
Feedback Systems

An Introduction for Scientists and Engineers
SECOND EDITION

Karl Johan Åström
Richard M. Murray

Version v3.0i (2017-08-28)

This is the electronic edition of *Feedback Systems* and is available from <http://www.cds.caltech.edu/~murray/FBS>. 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 Fifteen

Architecture and System Design

The architects two most important tools are the eraser in the drafting room and the wrecking bar on the site.

Frank Lloyd Wright [?].

In this chapter we place the relatively simple feedback loops that have been the focus of the previous chapters in the context of overall system design. We outline a typical design process and discuss the role of architecture and how it can be approached from top-down and bottom-up perspectives. Interaction and adaptation are then reviewed and the chapter ends with a brief overview of control design in some major industrial fields.

15.1 Introduction

So far we have dealt with relatively simple feedback systems. We will now give a glimpse of how they appear as components of real systems and how they are designed. All control systems have sensors, actuators, communication, computers, and operator interfaces, but they can have dramatically different sizes and shapes and very different user communities. It is surprising that such a variety of systems can be analyzed and designed using the same engineering framework.

The system to be controlled is often designed before control is considered. There are, however, significant advantages to designing a process and its control system jointly, so-called *co-design*. Care can be taken to ensure that the system is easy to control, for example by avoiding non-minimum phase dynamics. Time delays can be avoided by proper positioning of sensing and actuation. Use of feedback gives an extra degree of freedom to the designer; an extreme example is that a system can be made more maneuverable by making it unstable and then stabilizing it with a controller. The system itself and its physical and operational environment are key elements together with requirements, analysis, and testing.

Architecture, from the Greek word *αρχιτεκτων* (*αρχι* chief and *τεκτων* builder, carpenter, mason), is the process of planning, designing, and constructing buildings and other objects. It is also used to describe the structure of practically anything. In the context of control systems, the elements consist of the process, sensors, actuators, computers, communication devices, human machine interfaces, algorithms, and software. The control system interacts with the operational environment, it observes the process by sensors, and it interacts with the process through actuators and with the users through a range of interfaces. Architecture

describes how the system components are connected and how they interact. There is a growing awareness that architecture is important in all engineering fields and today we have software, hardware, and systems architects.

15.2 System and Control Design

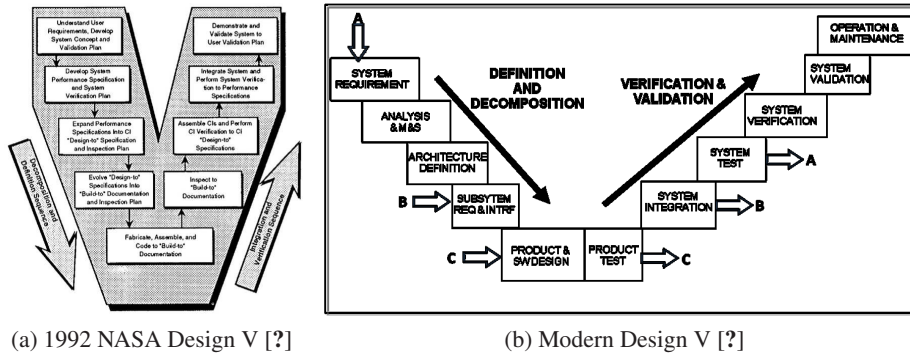
System design starts by developing an understanding of the system and its environment. It includes analysis of static and dynamic properties of the physical system and its sensors and actuators, bounds for safe operation, and characterization of the nature of the disturbances and the users of the system. There are a wide range of problems. Sometimes the process is given *a priori* and the task is to design a controller for a given process. In other cases the process and the controller are designed jointly. Co-design has many advantages because performance could be optimized. Sometimes it is an enabler, as was illustrated by the Wright Flyer, which was discussed in Section 1.5. We quote from the 43rd Wilbur Wright Memorial Lecture by Charles Stark Draper [69]:

The Wright Brothers rejected the principle that aircraft should be made inherently so stable that the human pilot would only have to steer the vehicle, playing no part in stabilization. Instead they deliberately made their airplane with negative stability and depended on the human pilot to operate the movable surface controls so that the flying system—pilot and machine—would be stable. This resulted in increased maneuverability and controllability.

If the stabilization of an unstable airframe is done by an automatic control system, there are very strong requirements on the reliability of the control system. Design of the X-29, which was discussed in Example 14.4, is a similar case. A more recent example, which deals with difficulties caused by insufficient actuator authority, is presented in [?]. It was attempted to reduce the risk for rotating stall in a jet engine by feedback, but actuators with the required bandwidth were not available. Analysis showed that the problem could instead be alleviated by introducing small asymmetries in the turbine.

Design of complex systems is a major effort where many people are involved. A variety of methods have been developed for efficient design. The so-called *V-model* in Figure 15.1, dating back to NASA's Apollo program, is a design pattern for hardware and software [?]. It appears in many different forms: one version is part of the official project management methodology of software for the German government [?].

The V represents the different steps in product development, progressing down the left side and up the right side of the V. The left side shows how design proceeds from understanding of the process and the requirements to development of a systems architecture and partitioning into subsystems, and documentation of requirements and tests. The bottom captures fabrication of subsystems. The right hand side shows how the complete system is implemented and tested. Verification



(a) 1992 NASA Design V [?]

(b) Modern Design V [?]

Figure 15.1: The V-model for system design. The left side of the V represents the decomposition of requirements, and creation of system specifications. The right side represents the activities in implementation including validation (building the right thing) and verification (building it right).

(doing things right) and validation (doing the right thing) are important elements. Commissioned, operation, and upgrading are also part of the right side of the V.

The V model has proven useful, particularly for project management, because it provides a structured process with emphasis on requirements, documentation, reviews, and sign off. There has, however, been criticism of the V-model. A major drawback is the long loop from requirement to a final product. The loop can be shortened by the W-model which contains a parallel V with testing. It is particularly attractive to use mathematical models and hardware for virtual testing using hardware-in-the-loop simulation. It is also difficult to change requirements along the project, which is particularly serious for products with short lead time.

An alternative is the so-called *agile development* model, which has been driven by software developers for products with short time to market, where requirements change and close interaction with customers is required. The method is characterized by the *Agile Manifesto* [?], which values individuals and interactions over processes and tool, working software over comprehensive documentation, customer collaboration over contract negotiation and responding to change over following a plan. When choosing a design methodology it is also important to keep in mind that products involving hardware are more difficult to change than software.

Control system design is a subpart of system design that includes many activities, starting with requirements and system modeling and ending with implementation, testing, commissioning, operation, and upgrading. In between are the important steps of detailed modeling, architecture selection, analysis, design, and simulation. The V-model used in an iterative fashion is well suited to control design, particular if it is supported by a tool chain that admits a combination of modeling, control design, and simulation. Testing is done iteratively at every step of the design using models of different granularity as virtual systems. Hardware in the loop simulations are also used when they are available. A scripting language is helpful to execute the design.




