

Feedback Systems:
An Introduction for Scientists and Engineers

Karl Johan Åström
Department of Automatic Control
Lund Institute of Technology

Richard M. Murray
Control and Dynamical Systems
California Institute of Technology

DRAFT v0.915, 26 September 2004
© 2004 Karl Johan Åström and Richard Murray
All rights reserved.

This manuscript is for review purposes only and may not be reproduced, in whole
or in part, without written consent from the authors.

Preface

This book provides an introduction to the basic principles and tools for design and analysis of feedback systems. It is intended to serve a diverse audience of scientists and engineers who are interested in understanding and utilizing feedback in physical, biological, information, and economic systems. To this end, we have chosen to keep the mathematical pre-requisites to a minimum while being careful not to sacrifice rigor in the process. Advanced sections, marked by the “dangerous bend” symbol shown to the right indicate material that is of a more advanced nature and can be skipped on first reading.



This book was originally developed for use in an experimental course at Caltech involving undergraduates and graduate students from a wide variety of disciplines. The course included undergraduates at the junior and senior level in traditional engineering disciplines, as well as first and second year graduate students in engineering and science. This included graduate students in biology, computer science, and economics, requiring a broad approach that emphasized basic principles and did not focus on applications in a given area.

A detailed web site has been prepared as a companion to this text:

<http://www.cds.caltech.edu/~murray/books/am04>

The web site contains the MATLAB and other source code for every example in the book, as well as MATLAB libraries to implement the techniques described in the text.

This book is intended to serve a broad spectrum of audiences and is organized in a slightly unusual fashion compared to many other books on feedback in control. In particular, we introduce a number of concepts in the text which are normally reserved for second year courses on control (and hence often not available to students who are not control majors). This has been done at the expense of certain “traditional” topics, which we felt that the astute student could learn on their own (and are often

explored through the exercises). Examples of topics that we have included are nonlinear behavior, Lyapunov stability, reachability and observability, and fundamental limits of performance and robustness. Topics that we have de-emphasized include root locus techniques, lead/lag compensation (although this is essentially covered in the Chapter 8, on PID control), and detailed rules for generating Bode and Nyquist plots by hand.

The first half of the book focused almost exclusively on so-called “state-space” control systems. We begin in Chapter 2 with a description of modeling of physical, biological and information systems using ordinary differential equations and difference equations. Following this, Chapter 3 looks at the dynamic behavior of models, including definitions of stability and more complicated nonlinear behavior. We provide advanced sections in this chapter on Lyapunov stability, because we find that it is useful in a broad array of applications (and frequently a topic that is not introduced until much later in ones studies). Chapter 4 looks at the input/output behavior of dynamical systems, focusing primarily on linear systems (for which the concepts are best defined). Finally, in Chapter 5, we introduce our first feedback systems by demonstrating how state space control laws can be designed, including a short description of observers. Chapter 5 introduces the key concepts of controllability and observability, which give tremendous insight into the choice of actuators and sensors, whether for engineered or natural systems.

The second half of the book presents tools that are often considered to be from the field of “classical control”. This includes the transfer function, introduced in Chapter 6, which is a fundamental tool for understanding feedback systems. Using transfer functions, one can begin to analyze the stability of feedback systems using loop analysis, which allows us to reason about the closed loop behavior (stability) of a system from its open loop characteristics. This is the subject of Chapter 7, which revolves around the Nyquist stability criterion. In Chapter 8, we again look at the design problem, focusing on proportional-integral-derivative (PID) controllers. PID control is by far the most common design technique in control systems and a useful tool for any student. In Chapter 9, we pull together the results from the second half of the book to analyze the fundamental tradeoffs between robustness and performance. This is also a key chapter illustrating the power of the techniques that have been developed. Finally, in Chapter 10 we provide some insights into the implementation of control systems, primarily to illustrate how some of the concepts of the course are applied to practical problems.

The book is designed for use in a 10–12 week course in feedback systems that can serve to introduce many of the key concepts that are needed in a

variety of disciplines. For a ten-week course, each chapter can be covered in a week's time, with some dropping of topics from the final three chapters of the book. A more leisurely course, spread out over 12–14 weeks, could spend two weeks on modeling (Chapter 2)—particularly for students without much background in ordinary differential equations—and two weeks each on loop analysis (Chapter 7) and robustness and performance (Chapter 9).

In choosing this set of topics and ordering, we necessarily left out some tools which will cause many control systems experts to raise their eyebrows (or choose another textbook). Overall, we believe that the early focus on state space systems, including the concepts of controllability and observability, are of such importance to justify trimming other topics to make room for them. We also included some relatively advanced material on fundamental tradeoffs and limitations of performance, feeling that these provided such insight into the principles of feedback that they could not be left for later. Throughout the text, we have attempted to maintain a balanced set of examples that touch many disciplines, relying on the supplements for more discipline specific examples and exercises.

One additional choice that we felt was very important was the decision not to make use of Laplace transforms in this book. While this is by far the most common approach to teaching feedback systems in engineering, many students in natural and information sciences may lack the necessary background. Since Laplace transforms are not required in any essential way, we have only made a few remarks to tie things together for students with that background. Of course, we make tremendous use of *transfer functions*, which we introduce through the notion of response to exponential inputs, an approach we feel is much more accessible to a broad array of scientists and engineers.

Contents

1	Introduction	1
1.1	What is Feedback?	1
1.2	What is Control?	4
1.3	Feedback Examples	6
1.4	Feedback Principles	18
1.5	Control Tools	25
1.6	Further Reading	28
1.7	Exercises	29
2	System Modeling	31
2.1	Modeling Concepts	31
2.2	Modeling from Experiments	37
2.3	State Space Models	38
2.4	Schematic Diagrams	45
2.5	Examples	49
2.6	Other Types of Models	58
2.7	Further Reading	59
2.8	Exercises	60
3	Dynamic Behavior	65
3.1	Solving differential equations	65
3.2	Qualitative analysis	71
3.3	Stability	76
3.4	Local and Global Behavior	80
3.5	Shaping Dynamic Behavior	82
3.6	System Performance Measures	85
3.7	Further Reading	87
3.8	Exercises	87

4	Input/Output Behavior	93
4.1	Introduction	93
4.2	Properties of Linear Systems	96
4.3	Second Order Systems	103
4.4	Linearization	106
4.5	Discrete Time Linear Systems	109
4.6	Further Reading	109
5	State and Output Feedback	111
5.1	Introduction	111
5.2	Reachability	112
5.3	State Feedback	116
5.4	Observers	123
5.5	Output Feedback	132
5.6	Integral Action	135
5.7	A General Controller Structure	137
5.8	A Design Example	139
5.9	Further Reading	139
5.10	Exercises	140
6	Transfer Functions	143
6.1	Introduction	143
6.2	The Transfer Function	143
6.3	Frequency Response	149
6.4	The Bode Plot	151
6.5	Frequency Responses from Experiments	155
6.6	Block Diagrams and Transfer Functions	160
6.7	The Kalman Decomposition	164
6.8	Laplace Transforms	167
6.9	Further Reading	170
6.10	Exercises	170
7	Loop Analysis	175
7.1	Introduction	175
7.2	The Basic Idea	176
7.3	Nyquist's Stability Theorem	179
7.4	Stability Margins	185
7.5	Bode's Relations	186
7.6	Loop Shaping	191
7.7	Fundamental Limitations	198

