rudder in the front of the airplane, which made the plane very maneuverable. A disadvantage was the necessity for the pilot to keep adjusting the rudder to fly the plane: if the pilot let go of the stick, the plane would crash. Other early aviators tried to build stable airplanes. These would have been easier to fly, but because of their poor maneuverability they could not be brought up into the air. By using their insight and skillful experiments the Wright brothers made the first successful flight at Kitty Hawk in 1903.

Since it was quite tiresome to fly an unstable aircraft, there was strong motivation to find a mechanism that would stabilize an aircraft. Such a device, invented by Sperry, was based on the concept of feedback. Sperry used a gyro-stabilized pendulum to provide an indication of the vertical. He then arranged a feedback mechanism that would pull the stick to make the plane go up if it was pointing down, and vice versa. The Sperry autopilot was the first use of feedback in aeronautical engineering, and Sperry won a prize in a competition for the safest airplane in Paris in 1914. Figure 1.5 shows the Curtiss seaplane and the Sperry autopilot. The autopilot is a good example of how feedback can be used to stabilize




Figure 1.5: Aircraft autopilot system. The Sperry autopilot (left) contained a set of four gyros coupled to a set of air valves that controlled the wing surfaces. The 1912 Curtiss used an autopilot to stabilize the roll, pitch and yaw of the aircraft and was able to maintain level flight as a mechanic walked on the wing (right) [Hug93].

an unstable system and hence “design the dynamics” of the aircraft.

One of the other advantages of designing the dynamics of a device is that it allows for increased modularity in the overall system design. By using feedback to create a system whose response matches a desired profile, we can hide the complexity and variability that may be present inside a subsystem. This allows us to create more complex systems by not having to simultaneously tune the responses of a large number of interacting components. This was one of the advantages of Black’s use of negative feedback in vacuum tube amplifiers: the resulting device had a well-defined linear input/output response that did not depend on the individual characteristics of the vacuum tubes being used.

Creating Modularity

Feedback can be used to create modularity and shape well-defined relations between inputs and outputs in a structured hierarchical manner. A typical example is the electrical drive system shown in Figure 1.6, which has an architecture with three cascaded loops. The innermost loop is a current loop, the controller CC drives the amplifier so that the current to the motor follows the set point. The velocity loop with the controller VC drives the set point of the current controller so that velocity follows the set point of VC. The outer loop drives the set point of the velocity loop to follow the set point of the position controller.

The control architecture with nested loops shown in Figure 1.6 is common. It simplifies both design, commissioning and operation. Consider for example the design of the velocity loop. With a well-designed current controller the motor current follows the set point of the controller independent of friction and other effects. Since the motor velocity is proportional to the current the dynamics relating velocity to the input of the current controller is approximately an integrator, because force is proportional to current and angular acceleration is proportional to force.




Figure 1.6: Block diagram of a system for position control. The system has three cascaded loops for control of current, velocity and position.

The design of the velocity loop is then simple. With a well-designed velocity loop, the design of the position loop is also simple. The loops can also be tuned sequentially starting with the inner loop. The architecture illustrates how feedback can be used to simplify modeling and create modular systems.

Drawbacks of Feedback

While feedback has many advantages, it also has some drawbacks. Chief among these is the possibility of instability if the system is not designed properly. We are all familiar with the effects of *positive feedback* when the amplification on a microphone is turned up too high in a room. This is an example of feedback instability, something that we obviously want to avoid. This is tricky because we must design the system not only to be stable under nominal conditions but also to remain stable under all possible perturbations of the dynamics.

In addition to the potential for instability, feedback inherently couples different parts of a system. One common problem is that feedback often injects measurement noise into the system. Measurements must be carefully filtered so that the actuation and process dynamics do not respond to them, while at the same time ensuring that the measurement signal from the sensor is properly coupled into the closed loop dynamics (so that the proper levels of performance are achieved).