Figure 15.2: Requirements can be tested by injecting signals at test points δ_k and measuring responses at s_{ij} . Compare with Figure 12.1

Control system requirements are typically given by large and small signal behavior of the closed loop system. Large signal behavior is characterized by limitations in actuation power and its rate, small signal behavior is typically caused by measurement noise, friction and resolution of A/D and D/A converters. Requirements for control systems include the ability to deal with disturbances, robustness to process variations and uncertainty, and the ability to follow reference signals. Many control system specifications can be captured by linear models and they can be expressed in terms of properties of the Gangs of Four and Six, discussed in Section 12.1. Referring to the block diagram in Figure 15.2, load disturbance attenuation can be characterized by the transfer function G_{yv} from load disturbance v to process output y . Measurement noise w generates undesired control actions, the effect can be captured by transfer function G_{uw} from measurement noise w to control action u . Robustness to parameter variations and process uncertainty can be captured by the sensitivity functions S and T . Command signal response can be shaped independently of response to disturbances and robustness for systems with two degrees of freedom. It is characterized by the transfer functions TF and CSF . Systems with error feedback are more restricted because the response to command signals is characterized by the complementary transfer function T and a compromise must be made between command signal response and the other requirements.

Since many requirements are expressed in terms of properties of the transfer functions in the Gang of Six, it is important to measure these transfer functions on simulated models and on real hardware. To do this the system must be provided with test points for injecting and measure signals, as indicated by the dashed arrows in Figure 15.2. The transfer function G_{yv} , which characterizes response to load disturbances, can be found by injecting a signal at δ_1 and measuring the output s_{21} . Chirp signals are convenient for measuring frequency responses.

Models of the process and its environment can be obtained from physics, from experiments, or from a combination. Experiments are typically done by changing the control signal and measuring the response. The signals can range from simple step tests to signals that are designed to give optimal information with limited process perturbations. System identification methods and software provide useful

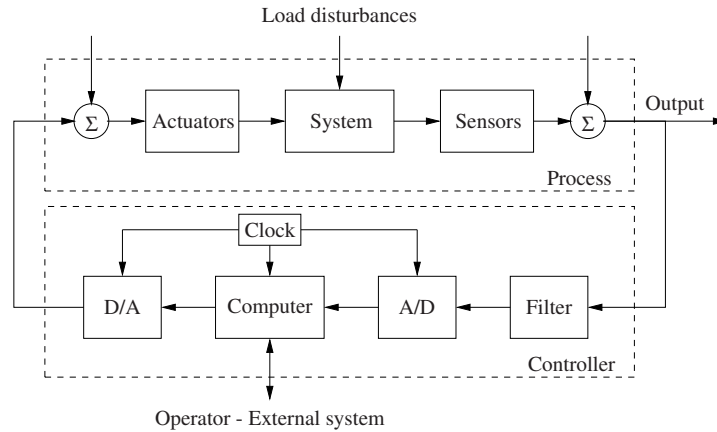


Figure 15.3: Schematic diagram of a control system with sensors, actuators, communications, computer, and interfaces.

tools. The models used typically have different fidelity, cruder in the beginning and more accurate as the design progresses.

A few standard design methods have been discussed in Chapters 7, 8, 11, and 12, but there are many more methods in the literature [85, 94]. Many design methods are based on linear models, however, when environmental conditions change significantly it is necessary to use gain scheduling, nonlinear control, or adaptation. Receding horizon control (also called model predictive control) is another common approach, especially useful when there are constraints on the inputs or states.

Today most control systems are implemented using computer control. Implementation then involves selection of hardware for signal conversion, communication, and computing. A block diagram of a system with computer control is shown in Figure 15.3. The overall system consists of sensors, actuators, analog-to-digital and digital-to-analog converters, and computing elements. The filter before the A/D converter is necessary to ensure that high frequency disturbances do not appear as low frequency disturbances after sampling because of aliasing. The operations of the system are synchronized by a clock.

Real-time operating systems that coordinate sensing, actuation, and computing have to be selected, and algorithms that implement the control laws must be generated. The sampling period and the anti-alias filter must be chosen carefully. Since a computer can only do basic arithmetic, the control algorithms have to be represented as difference equations. They can be obtained by approximating differential equations, as was illustrated in Section 8.5, but there are also design methods that automatically give controllers in the form of difference equations. Code can be generated automatically. It must also be ensured that computational delays and synchronization of algorithms do not create problems.

When the design is implemented and tested the system must be commissioned. This step may involve adjustment of controller parameters, and automatic tuning,

discussed in Section 11.3, can be very beneficial at this stage. During operation it is important to monitor the behavior of the system to ensure that requirements are still satisfied. It may be necessary to upgrade the system when it has been operating. Requirements may also be modified due to operational experiences.

It is highly desirable to have a suite of test programs that can be used throughout the design and operation stages to ensure that requirements are satisfied.

15.3 Top-Down Architectures

When system design is approached systematically using the design V as described in Section 15.2, it is natural to use a top-down procedure for control design starting with the requirements and decomposing the overall control design problem into an interlinked set of control problems at different layers of abstraction. At each layer of abstraction, we make assumptions about the interactions with the higher and lower layers in order to simplify the control design problem. In this section we give a brief introduction to some of the organizing principles of top-down architectures for control and its connections to some of the techniques described in the text, along with references to the literature for those interested in further details.

There are many other aspects of control architectures that are part of control systems design. These include such topics as cooperative control [?, ?], diagnostics and health monitoring [?, ?], fault recovery and system reconfiguration [?, ?], and game-theoretic approaches to control [?]. We focus here on a small subset of the problem, with an emphasis on multi-layer approaches to control.

Layered Architectures for Control

For complex control systems, it is often useful to break down the control problem into a hierarchy of control problems, each solved at a different layer of abstraction, as illustrated in Figure 15.4. Different types of specifications are used at each layer to determine the control functionality that will be implemented.

The specific abstraction layers in a control architecture depend on the problem domain. In Figure 15.4, we have used a decomposition that is common in many motion control problems, including robotics, self-driving cars, and flight control. Similar decompositions also appear in application domains such as manufacturing, process control, and computing systems. At the top layer of abstraction, we care about discrete modes of behavior, which could correspond to different phases of operation (takeoff, cruise, landing) or different environment assumptions (highway driving, city streets, parking lot). The next layer of abstraction reasons about trajectories of the system, often using an optimization-based approach. At this layer, we often take into account the constraints on the system operating state and inputs, as well as system-level descriptions of performance. Finally, at the lowest layer of the abstraction hierarchy we have the feedback control design that has been the main topic in this text thus far, where we may use a linearized model based on the current operating point (along a trajectory).




Figure 15.4: Layered decomposition of a control system.

Note that at each abstraction level we must consider not only the control design but also the way that sensory information is processed and represented. At the lowest levels of abstraction we may use individual sensor signals, perhaps with some filtering. As we move up the abstraction hierarchy, it will be necessary to fuse multiple measurements to obtain a more integrated view of the system state, with the highest level of abstraction often requiring sophisticated methods of reasoning about the state of the environment and the predicted interactions with other entities in that environment.

The architecture in Figure 15.4 is suitable for systems of moderate complexity where the users interact with the system by changing modes and references. More layers are used for complex systems with complicated user interaction. Batch control is a typical example, where a complex manufacturing system is used to control different batches of chemicals and where the material flow through the factory changes. In this case there are two additional layers—procedural control and coordination control—on top of those shown in Figure 15.4. Procedural control executes the sequence of actions necessary to carry out a process oriented task, such as charging a reactor with a specific amount and type of raw material and reporting the result. Coordination control directs, initiates, or modifies the execution of procedural control and the utilization of equipment entities.