7.8	The Small Gain Theorem	205
7.9	Further Reading	206
7.10	Exercises	206
8	PID Control	209
8.1	Introduction	209
8.2	The PID Controller	210
8.3	Integrator Windup	215
8.4	Tuning	219
8.5	Computer Control	227
8.6	Further Reading	230
9	Robustness and Performance	231
9.1	Introduction	231
9.2	An Example	233
9.3	The Basic Feedback Loop	238
9.4	The Gangs of Four and Six	241
9.5	Disturbance Attenuation	247
9.6	Robustness to Process Variations	250
9.7	The Sensitivity Functions	255
9.8	Implications for Design	259
9.9	When are Two Processes Similar?	263
9.10	Specifications	266
9.11	Further Reading	275
9.12	Exercises	275
10	Implementation	277
10.1	Introduction	277
10.2	Sensing Actuation Computing and Communication	278
10.3	Analog Control	285
10.4	Computer Control	287
10.5	Biological Control	296
10.6	Summary	303
10.7	References	304
A	Linear Algebra	305
A.1	Linear spaces and operators	305
A.2	Jordan Form	305
B	Linear Differential Equations	307

C Laplace Transforms	311
C.1 Basic Concepts	311
C.2 Additional Properties	313
C.3 Laplace Transforms	314
D General Theory for Dynamical Systems	319

Chapter 1

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.


Mahlon B. Hoagland and B. Dodson, from *The Way Life Works*, 1995 [12].

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 adopted from [20] and the authors gratefully acknowledge the contributions of Roger Brockett and Gunter Stein for portions of this chapter.


1.1 What is Feedback?

The term *feedback* is used to refer 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 such a system is difficult because the first system influences the second and the second system influences the first, leading to a circular argument. This makes reasoning based on cause and effect tricky and it is necessary to analyze the system as a whole. A consequence of this is that the behavior of a feedback system is often counterintuitive and it is therefore necessary to resort to formal methods to understand them.

As the quote at the beginning of this chapter illustrates, a major source of examples for feedback systems is from biology. Biological systems make



(a) Closed loop



(b) Open loop

Figure 1.1: Components of a modern control system.

use of feedback in an extraordinary number of ways, on scales ranging from molecules to microbes 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. The interplay between insulin and glucagon secretions throughout the day help 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 the 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 1.2.¹ The system is designed so that as the speed of the engine increases (perhaps due to 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. When properly designed, the flyball governor maintains a constant speed of the engine, roughly independent of the loading conditions.

¹The centrifugal governor is often called the "Watt governor", because James Watt popularized its use on steam engines. However, contrary to common belief, James Watt did not invent the flyball governor itself.

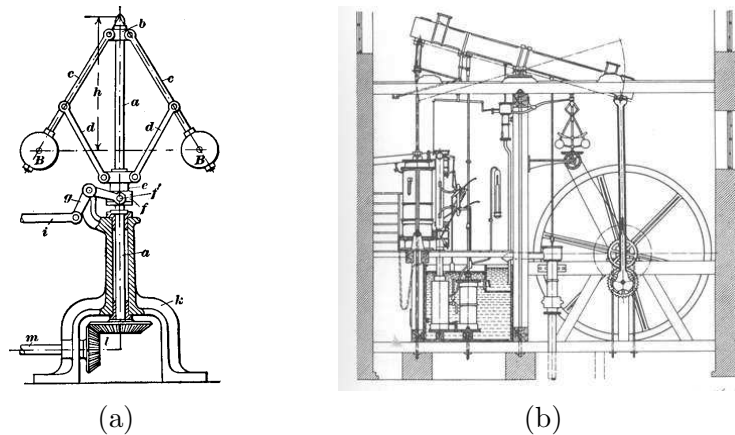


Figure 1.2: The centrifugal governor (a), developed in the 1780s, was an enabler of the successful Watt steam engine (b), which fueled the industrial revolution. Figures courtesy Richard Adamek (copyright 1999) and Cambridge University.

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 very resilient towards external influences. It can also be possible to create linear behavior out of nonlinear components, a common approach in electronics. More generally, feedback allows a system to be very insensitive both to external disturbances and to variations in its individual components. Feedback has one major disadvantage: it may create dynamic instabilities in a system, causing oscillations or even runaway behavior. It is for this reason that a substantial portion of the study of feedback systems is devoted to developing an understanding of dynamics and mastery of techniques in dynamical systems.

Feedback systems are ubiquitous in both natural and engineered systems. Homeostasis in biological systems maintains thermal, chemical, and biological conditions through feedback. Global climate dynamics depend on the feedback interactions between the atmosphere, oceans, land, and the sun. Ecologies are filled with examples of feedback, resulting in complex interactions between animal and plant life. 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, set point 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.

A modern controller senses the operation of a system, compares that 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 the desired behavior (good disturbance rejection, fast responsiveness to changes in operating point, etc). These properties are established using a variety of modeling and analysis techniques that capture the essential physics of the system and permit the exploration of possible behaviors in the presence of uncertainty, noise, and component failures.

A typical example of a modern 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 physics, and uncertain dynamics in the physical system (parameter errors, unmodeled effects, etc).

Control engineering relies on and shares tools from physics (dynamics and modeling), computer science (information and software) and operations research (optimization 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 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

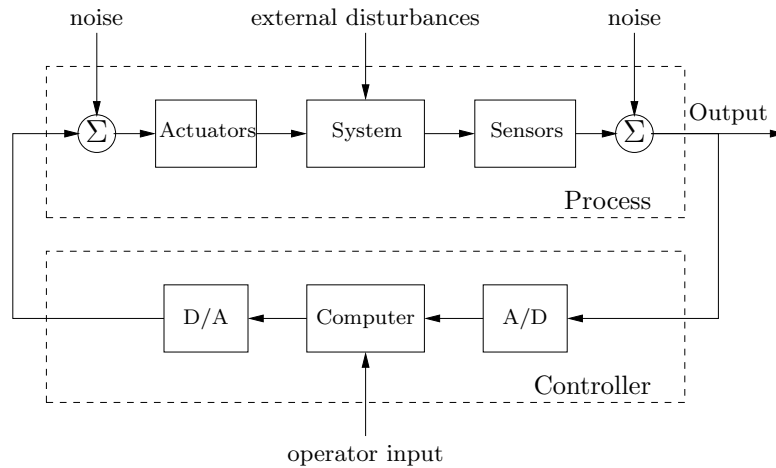


Figure 1.3: Components of a modern control system.

which interactions between subsystems (components) are represented. Control relies on input/output modeling that allows many new insights into the behavior of systems, such as disturbance rejection and stable interconnection. Model reduction, where a simpler (lower-fidelity) description of the dynamics is derived from a high fidelity model, is also very 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 are implemented in software. However, control algorithms and software are very different from traditional computer software. The physics (dynamics) of the system are paramount in analyzing and designing them and their (hard) real-time nature dominates issues of their implementation. From a software-centric perspective, an aircraft is simply another peripheral, while from a control-centric perspective, the computer is just another implementation medium for the feedback law. Neither of these are adequate abstractions, and this is one of the key areas of current research in the field.