Another potential drawback of control is the complexity of embedding a control system in a product. While the cost of sensing, computation and actuation has decreased dramatically in the past few decades, the fact remains that control systems are often complicated, and hence one must carefully balance the costs and benefits. An early engineering example of this is the use of microprocessor-based feedback systems in automobiles. The use of microprocessors in automotive applications began in the early 1970s and was driven by increasingly strict emissions standards, which could be met only through electronic controls. Early systems were expensive and failed more often than desired, leading to frequent customer dissatisfaction. It was only through aggressive improvements in technology that the performance, reliability and cost of these systems allowed them to be used in a transparent fashion. Even today, the complexity of these systems is such that it is difficult for an individual car owner to fix problems.

Feedforward

Feedback is reactive: there must be an error before corrective actions are taken. However, in some circumstances it is possible to measure a disturbance before it enters the system, and this information can then be used to take corrective action before the disturbance has influenced the system. The effect of the disturbance is thus reduced by measuring it and generating a control signal that counteracts it. This way of controlling a system is called *feedforward*. Feedforward is particularly useful in shaping the response to command signals because command signals are always available. Since feedforward attempts to match two signals, it requires good process models; otherwise the corrections may have the wrong size or may be badly timed.

The ideas of feedback and feedforward are very general and appear in many different fields. In economics, feedback and feedforward are analogous to a market-based economy versus a planned economy. In business, a feedforward strategy corresponds to running a company based on extensive strategic planning, while a feedback strategy corresponds to a reactive approach. In biology, feedforward has been suggested as an essential element for motion control in humans that is tuned during training. Experience indicates that it is often advantageous to combine feedback and feedforward, and the correct balance requires insight and understanding of their respective properties.

Positive Feedback

In most of this text, we will consider the role of *negative feedback*, in which we attempt to regulate the system by reacting to disturbances in a way that decreases the effect of those disturbances. Positive feedback plays an important role in social and biological systems, *positive feedback*. can play an important role. In a system with positive feedback, the increase in some variable or signal leads to a situation in which that quantity is further increased through its dynamics. This has a destabilizing effect and is usually accompanied by a saturation that limits the growth of the quantity. Although often considered undesirable, this behavior is used in biological (and engineering) systems to obtain a very fast response to a condition or signal. Another common use of positive feedback is the design of systems with oscillatory dynamics.

One example of the use of positive feedback is to create switching behavior, in which a system maintains a given state until some input crosses a threshold. Hysteresis is often present so that noisy inputs near the threshold do not cause the system to jitter. This type of behavior is called *bistability* and is often associated with memory devices.

1.4 Simple Forms of Feedback

The idea of feedback to make corrective actions based on the difference between the desired and the actual values of a quantity can be implemented in many differ-




Figure 1.7: Input/output characteristics of on-off controllers. Each plot shows the input on the horizontal axis and the corresponding output on the vertical axis. Ideal on-off control is shown in (a), with modifications for a dead zone (b) or hysteresis (c). Note that for on-off control with hysteresis, the output depends on the value of past inputs.

ent ways. The benefits of feedback can be obtained by very simple feedback laws such as on-off control, proportional control and proportional-integral-derivative control. In this section we provide a brief preview of some of the topics that will be studied more formally in the remainder of the text.

On-Off Control

A simple feedback mechanism can be described as follows:

$$u = \begin{cases} u_{\max} & \text{if } e > 0 \\ u_{\min} & \text{if } e < 0, \end{cases} \quad (1.1)$$

where the *control error* $e = r - y$ is the difference between the reference signal (or command signal) r and the output of the system y and u is the actuation command. Figure 1.7a shows the relation between error and control. This control law implies that maximum corrective action is always used.

The feedback in equation (1.1) is called *on-off control*. One of its chief advantages is that it is simple and there are no parameters to choose. On-off control often succeeds in keeping the process variable close to the reference, such as the use of a simple thermostat to maintain the temperature of a room. It typically results in a system where the controlled variables oscillate, which is often acceptable if the oscillation is sufficiently small.

Notice that in equation (1.1) the control variable is not defined when the error is zero. It is common to make modifications by introducing either a dead zone or hysteresis (see Figure 1.7b and 1.7c).