In addition to these additional layers, a production facility typically operates in different modes: normal, maintenance, and manual. The maintenance mode has its own control algorithms and safety procedures to ensure that the system does not react in an unsafe manner during maintenance. The manual mode is typically used for equipment maintenance and debugging. Different parts of a manufacturing system can be in different modes. An example of an architecture for distributed

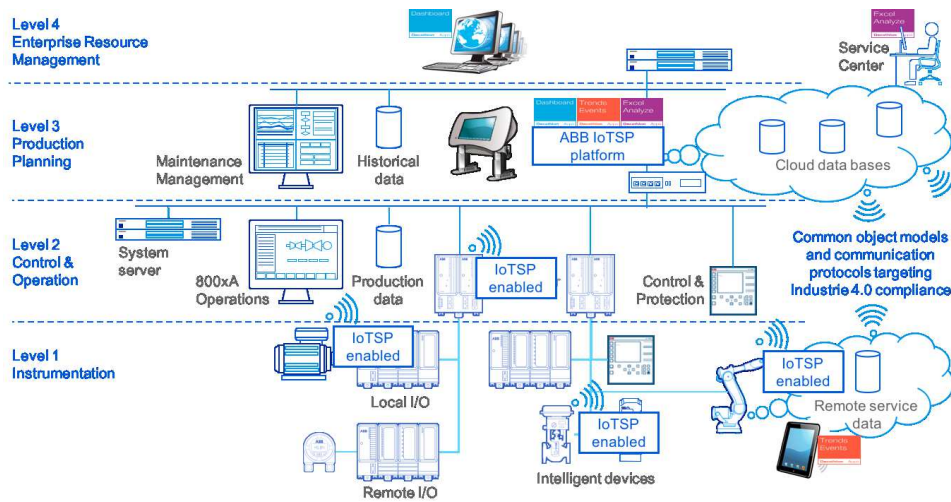



Figure 15.5: Functional architecture of process control system, implemented as a distributed control system (DCS). Figure courtesy of ABB, Inc.

control system (DCS), typical for complex manufacturing systems, is shown in Figure 15.5.

An important feature of many control systems architectures is the modularity of the control software, enabling parallel development of components and the ability to upgrade components without having to redesign the entire system. Figure 15.6 shows two types of features that are common in architectures: a “bowtie” pattern and an “hourglass” pattern.

The “bowtie” pattern refers to the use of a common interface within a layer of abstraction that enables many different subsystems to connect together across the interface. As an example, in the context of the sensing system for an autonomous vehicle, a common representation of map data allows many types of sensors to



(a) Bowtie pattern

(b) Hourglass pattern

Figure 15.6: Architectural patterns that support modularity. The bowtie pattern is used to connect subsystems at the same level. Proper interface design supports independence of the subsystems. The hourglass pattern is used to connect subsystem at different levels using protocols. Proper protocols supports independence of the subsystems at different levels.




Figure 15.7: DARPA Grand Challenge. “Alice,” Team Caltech’s entry in the 2005 and 2007 competitions and its networked control architecture [56].

feed information into the map and different layers of controller to extract (fused) information from the map. Through this common interface, new sensors or control functionality can be added on either side of the interface without having to redesign the rest of the system.

The “hourglass” pattern represents a hierarchy of control functions that uses a common interface to enable changes above and below that interface to be changed independently of each other. The Open Systems Interconnection model (OSI model) uses seven standardized layers—applications at the top and the physical layer at the bottom—and has been a key to obtain interoperability in communication systems. Another example of an hourglass pattern is used for the planning system in an autonomous vehicle. Trajectories are used to connect higher and lower levels of the navigation system. Any high-level function that eventually leads to a trajectory is completely compatible with the lower level controllers that will track that trajectory, allowing the higher levels of decision making to be changed without having to modify the trajectory tracking code. Similarly, the lower level controllers can be changed without having to redesign the high level decision-making, as long as they properly perform the function of tracking a given trajectory.

The bowtie and hourglass patterns shown here can appear multiple times in a given architecture, so that we obtain appropriate “stacks” of sensing and control functionality.

Example 15.1 Autonomous driving

As an example of a top-down architecture for control, we consider a control system for an autonomous vehicle, shown in Figure 15.7. This control system is designed for driving in urban environments. The feedback system fuses data from road and traffic sensors (cameras, laser range finders, and radar) to create a multi-layer “map” of the environment around the vehicle. This map is used to make decisions about actions that the vehicle should take (drive, stop, change lanes) and plan a specific path for the vehicle to follow. An optimization-based planner is used to compute the trajectory for the vehicle to follow, which is passed to a trajectory tracking module. A supervisory control module performs higher-level tasks such as mission planning and contingency management (if a sensor or actuator fails).




Figure 15.8: The idea of receding horizon control.

We see that this architecture has the basic features shown in Figure 15.4. The control layers are shown in the navigation block, with the mission planner and traffic planner representing two levels of discrete decision-making logic, the path planner representing a trajectory optimization function and then the lower layers of control. Similarly, there are multiple layers of sensing, with low level information, such as vehicle speed and position in the lane, being sent to the trajectory tracking controller, while higher level information about other vehicles on the road and their predicted motions is sent to the trajectory, traffic, and mission planners. ∇

Online Optimization

The use of real-time trajectory generation techniques enables a much more sophisticated approach to the design of control systems, especially those in which constraints must be taken into account. The fact that such trajectories can be computed quickly enables us to use a *receding horizon control* technique: a (optimal) feasible trajectory is computed from the current state to the desired state over a finite time T horizon, used for a short period of time $\delta < T$, and then recomputed based on the new system state starting at time $t + \delta$ until time $t + T + \delta$, as shown in Figure 15.8. Development and application of receding horizon control (also called model predictive control, or MPC) originated in process control industries where the processes being controlled are often sufficiently slow to permit its implementation. An overview of the evolution of commercially available MPC technology is given in [?] and a survey of the state of stability theory of MPC is given in [?].