Figure 1.4: Early control devices: (a) Honeywell T86 thermostat, originally introduced in 1953, (b) Chrysler cruise control system, introduced in the 1958 Chrysler Imperial (note the centrifugal governor) [13].

1.3 Feedback Examples

Feedback control systems are all around us in the modern technological world. They 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 processors, executing their functions accurately and reliably. Nevertheless, their existence is a major intellectual and engineering accomplishment that is still evolving and growing, promising ever more important consequences to society.

Early Technological Examples

The proliferation of control in engineered systems has occurred primarily in the latter half of the 20th Century. There are some familiar exceptions, such as the Watt governor described earlier and the thermostat (Figure 1.4a), designed at the turn of the century to regulate temperature of buildings.

The thermostat, in particular, is often cited as a simple example of feedback control that everyone can understand. Namely, the device measures the temperature in a building, compares that temperature to a desired set point, and uses the “feedback error” between these two to operate the heating plant, e.g., to turn heating on when the temperature is too low and to

turn it off when temperature is too high. This explanation captures the essence of feedback, but it is a bit too simple even for a basic device such as the thermostat. Actually, because lags and delays exist in the heating plant and sensor, a good thermostat does a bit of anticipation, turning the plant off before the error actually changes sign. This avoids excessive temperature swings and cycling of the heating plant.

This modification illustrates that, even in simple cases, good control system design is not entirely trivial. It must take into account the dynamic behavior of the object being controlled in order to do a good job. The more complex the dynamic behavior, the more elaborate the modifications. In fact, the development of a thorough theoretical understanding of the relationship between dynamic behavior and good controllers constitutes the most significant intellectual accomplishment of the control community, and the codification of this understanding into powerful computer aided engineering design tools makes all modern control systems possible.

There are many other control system examples, of course, that have developed over the years with progressively increasing levels of sophistication and impact. An early system with broad public exposure was the “cruise control” option introduced on automobiles in 1958 (see Figure 1.4b). With cruise control, ordinary people experienced the dynamic behavior of closed loop feedback systems in action—the slowdown error as the system climbs a grade, the gradual reduction of that error due to integral action in the controller, the small (but unavoidable) overshoot at the top of the climb, etc. More importantly, by experiencing these systems operating reliably and robustly, the public learned to trust and accept feedback systems, permitting their increasing proliferation all around us. Later control systems on automobiles have had more concrete impact, such as emission controls and fuel metering systems that have achieved major reductions of pollutants and increases in fuel economy.

In the industrial world, control systems have been key enabling technologies for everything from factory automation (starting with numerically controlled machine tools), to process control in oil refineries and chemical plants, to integrated circuit manufacturing, to power generation and distribution. They now also play critical roles in the routing of messages across the Internet (TCP/IP) and in power management for wireless communication systems.



Figure 1.5: The F-18 aircraft, one of the first production military fighters to use “fly-by-wire” technology, and the X-45 (UCAV) unmanned aerial vehicle. Photographs courtesy of NASA Dryden Flight Research Center.

Aerospace and Transportation

Aerospace and transportation encompasses a collection of critically important application areas where control is a key enabling technology. These application areas represent a significant part of the modern world’s overall technological capability. They are also a major part of its economic strength, and they contribute greatly to the well being of its people.

In aerospace, specifically, control has been a key technological capability tracing back to the very beginning of the 20th century. Indeed, the Wright brothers are correctly famous not simply for demonstrating powered flight—they actually demonstrated *controlled* powered flight. Their early Wright Flyer incorporated moving control surfaces (vertical fins and canards) and warpable wings that allowed the pilot to regulate the aircraft’s flight. In fact, the aircraft itself was not stable, so continuous pilot corrections were mandatory. This early example of controlled flight is followed by a fascinating success story of continuous improvements in flight control technology, culminating in the very high performance, highly reliable automatic flight control systems we see on modern commercial and military aircraft today.

Similar success stories for control technology occurred in many other application areas. Early World War II bombsights and fire control servo systems have evolved into today’s highly accurate radar-guided guns and precision-guided weapons. Early failure-prone space missions have evolved into routine launch operations, manned landings on the moon, permanently manned space stations, robotic vehicles roving Mars, orbiting vehicles at the outer planets, and a host of commercial and military satellites serving var-

ious surveillance, communication, navigation, and earth observation needs. Cars have advanced from manually tuned mechanical/pneumatic technology to computer-controlled operation of all major functions, including fuel injection, emission control, cruise control, braking, and cabin comfort.

Despite its many successes, the control needs of some engineered systems today and those of many in the future outstrip the power of current tools and theories. Design problems have grown from so-called “inner loops” in a control hierarchy (e.g. regulating a specified flight parameter) to various “outer loop” functions that provide logical regulation of operating modes, vehicle configurations, payload configurations, health status, etc. [3]. For aircraft, these functions are collectively called “vehicle management.” They have historically been performed by pilots or other human operators, but today that boundary is moving, and control systems are increasingly taking on these functions.

Today’s engineering methods for designing the upper layers of this hierarchy are far from formal and systematic. In essence, they consist of collecting long lists of logical if-then-else rules from experts, programming these rules, and simulating their execution in operating environments. Because few analytic techniques exist for reasoning about such systems, simulation is the primary tool for evaluation and only exhaustive simulation can guarantee good design properties. Clearly, this is an unacceptable circumstance—one where the strong system-theoretic background and the tradition of rigor held by the control community can make substantial contributions.

Another dramatic trend on the horizon is a change in dynamics to large collections of distributed entities with local computation, global communication connections, very little regularity imposed by the laws of physics, and no possibility of imposing centralized control actions. Examples of this trend include the national airspace management problem, automated highway and traffic management, and the command and control for future battlefields.

Information and Networks

The rapid growth of communication networks provides several major opportunities and challenges for control. Although there is overlap, we can divide these roughly into two main areas: control of networks and control over networks.