PID Control

The reason why on-off control often gives rise to oscillations is that the system overreacts since a small change in the error makes the actuated variable change over the full range. This effect is avoided in *proportional control*, where the characteristic of the controller is proportional to the control error for small errors. This

can be achieved with the control law

$$u = \begin{cases} u_{\max} & \text{if } e \geq e_{\max} \\ k_p e & \text{if } e_{\min} < e < e_{\max} \\ u_{\min} & \text{if } e \leq e_{\min}, \end{cases} \quad (1.2)$$

where k_p is the controller gain, $e_{\min} = u_{\min}/k_p$ and $e_{\max} = u_{\max}/k_p$. The interval (e_{\min}, e_{\max}) is called the *proportional band* because the behavior of the controller is linear when the error is in this interval:

$$u = k_p(r - y) = k_p e \quad \text{if } e_{\min} \leq e \leq e_{\max}. \quad (1.3)$$

While a vast improvement over on-off control, proportional control has the drawback that the process variable often deviates from its reference value. In particular, if some level of control signal is required for the system to maintain a desired value, then we must have $e \neq 0$ in order to generate the requisite input.

This can be avoided by making the control action proportional to the integral of the error:

$$u(t) = k_i \int_0^t e(\tau) d\tau. \quad (1.4)$$

This control form is called *integral control*, and k_i is the integral gain. It can be shown through simple arguments that a controller with integral action has zero steady-state error (Exercise 1.5). The catch is that there may not always be a steady state because the system may be oscillating. In addition, if the control action has magnitude limits, as in equation (1.2), an effect known as “integrator windup” can occur and may result in poor performance unless appropriate “anti-windup” compensation is used.

An additional refinement is to provide the controller with an anticipative ability by using a prediction of the error. A simple prediction is given by the linear extrapolation

$$e(t + T_d) \approx e(t) + T_d \frac{de(t)}{dt},$$

which predicts the error T_d time units ahead. Combining proportional, integral and derivative control, we obtain a controller that can be expressed mathematically as

$$u(t) = k_p e(t) + k_i \int_0^t e(\tau) d\tau + k_d \frac{de(t)}{dt}. \quad (1.5)$$

The control action is thus a sum of three terms: the past as represented by the integral of the error, the present as represented by the proportional term and the future as represented by a linear extrapolation of the error (the derivative term). This form of feedback is called a *proportional-integral-derivative (PID) controller* and its action is illustrated in Figure 1.8.

A PID controller is very useful and is capable of solving a wide range of control problems. More than 95% of all industrial control problems are solved by PID control, although many of these controllers are actually *proportional-integral (PI) controllers* because derivative action is often not included [DM02]. There are




Figure 1.8: Action of a PID controller. At time t , the proportional term depends on the instantaneous value of the error. The integral portion of the feedback is based on the integral of the error up to time t (shaded portion). The derivative term provides an estimate of the growth or decay of the error over time by looking at the rate of change of the error. T_d represents the approximate amount of time in which the error is projected forward (see text).

also more advanced controllers, which differ from PID controllers by using more sophisticated methods for prediction.

1.5 Combining Feedback with Logic

The PID controller is a continuous time system. The on-off controller can be viewed both as a controller and a logic system. Continuous control is often combined with logic to cope with different operating conditions. Logic is typically related to changes in operating conditions, equipment protection, manual interaction and saturating actuators. One situation is when there is one variable that is of primary interest, but other variables may have to be controlled for equipment protection. For example, when controlling a compressor the outflow is the primary variable but it may be necessary to switch to a different mode to avoid compressor stall, which may damage the compressor. We illustrate some ways in which, logic and feedback are combined by a few examples.

Collision Avoidance

Path following is a primary criterion for controlling a robot or an autonomous vehicle but these systems also have a system for collision detection which will override path following if there is a risk for collision with another object. It is thus necessary to have logic for choosing between path following and collision avoidance. Since controllers have state it is necessary to handle the controller states properly when switching between different control objectives as will be discussed further in Section 2.3.