Figure 15.9 shows a typical setup for a receding horizon control problem. In this formulation, the trajectory generation block solves the following constrained

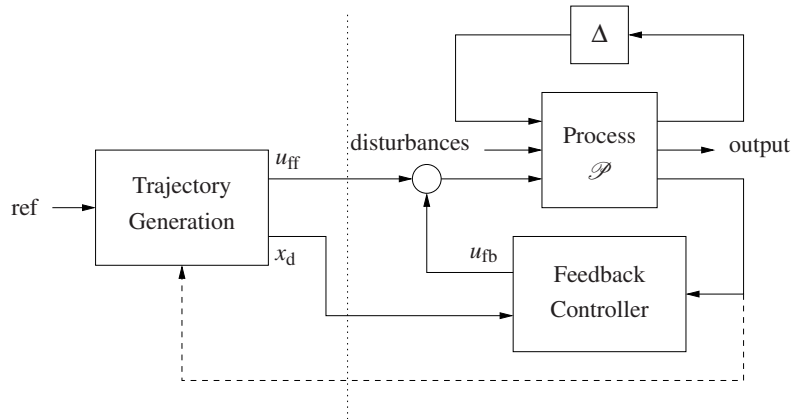


Figure 15.9: Two degree-of-freedom controller design for a process P with uncertainty Δ . The controller consists of a trajectory generator and feedback controller. The trajectory generation subsystem computes a feedforward command u_d along with the desired state x_d . The state feedback controller uses the measured (or estimated) state and desired state to compute a corrective input u_{fb} . Uncertainty is represented by the block Δ , representing unmodeled dynamics, as well as disturbances and noise.

trajectory generation problem at each time step:

$$\min_{(x,u)} = \int_t^{t+T} L(x,u) d\tau + V(x(t+T)) \quad \text{subject to} \quad \begin{cases} x(t) = \text{current state} \\ \dot{x} = f(x) + g(x)u, \\ lb \leq c(x,u) \leq ub. \end{cases} \quad (15.1)$$

One of the challenges of properly implementing receding horizon control is that instabilities can result if the problem is not specified correctly. In particular, because we optimize the system dynamics over a finite horizon T , it can happen that choosing the optimal short term behavior can lead us *away* from the long term solution (see [?] for a more detailed discussion). To address this problem, the terminal cost $V(x(T))$ must have certain properties to ensure stability (see [?] for details).

One of the chief challenges in implementing receding horizon control is the need for fast computation of feasible trajectories. One class of systems for which this is easier are *differentially flat* systems, defined briefly in Section 8.5.

Discrete-decision making and supervisory control

Design of control systems involves the analysis and synthesis of feedback controllers at multiple levels of abstraction, from fast feedback loops around actuators and subsystems, to higher level decision-making logic in supervisory controllers and autonomous systems. One of the major challenges in design of complex networked control systems—such as those arising in aerospace, computing, robotics, and critical infrastructure—is insuring that the combination of dynamical behavior and logical decision-making satisfies safety and performance specifications. In many of these areas, verification and validation are now dominant drivers of sched-

ule and cost, and the tools available for design of such systems are falling behind the needs of systems and control engineers, particularly in the area of systematic design of the mixed continuous and discrete control laws for networked systems.

We consider systems consisting of subsystems/agents whose dynamics are described by ordinary differential equations of the form

$$\dot{x}^i = f^i(x^i, \alpha^i, u^i), \quad y^i = h(x^i, \alpha^i),$$

where $x^i \in \mathbb{R}^{n_i}$ is the continuous state of the i th subsystem, $\alpha \in \mathcal{A}$ is the discrete state, $u^i \in \mathbb{R}^{m_i}$ is a control input and $y^i \in \mathbb{R}^{p_i}$ is the measured output of subsystem. The discrete state evolves according to a set of “guarded commands”, in which the discrete state α is updated to a new value only if a guard $g^p(x, \alpha)$ is true:

$$g_j^p(x, \alpha) \implies \alpha' = r_j^p(x, \alpha).$$

This specification allows the discrete state to evolve in an asynchronous way (e.g. for modeling failures) or to depend on the system state (e.g. to model nonlinearities or changes in connectivity). A model of this type is called a *discrete transition system*. The overall system, consisting of both continuous dynamics and discrete (state) dynamics, is called a *hybrid system*.

A controller for the system is a combination of a continuous control law and a discrete control *protocol*:

$$u = k(x, \alpha), \quad g_j^c(x, \alpha) \implies \alpha' = r_j^c(x, \alpha). \quad (15.2)$$

The control protocol is in the form of discrete transition system, which controls some subset of the discrete states. The discrete state is assumed to be updated by a periodically controlled process that examines the guards and updates appropriate rules. This model for the control allows for the possibility of distributed computation in which different systems (or subsystems) execute on loosely regulated clocks.

The system specification for a hybrid system is often composed of both a continuous performance specification and a discrete performance specification. For the continuous portion of the specification, a typical form is to use a cost function J that is written as a finite horizon cost

$$J = \int_0^T L(x, \alpha, u) dt + V(x(T)). \quad (15.3)$$

This function, a variant of which we have just seen in the context of receding horizon control, uses an integral cost over a fixed horizon T along with a terminal cost $V(x(T))$, where V is an appropriate positive function.

For the discrete performance specification, we make use of *temporal logic* formulas. For a discrete transition system, a temporal logic formula describes conditions on the sequence of events. One mathematical language that is widely used is *linear temporal logic* (LTL), which makes use of two temporal operators: always (\square) and eventually (\diamond). Given a logical formula $\varphi(\alpha)$ that evaluates to true or false for a given (discrete) state α , we can define a temporal logic formula $\square\varphi$, which is interpreted as meaning that $\varphi(\alpha)$ should be true at all times in the future. Similarly,

the formula $\diamond\varphi$ represents the temporal logic statement that the logical formula φ is true at some future state. By combining these temporal operators with standard logical operators, we can obtain more complex formulas. For example, the formula

$$\square(\varphi \implies \diamond\psi)$$

can be interpreted as saying that at all times, if the formula φ is satisfied (evaluates to true), then eventually (at some time in the future) the formula ψ is true. This formula is a typical form for specifying that a system should respond to a certain event (captured by φ becoming true) by taking a certain action (captured by ψ becoming true). The always (\square) operator at the outer level of the formula describes the specification that this condition should be satisfied at all times, which means that the system should respond over and over again if the event condition occurs repeatedly.

A typical LTL formula for a supervisory control system has the form

$$\varphi_{\text{init}} \wedge \square\varphi_{\text{env}} \implies \square\varphi_{\text{safe}} \wedge \diamond\varphi_{\text{goal}}, \quad (15.4)$$

where φ_{init} represents a proposition describing the initial state of the system, φ_{env} describes the possible actions of the environment, φ_{safe} is a safety requirement, and φ_{goal} is a progress requirement. The environmental description φ_{env} , safety requirement φ_{safe} , and progress requirement φ_{goal} are typically described using LTL or other temporal logics (including computational tree logic, CTL, or one of its variants, TCTL, pCTL, etc). These languages allow various specifications on sequences of actions, such as requiring that a certain condition hold until another condition is satisfied, or requiring that a certain condition occurs on a regular basis.

The control design problem for a supervisory system consists of finding a control law of the form in equation (15.2) that satisfies the system specification (15.4) while minimizing the cost function (15.3). Traditional approaches to this problem involve synthesis of the continuous control law using optimization-based approaches, as described earlier in this section, and manual design of the discrete control protocols. The system is then checked against the specification by running (many) repeated simulations and checking that in each case the system specification is satisfied.

Linking Continuous and Discrete Controllers

In Section 15.3 we saw techniques for finding (optimal) controllers for continuous systems and in the previous section we have talked about techniques for design of discrete controllers. Going back to our initial top-down architecture in Figure 15.4, we now consider the problem of how to link these two different layers of the control system.

One approach to linking the two layers is to develop control techniques that can handle both the continuous and discrete dynamics in a common framework. One such approach is to make use of the notion of *mixed logical dynamical* (MLD) systems, introduced by Bemporad and Morari [?]. In this formulation, we extend

the trajectory generation and receding horizon frameworks to handle discrete variables in the underlying constrained optimization problem. This requires the use of so-called *mixed integer solvers*, which allow optimization for problems with both continuous and discrete (integer valued) variables. As computers become increasingly powerful, the size and complexity of problems that we can handle with these tools have increased and these techniques are more and more commonly used.