Control of networks is a large area, spanning many topics, including congestion control, routing, data caching, and power management. Several features of these control problems make them very challenging. The dominant feature is the extremely large scale of the system; the Internet is




Figure 1.6: UUNET network backbone for North America. Figure courtesy of WorldCom.

probably the largest feedback control system man has ever built. Another is the decentralized nature of the control problem: local decisions must be made quickly and based only on local information. Stability is complicated by the presence of varying time lags, as information about the network state can only be observed or relayed to controllers after a delay, and the effect of a local control action can be felt throughout the network only after substantial delay. Uncertainty and variation in the network, through network topology, transmission channel characteristics, traffic demand, available resources, and the like, may change constantly and unpredictably. Other complicating issues are the diverse traffic characteristics—in terms of arrival statistics at both the packet and flow time scales—and the different requirements for quality of service that the network must support.

Resources that must be managed in this environment include computing, storage and transmission capacities at end hosts and routers. Performance of such systems is judged in many ways: throughput, delay, loss rates, fairness, reliability, as well as the speed and quality with which the network adapts to changing traffic patterns, changing resource availability, and changing network congestion.

While the advances in information technology to date have led to a global Internet that allows users to exchange information, it is clear that the next

phase will involve much more interaction with the physical environment and the increased use of control over networks. Networks of sensor and actuator nodes with computational capabilities, connected wirelessly or by wires, can form an orchestra that controls our physical environment. Examples include automobiles, smart homes, large manufacturing systems, intelligent highways and networked city services, and enterprise-wide supply and logistics chains. Thus, this next phase of the information technology revolution is the convergence of communications, computing, and control.

As existing networks continue to build out, and network technology becomes cheaper and more reliable than fixed point-to-point connections, even in small localized systems, more and more control systems will operate over networks. We can foresee sensor, actuator, diagnostic, and command and coordination signals all traveling over data networks. The estimation and control functions can be distributed across multiple processors, also linked by data networks. For example, smart sensors can perform substantial local signal processing before forwarding relevant information over a network.

Current control systems are almost universally based on synchronous, clocked systems, so they require communication networks that guarantee delivery of sensor, actuator, and other signals with a known, fixed delay. Although current control systems are robust to variations that are included in the design process (such as a variation in some aerodynamic coefficient, motor constant, or moment of inertia), they are not at all tolerant of (unmodeled) communication delays or dropped sensor or actuator packets. Current control system technology is based on a simple communication architecture: all signals travel over synchronous dedicated links, with known (or worst-case bounded) delays and no packet loss. Small dedicated communication networks can be configured to meet these demanding specifications for control systems, but a very interesting question is whether we can develop a theory and practice for control systems that operate in a distributed, asynchronous, packet-based environment.

Robotics and Intelligent Machines

Robotics and intelligent machines refers to a collection of applications involving the development of machines with human-like behavior. Whereas early robots were primarily used for manufacturing, modern robots include wheeled and legged machines capable of competing in robotic competitions and exploring planets, unmanned aerial vehicles for surveillance and combat, and medical devices that provide new capabilities to doctors. Future applications will involve both increased autonomy and increased interaction with




Figure 1.7: The Mars Sojourner and Sony AIBO Entertainment Robot. Photographs courtesy of Jet Propulsion Laboratory and Sony.

humans and with society. Control is a central element in all of these applications and will be even more important as the next generation of intelligent machines are developed.

The goal of cybernetic engineering, already articulated in the 1940s and even before, has been to implement systems capable of exhibiting highly flexible or “intelligent” responses to changing circumstances. In 1948, the MIT mathematician Norbert Wiener gave a widely read account of cybernetics [26]. A more mathematical treatment of the elements of engineering cybernetics was presented by H.S. Tsien at Caltech in 1954, driven by problems related to control of missiles [25]. Together, these works and others of that time form much of the intellectual basis for modern work in robotics and control.

Two accomplishments that demonstrate the successes of the field are the Mars Sojourner robot and the Sony AIBO Entertainment Robot, shown in Fig. 1.7. Sojourner successfully maneuvered on the surface of Mars for 83 days starting in July 1997 and sent back live pictures of its environment. The Sony AIBO robot debuted in June of 1999 and was the first “entertainment” robot to be mass marketed by a major international corporation. It was particularly noteworthy because of its use of AI technologies that allowed it to act in response to external stimulation and its own judgment.

It is interesting to note some of the history of the control community in robotics. The IEEE Robotics and Automation Society was jointly founded in the early 1980s by the Control Systems Society and the Computer Society, indicating the mutual interest in robotics by these two communities. Un-

fortunately, although many control researchers were active in robotics, the control community did not play a leading role in robotics research throughout much of the 1980s and 90s. This was a missed opportunity, since robotics represents an important collection of applications that combine ideas from computer science, artificial intelligence, and control. New applications in (unmanned) flight control, underwater vehicles, and satellite systems are generating renewed interest in robotics, and many control researchers are becoming active in this area.

Despite the enormous progress in robotics over the last half century, the field is very much in its infancy. Today's robots still exhibit extremely simple behaviors compared with humans, and their ability to locomote, interpret complex sensory inputs, perform higher level reasoning, and cooperate together in teams is limited. Indeed, much of Wiener's vision for robotics and intelligent machines remains unrealized. While advances are needed in many fields to achieve this vision—including advances in sensing, actuation, and energy storage—the opportunity to combine the advances of the AI community in planning, adaptation, and learning with the techniques in the control community for modeling, analysis, and design of feedback systems presents a renewed path for progress.

Materials and Processing

The chemical industry is among the most successful industries in the United States, producing \$400 billion of products annually and providing over one million U.S. jobs. Having recorded a trade surplus for 40 consecutive years, it is the country's premier exporting industry: chemical industry exports totaled \$72.5 billion in 2000, accounting for more than 10% of all U.S. exports, and generated a record trade surplus in excess of \$20 billion in 1997.

Process manufacturing operations will require a continual infusion of advanced information and process control technologies if the chemical industry is to maintain its global ability to deliver products that best serve the customer reliably at the lowest cost. In addition, several new technology areas are being explored that will require new approaches to control to be successful. These range from nanotechnology in areas such as electronics, chemistry, and biomaterials to thin film processing and design of integrated microsystems to supply chain management and enterprise resource allocation. The payoffs for new advances in these areas are substantial, and the use of control is critical to future progress in sectors from semiconductors to pharmaceuticals to bulk materials.

There are several common features within materials and processing that




Figure 1.8: Intel Pentium IV wafer and die. Photographs courtesy of Intel.

pervade many of the applications. Modeling plays a crucial role, and there is a clear need for better solution methods for multidisciplinary systems combining chemistry, fluid mechanics, thermal sciences, and other disciplines at a variety of temporal and spatial scales. Better numerical methods for traversing these scales and designing, controlling, and optimizing under uncertainty are also needed. And control techniques must make use of increased in situ measurements to control increasingly complex phenomena.