Server Farms

Server farms consist of a large number of servers. Large farms may have thousands of processors. Power consumption for driving the servers and for cooling

them is a prime concern. The cost for energy can be more than 40% of the total cost for data centers, which is of the order of a million \$ per month for a large installation [EKR03]. The prime task of the farm is to respond to a strongly varying demand. There are constraints given by electricity consumption and the available cooling capacity. The throughput of an individual server depends on the clock rate, which can be changed by the voltage applied to the system. Increasing the supply voltage increases the energy consumption and more cooling is required. Capacity can be increased rapidly if a server is switched on simply by increasing the voltage to a server, but a server that is switched on consumes energy and requires cooling. Control of server farms is often performed by a combination of feedback and logic. To save energy it is advantageous to switch off servers that are not required, but it takes some time to switch on a new server. A control system for a server farm requires individual control of the voltage and cooling of each server and a strategy for switching servers on and off. Temperature is important. Overheating will reduce the life time of the system and may even destroy it. The cooling system is complicated because cooling air goes through the servers in series and parallel. The measured value for the cooling system is therefore the server with the highest temperature. Control of server farms is often performed with a combination of feedback and logic.

Air-Fuel Control

Air-fuel control is an important problem for ship boilers. The control system consists of two loops for controlling air and oil flow and a supervising controller that adjusts the air-fuel ratio. The ratio should be adjusted for optimal efficiency when the ships are on open sea but it is necessary to run the system with air excess when the ships are in the harbor, since generating black smoke will result in heavy penalties.

An elegant solution to the problem can be obtained by combining PI controllers with maximum and minimum selectors, as shown in Figure 1.9. A *selector* is a static system with several inputs and one output: a maximum selector gives an output which is the largest of the inputs, a minimum selector gives an output which is the smallest of the inputs. Consider the situation when the power demand is increased: the reference r to the air controller is selected as the commanded power level by the maximum selector, and the reference to the oil flow controller is selected as the measured airflow. The oil flow will lag the air flow and there will be air excess. When the commanded power level is decreased, the reference of the oil flow controller is selected as the power demand by the minimum selector and the reference for the air flow controller is selected as the oil flow by the the maximum selector. The system then operates with air excess when power is decreased.

A simulation of the control logic is shown in Figure 1.9b, based on the process model

$$\frac{dx_a}{dt} = -4x_a + 4u_a, \quad \frac{dx_o}{dt} = -x_o + u_o,$$

where x_a and x_o are the states representing air and oil dynamics. The air dynamics

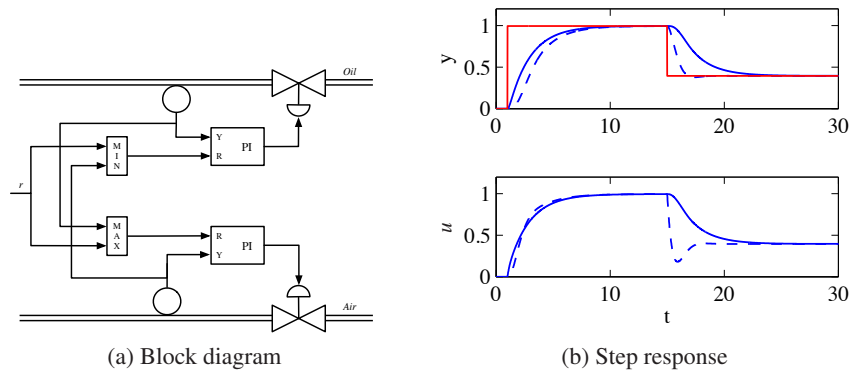


Figure 1.9: Air-fuel controller based on selectors. The left figure shows the system architecture. The letters r and y in the PI controller denotes the input ports for reference and measured signal respectively. The right figure shows a simulation where the power reference r (red) is changed at $t = 1$ and $t = 15$. Notice that the air flow (solid blue) is larger than the fuel flow (dashed) both for increasing and decreasing reference steps.

are faster than the oil dynamics. The PI controllers are described by

$$u_a = -k_{pa}x_a + k_{ia} \int^t (r_a - y_a)dt, \quad r_a = \max(r, x_0),$$

$$u_o = -k_{po}x_o + k_{io} \int^t (r_o - y_o)dt, \quad r_o = \min(r, x_a).$$

The controller gains used in are $k_{pa} = 1$, $k_{ia} = 1$, $k_{po} = 2$ and $k_{io} = 4$.