An alternative is to solve the supervisory control problem and the trajectory generation problem separately, with an appropriate simplified representation of the dynamics of the layers above and below the one that for which we are doing the design.

For example, a common approach in doing trajectory generation is to assume that the discrete state (which might represent an operational mode or an environmental context) transitions from one state to another and then remains fixed for a sufficiently long period of time. Under this assumption, we can focus our attention on the problem of trajectory generation with an initial condition that may represent the state of the system just prior to transition to the new mode and the assume that the mode stays constant for the duration of our planning horizon. We are then required to make sure that our supervisory controller imposes this restriction on the time between mode switches as part of its specification.

This type of linkage between two layers in our abstraction is called a *vertical contract* and can often be written in *assume-guarantee* form. The supervisory layer will assume that the trajectory generation layer maintains the system specification in a given mode given enough “settling” time, and then guarantees that it does not switch the system mode more quickly than the settling time. Similarly, the trajectory generation layer will assume that the switching is sufficiently long, but must then guarantee that it satisfies the system specification within the prescribed settling time.

Similar to the simplified model of the supervisory controller used at the trajectory generation layer (the mode is essentially constant), the supervisory controller must have an appropriate representation of the trajectory generation dynamics. Since the supervisory controller design is done using discrete transition system models, our representation of the dynamics of the lower layers of abstraction must be in terms of a discrete transition system model. One simple approach to such a representation is to break the continuous state space into a collection of discrete regions, such that each continuous state is contained in one discrete region. The dynamics of the trajectory generation layer can then be represented as a set of transitions between adjacent regions in the discretized state space, as illustrated in Figure 15.10a. In setting up the discrete representation, it is important that any trajectory in the discretized state space correspond to a feasible trajectory in the original continuous state space, and *vice versa*. This is required so that when the supervisory controller layer commands the system to move from one region to another then the trajectory generation layer is able to do so, and conversely when the continuous system executes a command and moves from one (continuous) state to another, it corresponds to a valid transition between the regions.




Figure 15.10: Representation of continuous dynamics as a discrete transition system

Model Checking and Program Synthesis

Given a system model, a controller design and performance specifications, we must verify that the controller satisfies the specifications. This is typically done by running many simulations. A problem in using simulation to try to check whether the design satisfies the specification is that it can be prohibitively time-consuming to simulate every possible sequence of events. For purely discrete state systems, in which the dynamics are completely specified by a set of discrete variables α , it turns out that there are more efficient techniques for verifying that a system is specified by making use of the structure of the logical formula and the discrete dynamics of the system. These verification techniques are referred to as *model checking*.

A block diagram that describes how model checking is performed is shown in Figure 15.11. The main idea of model checking is to make use of a model of the (discrete) system dynamics and the temporal logic specification to create a discrete transition system, known as the *product transition system*, where finding a path through this transition system represents a possible system execution (set of allowable state transitions) that *violates* the system specification. If no such path can be found, then the system specification is satisfied. But if a path is found, it represents a counter-example that can be used to update the controller design.

Model checking tools, such as nuSMV [?], PRISM [?], SPIN [?], and TLC [?] are capable of handling relatively large discrete state systems with quite general classes of specifications (beyond simple LTL formulas). They are now widely used in industry and are increasingly being applied in mission and safety critical applications, such as planetary exploration and aviation.

Despite their increasing power and applicability, a limitation in the use of model checking is that it does not necessarily provide any insight into how to redesign the system if the specification is not met. Rather, it provides a counterex-




Figure 15.11: Model checking.

ample indicating what can go wrong, and leaves it to the designer to understand why the controller is not correct and then redesign the system. The updated design can be re-verified using the model checker, and this process is iterated until a correct design is obtained.

An alternative to iterative manual design is to make use of results in *program synthesis* or *correct-by-construction* synthesis [?]. The basic idea in synthesis is to create an algorithm that takes a model of the discrete system dynamics along with a temporal logic specification and then *synthesize* a control protocol such that the resulting closed loop system satisfies the specification. This approach is conceptually similar to the LQR design technique outlined in Section 7.3: given a system specification (cost function to minimize) and system model (A, B, C matrices), create a controller ($u = -Kx$) that by construction stabilizes the system and optimizes the performance specification.

Model checking and program synthesis make use of a common set of information (system model and system specification), but solve different problems, as illustrated in Figure 15.12. The advantage of using synthesis is that it provides a control protocol that is guaranteed to satisfy the specification. That is, no matter what sequence of environmental events occur, the system will always respond in a manner that satisfies the specification.

Of course, correct-by-construction synthesis is not a panacea. It can only be used for certain types of specifications and it too can be overwhelmed by complexity. But it is an increasingly important tool in the systematic synthesis of control protocols, especially safety-critical systems. A practical way to obtain a safe system is to design the system in layers with algorithms and guards. The algorithms perform their ordinary tasks and the guards determine if the system is functioning properly. The innermost layer, which for flight control is called *flying home mode*, consists of a simple robust controller that provides the basic properties. If the loop




Figure 15.12: Verification versus synthesis.

is simple enough it can be designed to be correct-by-construction. The algorithms at the higher layers deliver better performance, and they have guards that move the system to a lower, safer layer if the system does not perform properly. The safety of the guards must of course also be guaranteed. Systems of this type have been designed by Sha [?], and examples of guards for adaptive control are found in [?].

15.4 Bottom-Up Architectures

An alternative to the approaches described in the previous section is to design controllers by interconnecting low-level control systems to create more sophisticated capabilities. This approach is referred to as “bottom up” design. The idea of building complex systems from standard parts has emerged in many branches of engineering. In design of mechanical devices it was very efficient to standardize nuts and bolts. In electronics, standards emerged for components, circuits, boards, and patterns for VLSI design.

To use bottom-up design of control systems, we must find the appropriate components or building blocks and the rules for interconnecting them. The building blocks are controllers (often PID), nonlinear functions, filters, logic, and finite state machines. They can either be separate pieces of hardware or function blocks implemented in software. The systems are built loop by loop by using structures or control principles such as feedback, feedforward, and gain scheduling, which have been discussed extensively in previous chapters. In this section we will introduce other architectural structures: cascade control, mid-range control, selector control, internal model control (IMC), Smith predictors, extremum seeking, and complementary filtering.




Figure 15.13: Block diagram of a system with cascade control. The dashed controller block has two controllers C_p and C_s in series. Both controllers have two degrees of freedom with two inputs: the reference (top) and the measured variable (bottom). The process has two blocks P_1 and P_2 in cascade. There are two loops: an *inner or secondary loop* and an *outer or primary loop*.

Bottom-up architectures can deal with systems having many inputs, many outputs, and constraints. An advantage is that the system can be designed, commissioned, and tuned loop by loop. The disadvantages are that it is not easy to judge the consequences of adding loops and that there may be difficulties when loops are interacting. Bottom-up architectures are easy to use for simple systems, but for complicated systems it may be better to use a top-down approach.