In addition to the continuing need to improve product quality, several other factors in the process control industry are drivers for the use of control. Environmental statutes continue to place stricter limitations on the production of pollutants, forcing the use of sophisticated pollution control devices. Environmental safety considerations have led to the design of smaller storage capacities to diminish the risk of major chemical leakage, requiring tighter control on upstream processes and, in some cases, supply chains. And large increases in energy costs have encouraged engineers to design plants that are highly integrated, coupling many processes that used to operate independently. All of these trends increase the complexity of these processes and the performance requirements for the control systems, making the control system design increasingly challenging.

As in many other application areas, new sensor technology is creating new opportunities for control. Online sensors—including laser backscattering, video microscopy, ultraviolet, infrared, and Raman spectroscopy—are becoming more robust and less expensive and are appearing in more manufacturing processes. Many of these sensors are already being used by current process control systems, but more sophisticated signal processing and con-

control techniques are needed to more effectively use the real-time information provided by these sensors. Control engineers can also contribute to the design of even better sensors, which are still needed, for example, in the microelectronics industry. As elsewhere, the challenge is making use of the large amounts of data provided by these new sensors in an effective manner. In addition, a control-oriented approach to modeling the essential physics of the underlying processes is required to understand fundamental limits on observability of the internal state through sensor data.

Feedback in Nature²

Many cutting edge problems in the natural sciences involve understanding aggregate behavior in complex large-scale systems. This behavior “emerges” from the interaction of a multitude of simpler systems, with intricate patterns of information flow. Representative examples can be found in fields ranging from embryology to seismology. Researchers who specialize in the study of specific complex systems often develop an intuitive emphasis on analyzing the role of feedback (or interconnection) in facilitating and stabilizing aggregate behavior, and it is often noted that one can only have hope of deep understanding if it is somehow possible for theories of collective phenomenology to be robust to inevitable uncertainties in the modeling of fine-scale dynamics and interconnection.

While sophisticated theories have been developed by domain experts for the analysis of various complex systems, the development of rigorous methodology that can discover and exploit common features and essential mathematical structure is just beginning to emerge. The range of applications for which tools in control, dynamics, and systems (CDS) are relevant is expanding rapidly as advances in science and technology create new understanding of the underlying dynamics and the importance of feedback in a wide variety of natural and technological systems, and new sensors and actuators allow investigation and manipulation of phenomena at heretofore unimagined levels of detail. We briefly highlight four application areas here.

Biological Systems. At a variety of levels of organization—from molecular to cellular to organismal to populational—biology is becoming more accessible to approaches that are commonly used in engineering: mathematical modeling, systems theory, computation, and abstract approaches to synthesis. Conversely, the accelerating pace of discovery in biological science

²This material was contributed by the Center for Dynamics and Feedback in Nature, www.cdfn.caltech.edu.




Figure 1.9: The wiring diagram of the growth signaling circuitry of the mammalian cell [11].

is suggesting new design principles that may have important practical applications in man-made systems. This synergy at the interface of biology and engineering offers unprecedented opportunities to meet challenges in both areas. The principles of control are central to many of the key questions in biological engineering and will play an enabling role in the future of this field.

A major theme currently underway in the biology community is the science of reverse (and eventually forward) engineering of biological control networks (such as the one shown in Figure 1.9). There are a wide variety of biological phenomena that provide a rich source of examples for control, including gene regulation and signal transduction; hormonal, immunological, and cardiovascular feedback mechanisms; muscular control and locomotion; active sensing, vision, and proprioception; attention and consciousness; and population dynamics and epidemics. Each of these (and many more) provide opportunities to figure out what works, how it works, and what we can do to affect it.

Ecosystems. In contrast to individual cells and organisms, emergent properties of aggregations and ecosystems inherently reflect selection mech-

organisms which act on multiple levels, and primarily on scales well below that of the system as a whole. Because ecosystems are complex, multiscale dynamical systems, they provide a broad range of new challenges for modeling and analysis of feedback systems. Recent experience in applying CDS tools to bacterial networks suggests that much of the complexity of these networks is due to the presence of multiple layers of feedback loops that provide robust functionality to the individual cell. Yet in other instances, events at the cell level benefit the colony at the expense of the individual. Systems level analysis can be applied to ecosystems with the goal of understanding the robustness of such systems and the extent to which decisions and events affecting individual species contribute to the robustness and/or fragility of the ecosystem as a whole.

Quantum Systems. While organisms and ecosystems have little to do with quantum mechanics in any traditional scientific sense, complexity and robustness issues very similar to those described above can be identified in the modern study of quantum systems. In large part, this sympathy arises from a trend towards wanting to control quantum dynamics and to harness it for the creation of new technological devices. At the same time, physicists are progressing from the study of elementary quantum systems to the study of large aggregates of quantum components, and it has been recognized that dynamical complexity in quantum systems increases exponentially faster with system size than it does in systems describable by classical (macroscopic) physics. Factors such as these are prompting the physics community to search broadly for new tools for understanding robust interconnection and emergent phenomena.

Modern scientific research is rapidly evolving a field of *quantum engineering*. Driven by technology goals in areas such as quantum information processing, nano-electromechanical sensing, chemical synthesis, trace gas detection, and ultrahigh-bandwidth optical communication, researchers are beginning to formulate strategies for achieving robust performance with physical devices or systems in which quantum mechanical effects are prominent. Mathematical tools from CDS for analysis and synthesis could have a profound impact on activities of this kind. A schematic diagram of a modern quantum control experiment is shown in Figure 1.10a.

Environmental Science. It is now indisputable that human activities have altered the environment on a global scale. Problems of enormous complexity challenge researchers in this area and first among these is to understand the feedback systems that operate on the global scale. One of the challenges in developing such an understanding is the multiscale nature of the problem, with detailed understanding of the dynamics of microscale

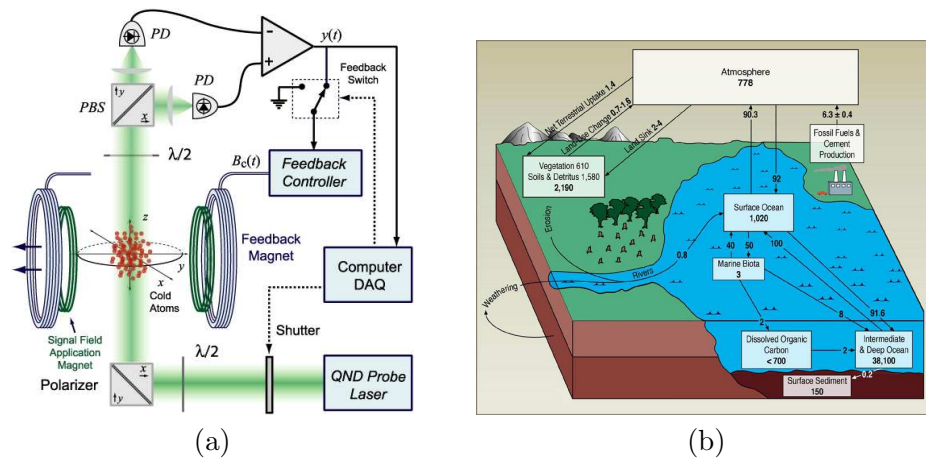


Figure 1.10: Examples of feedback systems in nature: (a) quantum control system and (b) global carbon cycle.

phenomena such as microbiological organisms being a necessary component of understanding global phenomena, such as the carbon cycle illustrated Figure 1.10b.