Selectors are commonly used to implement logic in engines and power systems. They are also used for systems that require very high reliability: by introducing three sensors and only accepting values where two sensors agree it is possible to guard for the failure of a single sensor.

Cruise control

The basic control function in cruise controller such as the one shown in Figure 1.10, is to keep the velocity constant. It is typically done with a PI controller. The controller normally operates in automatic mode but it is necessary to switch it off when braking, accelerating or changing gears. The cruise control system has a human-machine interface that allows the driver to communicate with the system. There are many different ways to implement this system; one version is illustrated in Figure 1.10. The system has four buttons: on-off, set/decelerate, resume/accelerate and cancel. The operation of the system is governed by a finite state machine that controls the modes of the PI controller and the reference generator; see Figure 1.10. The finite state machine has four modes: off, standby, cruise and hold. The state changes depending on actions of the driver who can brake, accelerate and operate with the buttons. The on/off switch moves the states between off and standby. From standby the system can be moved to cruise by pushing the resume/accelerate button. The velocity reference is set as the velocity of the car




Figure 1.10: Finite state machine for cruise control system. The figure on the left shows some typical buttons used to control the system. The controller can be in one of four modes, corresponding to the nodes in the diagram on the right. Transition between the modes is controlled by pressing one of the five buttons on the cruise control interface: on, off, set, resume or cancel.

when the button is released. In the cruise state the operator can change the velocity reference; it is increased by the button res/accelerate and decreased by the button coast/set. If the driver accelerates by pushing the gas pedal the speed increases but it will go back to the set velocity when the gas pedal is released. If the driver brakes the car brakes and the cruise controller goes into hold but it remembers the set point of the controller, it can be brought to the cruise state by pushing the res/accelerate button. The system also moves from cruise mode to standby if the cancel button is pushed. The reference for the velocity controller is remembered. The system goes into off mode by pushing off.

The PI controller is designed to have good regulation properties and to give good transient performance when switching between resume and control modes.

Simple logic can be dealt with using selectors as illustrated in Section 2.3, but more formal methods are needed in more complex situation. Since controllers are dynamical systems they have a state, which has to be dealt with properly when switching controllers.

1.6 Control System Architectures

Most of the control systems we are investigating in this book will be relatively simple feedback loops. In this section we will try to give a glimpse of the fact that in reality the simple loops combine to form a complex system which often has an hierarchical structure with controllers, logic and optimization in different combinations.

Freight Train Trip Optimizer

Typical requirements for operating a freight train is to arrive in time and to use as little fuel as possible. The key issue is to avoid unnecessary braking. Figure 1.11 illustrates a system developed by General Electric. At the low level the train has a speed regulator and a simple logic to avoid entering a zone where there is another train. The key disturbance for the speed control is the slope of the ground. The

GE picture from Paul Houpt

Figure 1.11: GE train picture

speed controller has a model of the track, a GPS sensor and an estimator. The set point for the speed controller is obtained from a trip optimizer, computes a driving plan that minimizes the fuel consumption while arriving at the desired arrival time. The arrival time is provided by a dispatch center which in turn may use some optimization.

Paper Factory

Figure 1.12 is a picture of a plant for making craft paper. The factory produces paper for sacks and container board from logs of wood. There are three fibre lines and six paper machines. The plant has a few dozen mechanical and chemical production processes that convert the logs to a slurry of fibres in different steps and six paper machines that convert the fibre slurry to paper. There are several dozen tanks for storage of intermediate products. Each production unit has PI(D) controllers that control, flow, temperature and tank levels. The loops typically operate in a time scales from fractions of seconds to minutes. There is logic to make sure that the process is safe and there is sequencing for start, stop and production changes. The setpoints of the low level control loops are determined from production rates and recipes, sometimes using optimization. The operation of the system is governed by a supervisory system that measures tank levels and sets the production rates of the different production unit. This system performs optimization based on demanded production, measurement of tank levels and flows. The optimization is performed at the time scale of minutes to hours and it is constrained by the production rates of the different production units. At a higher level there is a system for distributing the product and for bringing in raw material using supply chain management. The manufacturing system may also be connected to the business system at an even higher level as illustrated in Figure 1.12. There is also extensive communication because the production unit may cover an area of kilometers and he supply chains for raw material and customers a much larger range.