Cascade Control – Several Sensors

Cascade control is a scheme for using one actuator and two or more sensors. Versions of it were previously encountered in Figure 1.13 and in Example 12.9, where it was called inner-outer loop design. A block diagram of a closed loop system with cascade control is shown in Figure 15.13. The dashed process block has one control variable u and two measured signals: the primary output y and the secondary output y_s . The process is modeled by the blocks with transfer functions P_1 and P_2 , which capture how the measured signals are related to the control variable. The dashed controller block in the figure has two controllers C_p and C_s , which are connected in cascade. It has three inputs: the reference r and the measured signals y and y_s . The *primary controller* C_p attempts to keep the output y close to the reference r by manipulating the reference input r_s of the *secondary controller* C_s . The secondary controller C_s attempts to keep the secondary output y_s close to its reference r_s manipulating the control variable u .

The controllers C_p and C_s can be of any type, but PID controllers are most common. Design is done loop by loop starting with the inner loop. If integral action is used, it is necessary to have a scheme to avoid integral windup. Anti-windup for the secondary controller can be done in the conventional way, since the controller drives the actuator directly. To provide anti-windup for the primary controller it must be told when the secondary controller goes into anti-windup mode.

Cascade control is a convenient way to use extra sensors to improve control performance. It can be used with more than two sensors: the ultimate case is state




Figure 15.14: Block diagram of a closed loop system with mid-range control. The process has one output y and two inputs u_f and u_c . The input u_f provides fine control with limited range, the input u_c provides coarse control with wide range. The controller C_f attempts to keep the output y close to its reference r . The controller C_c controls u_c so that the variable u_f is in the middle of this range.

feedback when all states are measured. Cascade control is particularly useful when there is significant time delay or dynamics between the input u and the primary output y , but significantly less dynamics between u and the secondary output y_s . A tight feedback in the inner loop reduces effects of disturbances and uncertainties in the block P_1 , and simplifies the design of the outer loop.

Mid-Range Control – Many Actuators

Midranging is a control scheme that can be used when several control signals influence the same measured output. A typical example is a CD player with a fast actuator having a small actuation range and a slower actuator with a wide actuation range. The block diagram in Figure 15.14 shows an example. The process has two control signals u_f for fine control and u_c for coarse control. They influence the output y through dynamics described by the transfer functions P_f and P_c .

The controller C_f drives u_f and is the primary controller that attempts to keep the output y close to its reference r . The second controller C_c drives the subsystem P_c , which has wide actuation range. The measured signal for C_c is the output u_f of the controller C_f , and the reference is the middle range of u_f . The controller C_c manipulates the control variable u_c so that u_f is in the middle of its operating range and can handle moderately large disturbances. Both controllers will act for large disturbances.

The controllers are tuned loop by loop, starting with the C_f . Anti-windup is handled in the standard manner if controllers have integral action.

Selector Control – Equipment Protection

Selector control is used to control a primary variable while keeping auxiliary variables within given constraints for safety or for equipment protection. A selector is a function with many inputs and one output. The output of a *maximum* selector is the largest of the inputs and the output of a *minimum* selector is the smallest of the inputs.

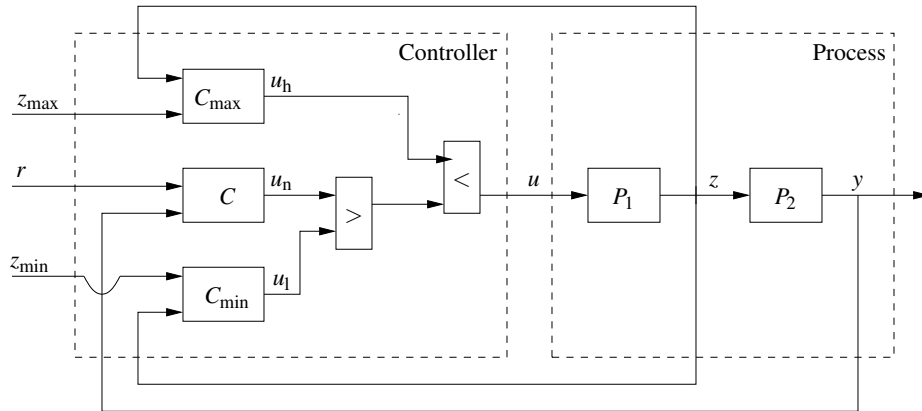


Figure 15.15: Block diagram of a system with selector control. The primary controller C is designed to keep y close to its reference r . The controllers C_{\max} and C_{\min} are controllers that ensures that the intermediary variable z is in its permissible range $z_{\min} < z < z_{\max}$. The block marked $<$ is a minimum selector, whose output equals the smallest input. The block marked $>$ is a maximum selector, whose output equals the largest input.

Selector control is illustrated in Figure 15.15. The primary controlled variable is the process output y . The primary controller C , with output u_n , attempts to keep y close to its reference r . To guarantee safe operation the auxiliary variable z should be kept in the range $z_{\min} < z < z_{\max}$. This is accomplished by the secondary controllers C_{\max} and C_{\min} , which have reference signals z_{\max} and z_{\min} .

The control signal u is generated by sending u_n , u_h and u_l through maximum and minimum selectors as shown in the figure. Under normal conditions the auxiliary variable z is larger than z_{\min} and smaller than z_{\max} . The system then acts as if the maximum and minimum controllers were not present and the input to the control system is $u = u_n$. If the variable z goes above its upper limit z_{\max} the error of the controller C_{\max} becomes negative and the minimum selector chooses u_h instead of u_n . Control is then executed by the controller C_{\max} , which attempts to drive z towards z_{\max} . A similar situation occurs if the variable z becomes smaller than z_{\min} . The switches between the controllers are determined by the limits z_{\min} and z_{\max} and the gains of the secondary controllers, which often are proportional controllers.

Selector control is commonly used to provide safety, for example to maintain temperature while ensuring that a pressure does not exceed certain limits or to avoid stall in compressor control. We have only discussed maximum and minimum selectors, but there are also *median selectors* and *two-out-of-three selectors* that are used for high integrity systems with multiple sensors.

Design of the controllers can be made loop by loop since only one of the controllers is in operation at each time. There may, however, be complications when switching between the controllers. With controllers having integral action, it is necessary to track the integral states of those controllers that are not in operation. Windup protections therefore requires care. Selector control will be complicated




Figure 15.16: Block diagram of a closed loop system with a controller based on the internal model control structure.

when there are many constraints. It is then safer to use multivariable architectures and design methods such as model predictive control.

Internal Model Control – Disturbance Observer

Figure 15.16 shows a controller architecture that is called *internal model control* (IMC) or inferential control. The basic idea is to create an estimate of the effect of the disturbances v and w on the output of the system.

If $\hat{P} = P$ it follows from the block diagram that the control signal u has no effect on the signal ε , hence $\varepsilon = w + Pv$. The signal ε is thus the net effect of the disturbances reduced to the process output, which explains the name disturbance observer. If $G_f = 1$ the block \hat{P}^{-1} generates a control signal that eliminates the disturbance. In reality, ideal disturbance rejection cannot be accomplished because the inverse P^{-1} is normally not realizable. Good disturbance attenuation can, however, be achieved by using an approximate inverse and a proper design of G_f .