Other Areas

The previous sections have described some of the major application areas for control. However, there are many more areas where ideas from control are being applied or could be applied. Some of these include: economics and finance, including problems such as pricing and hedging options; electromagnetics, including active electromagnetic nulling for stealth applications; molecular, quantum, and nanoscale systems, including design of nanostructured materials, precision measurement, and quantum information processing; energy systems, including load distribution and power management for the electrical grid; and manufacturing systems, including supply chains, resource management and scheduling, and factory automation.

1.4 Feedback Principles

Feedback is a powerful idea which, as we have seen, is used extensively in natural and technical systems. The principle of feedback is very simple: base correcting actions on the difference between desired and actual performance. In engineering, feedback has been rediscovered and patented many times in




Figure 1.11: A simple feedback system for controlling the speed of a vehicle.

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 the subsequent chapters.

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 actual 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 1.11. In this system, the speed of a vehicle is controlled by adjusting the amount of gas flowing to the engine. A simple “proportional plus integral” feedback is used to make the amount of gas depend on both the error between the current and 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 2), the steady state speed of the vehicle always approaches the desired speed. Thus the performance of the system is robust with respect to this uncertainty.

Another early example of the use of feedback to provide robustness was the negative feedback amplifier of Black. When telephone communications were developed, amplifiers were used to compensate for signal attenuation in long lines. The 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 development of line amplifiers for a long time. A major breakthrough was Black's invention of the feedback amplifier in 1927. Black used negative feedback which reduces the gain but makes the amplifier very insensitive to variations in tube characteristics. This invention made it possible to build stable amplifiers with linear characteristics despite 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 by Wilbur Wright to the Western Society of Engineers in 1901 [18], 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, Kitty Hawk, that was unstable but maneuverable.



Figure 1.12: The Curtiss-Sperry E in 1912 (left) and a closeup of the Sperry Autopilot (right).

Kitty Hawk 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 with Kitty Hawk in 1905.

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 is the first use of feedback in aeronautical engineering and Sperry won a prize in a competition for the safest airplane in Paris in 1912. Figure 1.12 shows the Curtiss-Sperry seaplane and the Sperry autopilot. The autopilot is a good example of how feedback can be used to stabilize an unstable system.

Higher Levels of Automation

A major trend in the use of feedback is its use in higher and higher levels of automation and decision making. A good example of this is the DARPA Grand Challenge, a competition sponsored by the US government to build a vehicle that could autonomously drive itself across the desert. Caltech competed in the 2004 Grand Challenge using a heavily modified 1996 Chevy

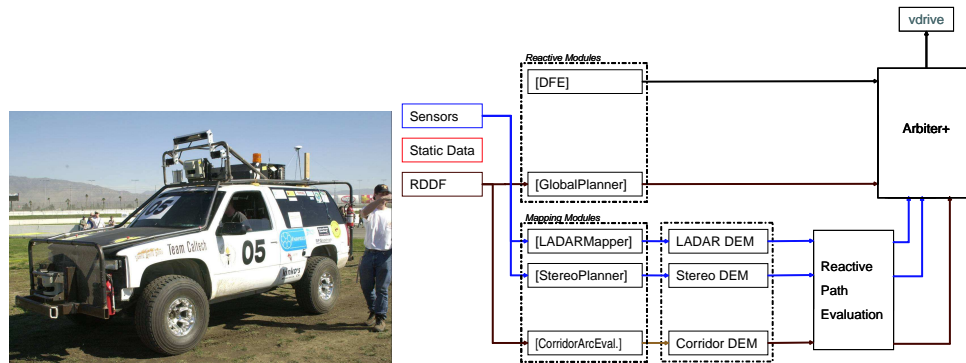


Figure 1.13: (a) Team Caltech’s entry in the 2004 DARPA Grand Challenge. (b) Arbiter Control architecture for Caltech’s entry.

Figure 1.14: Control architecture for Caltech’s entry in the DARPA Grand Challenge

Tahoe. It was fully automated, including electronically controlled steering, throttle, brakes, transmission, and ignition. Its sensing systems included 4 black and white cameras sampling at 30 Hz (arranged in two stereo pairs), 1 color camera, 2 LADAR (laser radar) units scanning at 10 Hz, and a GPS/IMU package capable of providing full position and orientation at 400 Hz time resolution. Computational resources included 8 high speed desktop computers connected together through a 1 Gbs ethernet switch. A picture of the vehicle is shown in Figure 1.13a.

Custom software was used to control the vehicle. The basic system architecture is shown in Figure 1.13b. An arbiter-based navigation framework, similar to that used by the Jet Propulsion Laboratory (JPL) for its Mars rover missions, was chosen due to its ability to allow multiple teams to work in parallel. In this architecture, a collection of *voters* evaluate the “goodness” of a set of arcs that are determined by setting the steering wheel to a pre-specified condition. Along each of these arcs, the voter uses its local data (for example, the terrain map generated by stereo vision) to determine which arcs can be followed for the longest distance. Arcs that have long runs are given high goodness values. A desired velocity along the arc is also commanded.

The arbiter takes the votes from multiple sources (six were used in the race) and combines them to determine which direction and at what speed the vehicle should be commanded to move. A pair of cameras used stereo

imaging to build up a map of the terrain right in front of the vehicle and looked for arcs that did not involve significant changes in height. Additional tests were made to make sure that the entire width of the vehicle would stay on terrain that was as flat as possible. In addition, two LADAR (laser radar) units were mounted on the vehicle: one at bumper height, pointing horizontally, and the other mounted on the roof, aimed at a line approximately 20 meters in front of the vehicle. The scans from these units gave the distance to points on the scan line and this was integrated into a terrain map of the same form as that used for stereo vision. A waypoint following voter kept track of the current waypoint target for the vehicle and generated votes that brought the vehicle onto the line connecting the waypoints. Similarly, a corridor evaluator was used to stay inside the corridor (specified as part of the route definition). Finally, a dynamic feasibility evaluator was used to ensure that the vehicle did not attempt to make a sharp turn at high speed and to keep the roll angle of the vehicle from getting too high. It used only the current state of the vehicle to command motion that was safe from a dynamic perspective.