Process for continuous production in the chemical and pharmaceutical industry are similar to the paper factory but the individual production units may be very different. At an abstract level the systems can be represented by Figure 1.13.

Autonomous Driving

The cruise controller in Figure 1.4 relieves the driver of one task to keep constant speed, but a driver still has many tasks to perform: plan the route, avoid collisions,

Billerud + Global Enterprise control

Figure 1.12: A paper plant. (May change the photo later)

Figure 1.13: Functional architecture of process control system.

decide the proper speed, plan the route, do lane changes, make turns, keep proper distance to the car ahead. Car manufacturers are continuously automating several of these functions going as far as automatic driving.

1.7 Further Reading

The material in this section draws heavily from the report of the Panel on Future Directions on Control, Dynamics and Systems [Mur03]. Several additional papers and reports have highlighted the successes of control [NS99] and new vistas in control [Bro00, Kum01, Wis07]. The early development of control is described by Mayr [May70] and in the books by Bennett [Ben79, Ben93], which cover the period 1800–1955. A fascinating examination of some of the early history of control in the United States has been written by Mindell [Min02]. A popular book that describes many control concepts across a wide range of disciplines is *Out of Control* by Kelly [Kel94]. There are many textbooks available that describe control systems in the context of specific disciplines. For engineers, the textbooks by Franklin, Powell and Emami-Naeini [FPEN05], Dorf and Bishop [DB04], Kuo and Golnaraghi [KG02] and Seborg, Edgar and Mellichamp [SEM04] are widely used. More mathematically oriented treatments of control theory include Sontag [Son98] and Lewis [Lew03]. The book by Hellerstein et al. [HDPT04] provides a description of the use of feedback control in computing systems. A number of books look at the role of dynamics and feedback in biological systems, including Milhorn [Mil66] (now out of print), J. D. Murray [Mur04] and Ellner and Guckenheimer [EG05]. The book by Fradkov [Fra07] and the tutorial article by Bechhoefer [Bec05] cover many specific topics of interest to the physics community.

Systems that combine continuous feedback with logic and sequencing are called hybrid system [RST12]. The theory required to properly model and analyze such systems is outside the scope of this book. It is, however, very common that practical control systems combine feedback control with logic sequencing and selectors; many examples are given in [ÅH05].

Exercises

1.1 (Eye motion) Perform the following experiment and explain your results: Holding your head still, move one of your hands left and right in front of your face, following it with your eyes. Record how quickly you can move your hand before you begin to lose track of it. Now hold your hand still and shake your head left to right, once again recording how quickly you can move before losing track of your hand.

1.2 Identify five feedback systems that you encounter in your everyday environment. For each system, identify the sensing mechanism, actuation mechanism and control law. Describe the uncertainty with respect to which the feedback system provides robustness and/or the dynamics that are changed through the use of feedback.

1.3 (Balance systems) Balance yourself on one foot with your eyes closed for 15 s. Using Figure 1.3 as a guide, describe the control system responsible for keeping you from falling down. Note that the “controller” will differ from that in the diagram (unless you are an android reading this in the far future).

1.4 (Cruise control) Download the MATLAB code used to produce simulations for the cruise control system in Figure 1.4 from the companion web site. Using trial and error, change the parameters of the control law so that the overshoot in speed is not more than 1 m/s for a vehicle with mass $m = 1000$ kg.

1.5 (Integral action) We say that a system with a constant input reaches steady state if all system variables approach constant values as time increases. Show that a controller with integral action, such as those given in equations (1.4) and (1.5), gives zero error if the closed loop system reaches steady state. Notice that there is no saturation in the controller.

1.6 Search the web and pick an article in the popular press about a feedback and control system. Describe the feedback system using the terminology given in the article. In particular, identify the control system and describe (a) the underlying process or system being controlled, along with the (b) sensor, (c) actuator and (d) computational element. If the some of the information is not available in the article, indicate this and take a guess at what might have been used.