To investigate the response to reference signals, we neglect the disturbances v and w . If $\hat{P} = P$ we have $\varepsilon = 0$ and the transfer function from reference to output becomes $G_{yr} = P\hat{P}^{-1}G_f = G_f$. In reality the inverse has to be substituted with an approximate inverse P^\dagger because P normally cannot be inverted. The response to command signals can be shaped by the transfer function $P^\dagger G_f$.

The block diagram in Figure 15.16 can be also be represented as a standard feedback loop with a process P and a controller C where

$$C = \frac{P^{-1}G_f}{1 - G_f}. \quad (15.5)$$

If $\hat{P} = P$, then the controller C cancels all process poles and zeros, which implies that it cannot be used for processes with unstable poles or zeros. The same observation can be made from the fact that the system has a parallel paths with two identical systems, which is a prototype for a system lacking observability and reachability.

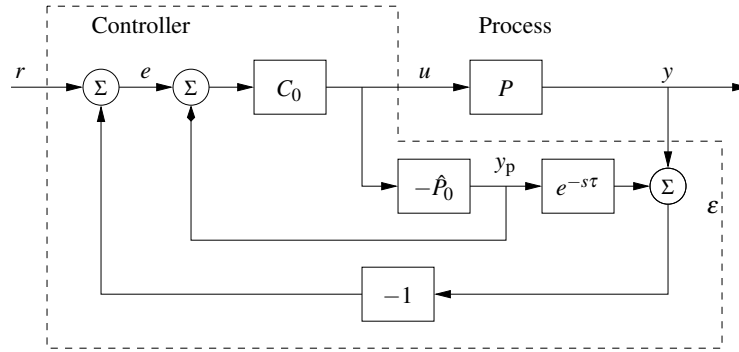


Figure 15.17: Block diagram of a closed loop system with a Smith predictor.

The Smith Predictor – Phase Advance

The Smith predictor is a special controller architecture for systems with time delays. A block diagram of the controller is shown in Figure 15.17. The controller is provided with a model $\hat{P} = \hat{P}_0 e^{-s\tau}$, in parallel with the process $P = P_0 e^{-s\tau}$. The parallel model provides the signal y_p , which is a proxy for the undelayed process output. Notice the similarity with the internal model controller architecture in Figure 15.16. Assuming that $\hat{P}_0 = P_0$ then the signal ε is zero for all u . Applying block diagram algebra then gives the following transfer function for the closed loop system:

$$G_{yr}(s) = \frac{P_0(s)C_0(s)}{1 + P_0(s)C_0(s)} e^{-s\tau}. \quad (15.6)$$

To obtain a desired response to reference signals we can design a controller C_0 for a process with the delay-free dynamics P_0 . Notice that the architecture has two parallel paths with identical dynamics, which is particularly serious if the transfer function P_0 has unstable poles. The Smith predictor cannot be used for processes with integrators without modifications [?].

To get additional insight into the properties of the Smith predictor, we observe that if $\hat{P} = P = P_0 e^{-s\tau}$ the block diagram Figure 15.17 can be redrawn as a conventional feedback loop with the controller

$$C = \frac{C_0}{1 + C_0 P_0 (1 - e^{-s\tau})} = C_0 C_{\text{pred}}, \quad C_{\text{pred}} = \frac{1}{1 + C_0 P_0 (1 - e^{-s\tau})}. \quad (15.7)$$

The controller can thus be viewed as a cascade connection of the conventional controller C_0 with the predictor C_{pred} . Notice that near the gain crossover frequency for $P_0 C_0$ we have $P_0 C_0 \approx -1$ and $C_{\text{pred}} \approx e^{sL}$, indicating that the transfer function C_{pred} has a significant phase advance.

An example is shown in Figure 15.18 for the case where

$$P_0 = \frac{1}{s+1}, \quad C_0 = \left(1 + \frac{1}{0.45s}\right), \quad L = 8. \quad (15.8)$$

The phase curve of C_{pred} shows that the phase lead is very close to that of an ideal




Figure 15.18: Bode plots of the predictor transfer function C_{pred} (solid curve) given by equation (15.7) and the ideal predictor $e^{s\tau}$ (dashed curves) for the transfer functions given in equation (15.8).

predictor $e^{s\tau}$ for the frequencies $\omega = 0.76$ and $\omega = 1.3$, where the phase lead is 360° and 720° . Also notice that the gain curve of the Bode plot has resonances at $\omega = 0.76$ and $\omega = 1.3$ and that the phase increases approximately 180° at those frequencies. This implies that the transfer function C_{pred} has two complex pole pairs in the right half plane.

The Smith predictor gives closed loop systems with good responses to command signals but the response to load disturbance responses are not much better than with PI control because the gain crossover frequency is limited by the time delay (around 0.1 for the system in Figure 15.18). This gain crossover frequency can also be obtained using a PI controller. The predictor C_{pred} is, however, a useful transfer function to provide large phase advance.

Complementary Filtering – Sensor Fusion

Complementary filtering is a technique that can be used to combine the information from different sensors, typically one sensor that is slow but accurate and another that is fast but drifting. One example is to fuse signals from a GPS with signals from gyroscopes and accelerometers to provide good estimates of position, velocity, and acceleration.

Consider the situation when we want to give a good estimate of the variable x and we have two sensors available that give the signals y_1 and y_2 where

$$y_1 = x + w_1, \quad y_2 = x + w_2.$$

The disturbance w_1 has zero mean but the disturbance w_2 may drift. Using exponential signals the complementary filter for recovering the signal x is given by

$$y_f = \frac{1}{s+1} y_1 + \frac{s}{s+1} y_2 = G_1(s) y_1 + G_2(s) y_2. \quad (15.9)$$

Notice that $G_1(s) + G_2(s) = 1$, which explains the name complementary filtering.




Figure 15.19: Self-optimizing control or extremum seeking. The steady state response of the the performance J is shown as a function of the reference r in (a), together with the effects of sinusoidal variations in r . A block diagram of the system is shown in (b).

Both complementary filtering and Kalman filtering can provide improved estimates by fusing information from several sensors, and they can also be optimized if information about the noise characteristics are available. Complementary filtering requires only a model of the sensor system but the Kalman filter requires a model of the complete system dynamics. The Kalman filter can, however, also exploit the control actions. Both methods are widely used both in simple and advanced systems.

Extremum Seeking or Self-Optimization

Another set of useful control structures are *extremum seeking* or *self-optimizing controllers*. Instead of keeping the process output close to a constant reference value, self-optimizing controllers attempt to change the reference so that an objective function is minimized. The idea of extremum seeking is to adjust the reference of a closed loop system so that a performance criterion is optimized. This concept is illustrated in Figure 15.19a. The reference value r of the controller is changed sinusoidally and the performance j is observed. The performance changes very little close to the optimum, see B in the figure. They are in phase with the changes of the reference if the reference is to the left of the maximum (A in the figure), and they are out of phase if the reference is to the right of the maximum (C in the figure). The phase difference can be used to find the how the reference should be changed to optimize the performance. A block diagram of an extremal seeker is shown in Figure 15.19b. The self-optimizer has a block PC that calculates the performance J from the process input u and output y . The signal generator SG generates a sinusoidal signal which is sent to the reference value of the closed loop system. The signal is correlated with the output J from the performance calculator and a low pass filtered version is sent to the reference value. The perturbation signal should be chosen sufficiently slow so that process dynamics can be neglected. There are many other more sophisticated schemes based on optimization and estimation [?]. They differ in how probing, analysis, and action are performed.