The software and hardware infrastructure that was developed enabled the vehicle to traverse relatively long distances at substantial speeds. In the qualification event for the grand challenge, Caltech's vehicle was able to complete the 1.25 mile course twice while avoiding a variety of obstacles. On the first successful run, the vehicle avoided all obstacles and completed in the course in just under 20 minutes (3.8 mph average speed). In the second run, the vehicle completed the course in just under 12 minutes (6.5 mph average speed), but was unable to stop quickly enough when a moving obstacle was put in its path at the end of the course. In both runs, the vehicle relied heavily on its sensing systems to avoid obstacles and navigate through openings in fences and barriers. No pre-planned trajectories were used.

In the race, Caltech's vehicle traversed 1.3 miles, but deviated from the specified corridor several times due to misinterpretation of sensor signals. During one of these deviations, the vehicle pushed up against a barbed wire fence in a way that stopped its motion, ending its journey. One of the clear lessons learned from the competition is the need to make the system more adaptable. A more sophisticated architecture would have several features to improve the overall performance. For example, a global reasoning system that could learn about poor choices in path planning and ensure that they are not chosen a second time might have helped the system avoid getting stuck on barbed wire.

The Magic of Feedback

Feedback has many interesting and useful properties. It makes it possible to design precise systems from imprecise components and to make physical variables 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.

A consequence of the nice properties of feedback is that it has had major impact on man-made systems. Drastic improvements have been achieved when feedback has been applied to an area where it has not been used before.

Drawbacks of Feedback

While feedback has many advantages, it also has some drawbacks. Chief among these is the possibility for 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 a feedback instability, something that we obviously want to avoid. This is tricky because of the uncertainty that feedback was introduced to compensate for: not only must we design the system to be stable with the nominal system we are designing for, but it must remain stable under all possible perturbations of the dynamics.

Another potential drawback of control is the complexity of embedding a control system into a product. While the cost of sensing, computation, and (to a lesser extent) actuation has decreased dramatically in the past few decades, the fact remains that control systems are often very 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 only be met 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, although the reliability has risen to the point that this is no longer a major issue.

Feedforward

When using feedback that there must be an error before corrective actions are taken. Feedback is thus reactive. In some circumstances it is possible to measure a disturbance before it enters the system and this information can 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 to shape 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 it 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 pure feedforward strategy corresponds to running a company based on extensive strategic planning while a feedback strategy corresponds to a pure reactive approach. The experience in control indicates that it is often advantageous to combine feedback and feedforward. Feedforward is particularly useful when disturbances can be measured or predicted. A typical example is in chemical process control where disturbances in one process may be due to processes upstream. The correct balance of the approaches requires insight and understanding of their properties.

1.5 Control Tools

Because of the wide use of feedback in a variety of applications, there has been substantial mathematical development in the field of control theory. We briefly describe some of the tools and concepts here.

System Modeling

Models play an essential role in analysis and design of feedback systems. Several sophisticated tools have been developed to build models that are suited for control.

Models can often be obtained from first principles and there are several modeling tools in special domains such as electric circuits and multibody systems. Since control applications cover such a wide domains it is also desirable to have modeling tools that cut across traditional discipline bound-

aries. Such modeling tools are now emerging, with the models obtained by cutting a system into subsystems and writing equations for balance of mass, energy and momentum and constitutive equations that describe material properties for each subsystem. Object oriented programming can be used very effectively to organize the work and extensive symbolic computation can be used to simplify the equations. Models and components can then be organized in libraries for efficient reuse. Modelica is an example of a modeling tool of this type.

Modeling from input/output data or system identification is another approach to modeling that has been developed in control. Direct measurement of the response to step input is commonly used in the process industry to tune proportional plus integral (PI) controllers. More accurate models can be obtained by measuring the response to sinusoidal signals, which is particularly useful for systems with fast response time. Control theory has also developed new techniques for modeling dynamics and disturbances. These include input/output representations of systems (how disturbances propagate through the system) and data-driven system identification techniques. The use of “forced response” experiments to build models of systems is well developed in the control field and these tools find application in many disciplines, independent of the use of feedback. A strong theory of modeling has also been developed, allowing rigorous definitions of model fidelity and comparisons to experimental data. The theory also tells how to select efficient forcing of the system. A number of model reduction techniques that make it possible to simplify complex models for specific purposes have also been developed.

The impact of modeling on engineering and control over the past 50 years has been profound. Today, entire products are designed using only models, with the first prototype being a fully working system. Doing this requires an enormous amount of infrastructure in simulation tools, detailed physical models, and experience in using models. Moving forward, this infrastructure becomes even more important as suppliers of components compete and sell their products based on detailed models of their systems which implement the specifications of their products sometimes before the system has even been built.

Analysis

Control theory has developed an extensive collection of theory and software tools for analysis of feedback systems. These tools include stability analysis for both linear and nonlinear systems, performance measures for

input/output systems, and evaluation of robustness. For robustness analysis, the tools are particularly sophisticated in the case of linear dynamical systems, where it is possible to analyze the stability and performance of the system in the presence of external disturbances, parametric uncertainty (e.g., unknown values of parameters such as mass or reaction rates), and unmodeled dynamics. In the case of unmodeled dynamics, one can even reason about the performance of a system without knowing the precise behavior of every subsystem, a very powerful concept.

In almost all cases, the theory used for analysis of feedback systems is implemented in software tools that can be used in computer aided design environments. Two very common environments for this analysis are MATLAB and Mathematica. In both cases, toolboxes are available that implement the common calculations described in this text (and many more) and these have become indispensable aides for modern design. More sophisticated tools, capable of constructing and analyzing very complex hierarchical models, are also available for more discipline specific environments.

An important point to remember about systems analysis is that it relies on models to describe the behavior of the system. In the simplest case, these models are simulated to provide information about how the system will respond to a given set of initial conditions, disturbances, and environment. But modern computational tools can do much more than just simulate the system, and can provide very sophisticated analyses that answer questions about the parametric behavior of systems, tradeoffs between different design factors, and even provide certificates (proofs) of performance. These tools are very well developed for linear systems, but recent advances in nonlinear analysis are quickly extending these results to larger and larger classes of systems.

Synthesis

In addition to tools for analysis of feedback systems, theory and software has also been developed for synthesizing feedback systems. As an example, consider the control system depicted in Figure 1.3. Given models of the process to be controlled, it is today possible to automatically synthesize a control algorithm that satisfies a given performance specification. A typical approach to doing this would involve first obtaining a nonlinear model for the process that captured the key behavior that one was interested in. This model would then be “linearized” around a desired operating point (this is described in Chapter 4) and a performance condition specified (usually as a function that one wishes to minimize). Given the linear model and the

control specification, a feedback law can be computed that is the optimal law for the given specification.

Modern implementation environments continue by allowing this control algorithm to be “autocoded”, so that programming code implementing the control logic is automatically generated. This has the advantage that the code generator can be carefully verified so that the resulting algorithm is correct (as opposed to hand coding the algorithm, which can lead to errors in translation). This autocoded control logic is then downloaded to a dedicated computing platform with the proper interfaces to the hardware and implemented. In addition to simple feedback algorithms, most modern control environments allow complex decision-making logic to be constructed via block diagrams and this is also automatically compiled and downloaded to a computer that is connected to the hardware.

This mechanism for generating control algorithms directly from specifications has vastly improved the productivity of control engineers and is now standard practice in many application areas. It also provides a clear framework in which new advances, such as real-time, optimization-based control, can be transitioned to applications quickly and efficiently through the generation of standard toolboxes.

1.6 Further Reading

The material in this section draws heavily from the report of the Panel on Future Directions on Control, Dynamics, and Systems [20] and a recent textbook by one of the authors [?]. Several reports and papers have highlighted future directions of research, including a 1988 panel report chaired by Fleming [10] and a 1987 *IEEE Transactions on Automatic Control* article [16], both of which provided a roadmap for many of the activities of the last decade and continue to be relevant. More recently, the European Commission sponsored a workshop on future control systems [9] and several other more focused workshops have been held over the last several years [1, 2, 22, 23]. Several recent papers and reports highlighted successes of control [24] and new vistas in control [7, 15].

A fascinating examination of some of the early history of control in the United States has been written by Mindell [19]. Additional historical overviews of the field have been prepared by Bennett [5, 6] and Mayr [17], which go back as early as the 1800s. A popular book that describes many control concepts across a wide range of disciplines is “Out of Control” by Kelly [14].

1.7 Exercises

1. Identify 5 feedback systems that you encounter in your everyday environment. For each system, identify the sensing mechanism, actuation mechanism, and control law. Describe the uncertainty that the feedback system provides robustness with respect to and/or the dynamics that are changed through the use of feedback. At least one example should correspond to a system that comes from your own discipline or research activities.
2. Perform the following experiment and explain your results: Holding your head still, move your right or left hand back and forth 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 your hand. Now hold your hand still and move your head back and forth, once again recording how quickly you can move before losing track.
3. Balance yourself on one foot with your eyes closed for 15 seconds. 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 the diagram (unless you are an android reading this in the far future).

Bibliography

- [1] P. Antsaklis, T. Basar, R. DeCarlo, N. H. McClamroch, M. Spong, and S. Yurkovich, editors. *NSF/CSS Workshop on New Directions in Control Engineering Education*. National Science Foundation and IEEE Control Systems Society, 1998. Available at <http://robot0.ge.uiuc.edu/~spong/workshop>.
- [2] M. Athans, J. J. S. Sentieiro, and M. A. Dahleh. Workshop on future directions in systems and control theory. Cascais, Portugal, June 2000.
- [3] S. Banda, J. C. Doyle, R. M. Murray, J. Paduano, J. Speyer, and G. Stein. Research needs in dynamics and control for uninhabited aerial vehicles. Panel Report, November 1997. Available at <http://www.cds.caltech.edu/~murray/notes/uav-nov97.html>.
- [4] M. B. Barron and W. F. Powers. The role of electronic controls for future automotive mechatronic systems. *IEEE Transactions on Mechatronics*, 1(1):80–89, 1996.
- [5] S. Bennett. *A History of Control Engineering: 1800–1930*. Peter Peregrinus, 1986.
- [6] S. Bennett. *A History of Control Engineering: 1930–1955*. Peter Peregrinus, 1986.
- [7] R. W. Brockett. New issues in the mathematics of control. In B. Engquist and W. Schmid, editors, *Mathematics Unlimited—2001 and Beyond*, pages 189–220. Springer Verlag, 2000.
- [8] F. Dyson. A meeting with enrico fermi. *Nature*, (247):297, 2004.
- [9] European Commission. Workshop on future and emerging control systems, November 2000. Available at ftp://ftp.cordis.lu/pub/ist/docs/ka4/report_controlws.pdf.

- [10] W. H. Fleming, editor. *Future Directions in Control Theory: A Mathematical Perspective*. SIAM, 1988.
- [11] D. Hanahan and R. A. Weinberg. The hallmarks of cancer. *Cell*, 100:57–70, 2000.
- [12] M. B. Hoagland and B. Dodson. *The Way Life Works*. Times Books, 1995.
- [13] F. Rowsone Jr. What it's like to drive an auto-pilot car. *Popular Science Monthly*, April 1958. Available at <http://www.imperialclub.com/ImFormativeArticles/1958AutoPilot>.
- [14] K. Kelly. *Out of Control*. Addison-Wesley, 1994. Available at <http://www.kk.org/outofcontrol>.
- [15] P. R. Kumar. New technological vistas for systems and control: The example of wireless networks. *Control Systems Magazine*, 21(1):24–37, 2001.
- [16] A. H. Levis, S. I. Marcus, W. R. Perkins, P. Kokotovic, M. Athans, R. W. Brockett, and A. S. Willsky. Challenges to control: A collective view. *IEEE Transactions on Automatic Control*, AC-32(4):275–285, 1987.
- [17] O. Mayr. *The Origins of Feedback Control*. MIT Press, 1970.
- [18] M. W. McFarland, editor. *The Papers of Wilbur and Orville Wright*. McGraw Hill, 1953.
- [19] D. A. Mindel. *Between Human and Machine: Feedback, Control, and Computing Before Cybernetics*. Johns Hopkins University Press, 2002.
- [20] R. M. Murray, editor. *Control in an Information Rich World: Report of the Panel on Future Directions in Control, Dynamics and Systems*. SIAM, 2003. To Appear. Available at <http://www.cds.caltech.edu/~murray/cdspanel>.
- [21] National Research Council. *Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers*. National Academy Press, 2001.
- [22] National Science Foundation. NSF workshop for high school teachers of mathematics and science. American Control Conference,

June 2000. Available at http://www.math.ukans.edu/ksacg/nsf_hswkshop.html.

- [23] National Science Foundation. NSF workshop on dynamic data-driven application systems, March 2000. Available at <http://www.cise.nsf.gov/eia/dddas>.
- [24] H. Nijmeijer and J. M. Schumacher. Four decades of mathematical system theory. In J. W. Polderman and H. L. Trentelman, editors, *The Mathematics of Systems and Control: From Intelligent Control to Behavioral Systems*, pages 73–83. Univ. of Groningen, 1999.
- [25] H. S. Tsien. *Engineering Cybernetics*. McGraw-Hill, 1954.
- [26] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. John Wiley, 1948.