15.5 Interaction

A drawback with building a system loop by loop is that there may be unintended interactions. It is therefore important to investigate when interactions arise and what can be done to reduce potential drawbacks. We will start by introducing *the relative gain array (RGA)*, which is a simple measure of interaction. We will then discuss parallel systems, which is a special form of interaction that occurs when several subsystems produce the same effect. A typical example is an electric car with motors on each wheel or a power system with many generators that are synchronized for frequency control.

The Relative Gain Array

To explore the effects of interactions we will investigate control of a system loop by loop. The first problem, which is a prototype for a system that lacks reachability and observability, is to decide if y_i should be controlled by u_j or by some other control signal. This is called the *pairing problem*. The second problem is to investigate if there will be interactions between the loops. It turns out that an understanding of the second problem also solves the first problem.

Consider a system with m inputs and m outputs. Let the transfer function matrix of the system be P . The open loop transfer function from input j to output i is ij th element of the transfer function matrix, which we denote P_{ij} . This transfer function will change when the other loops are controlled. The change depends on the controllers in the other loops in a complicated way. A simple situation is when all other loops are *perfectly controlled* in the sense that all outputs y_k are zero for $k \neq i$. To find the transfer function from u_j to u_i in this case we use exponential signals, which gives

$$u_j = (P^{-1})_{ji} y_i.$$

The closed loop transfer function from u_j to y_j is thus $1/(P^{-1})_{ji}$. The relative gain for the loop ij is defined as the ratio of the transfer functions from u_j to y_i under open loop and ideal closed loop control, hence

$$\lambda_{ij} = P_{ij}(P^{-1})_{ji} = P_{ij}(P^{-T})_{ij},$$

where P^{-T} denotes the transpose of P^{-1} . The transfer functions λ_{ij} can be combined into the matrix

$$\Lambda = P \circ P^{-T}. \quad (15.10)$$

where \circ denotes element by element multiplication of the matrices (the Hadamard product denoted $.*$ in MATLAB). The matrix Λ is called *the relative gain array (RGA)* or *Bristol's RGA* after its inventor [?]. It was originally derived for the steady state case ($s = 0$), which explains the name relative gain and it was later extended to dynamics. We illustrate with an example.

Example 15.2 Relative gain array for 2×2 systems

Consider a static system with two inputs and two outputs. The transfer function

and its inverse are

$$P = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix}, \quad P^{-1} = \frac{1}{p_{11}p_{22} - p_{12}p_{21}} \begin{pmatrix} p_{22} & -p_{12} \\ -p_{21} & p_{11} \end{pmatrix},$$

and the relative gain array (15.10) then becomes

$$\Lambda = \begin{pmatrix} 1 - \lambda & \lambda \\ \lambda & 1 - \lambda \end{pmatrix}, \quad \lambda = \frac{p_{11}p_{22}}{p_{11}p_{22} - p_{12}p_{21}}.$$

In this case the interaction can thus be characterized by a single number λ . Notice that the relative gain array Λ has a special structure: the diagonal elements are equal to $1 - \lambda$ and all row and column sums are one. ∇

The matrix Λ given by equation (15.10) has special properties. Since P^{-1} is the inverse of P we have

$$\sum_{k=1}^n P_{ik}(P^{-1})_{kj} = \delta_{ij},$$

and hence

$$\sum_{i=1}^n P_{ij}(P^{-1})_{ji} = \delta_{ii} = 1.$$

The row and column sums of Λ are thus all equal to one, which implies that the interactions can be characterized by $(n - 1)^2$ elements: one element for $m = 2$ as in Example 15.2 and four elements for $m = 3$.

The relative gain has a good physical interpretation in the static case. If $\lambda_{ij} = 1$ there is no interaction because the open and closed loop gains are the same. The interaction increases the gain if $\lambda_{ij} < 1$ and decreases the gain if $\lambda_{ij} > 0$. The interaction changes the sign of the gain if λ_{ij} is negative. The relative gain can also be used for pairing inputs and outputs as illustrated by the example.

Example 15.3 RGA and pairing

Consider a system where the static gain matrix and the RGA are

$$P = \begin{pmatrix} 2.662 & 8.351 & 8.351 \\ 0.382 & -0.559 & -0.559 \\ 0 & 11.896 & -0.351 \end{pmatrix}, \quad \Lambda = \begin{pmatrix} 0.32 & 0.02 & \mathbf{0.66} \\ \mathbf{0.68} & 0.01 & 0.31 \\ 0 & \mathbf{0.97} & 0.03 \end{pmatrix},$$

where the bold entries are the maximum entries in each row. In this case the pairing is straightforward because there is one largest relative gain in each row and each column, which gives the pairing $y_1:u_3$, $y_2:u_1$, and $y_3:u_2$. The relative gains are also reasonably large. ∇

The relative gain array is a simple measure of interaction, it is dimension free, easy to compute and it gives insight into interactions and pairing of variables. It was originally derived for static systems but analysis of the frequency response $\Lambda(i\omega)$ gives insight into the frequency dependence of the interactions. The RGA also gives information about the variables that should be grouped for multivariable control. Since it was derived under the special assumption of perfect control, it does not capture all aspects of interaction.

Parallel Systems

There are situations when several subsystems are used to control the same variable. Typical examples include temperature control using several cooling or heating devices and control of an electric car with one motor on each wheel. An extreme example is control of a power grid, which may have hundreds of energy sources that all contribute to maintain frequency and voltage of the net. Designing a system loop by loop requires care as is illustrated by the following example.

Example 15.4 Cruise control for electric car

Consider speed control of an electrical car with motors on each wheel. For simplicity we will consider linear motion with only two motors, and we will use the simple model (4.1) in Section 4.1. Neglecting all disturbance forces except the force due to gravity, $F_d = mg\theta$, the model (4.3) becomes

$$m \frac{dv}{dt} = F_1 + F_2 - mg\theta, \quad (15.11)$$

where v is the speed of the car, θ is the slope of the road, and F_1 and F_2 are the forces generated by each tire.

For simplicity we will neglect the dynamics of motors and wheels, so that the forces are simply the output of the wheel controllers. When both wheels have proportional controllers we have

$$F_1 = k_{p1}(r - v), \quad F_2 = k_{p2}(r - v), \quad (15.12)$$

where r is the speed reference. Combining equations (15.11) and (15.12) gives the following equation for the closed loop system:

$$m \frac{dv}{dt} = (k_{p1} + k_{p2})(r - v) - mg\theta = k_p(r - v) - mg\theta,$$

where $k_p = k_{p1} + k_{p2}$. If the slope θ is constant there will be a steady state error $e_{ss} = r - v_{ss} = mg\theta/k_p$ and the steady-state forces are then

$$F_{1,ss} = \frac{k_{p1}}{k_{p1} + k_{p2}} mg\theta, \quad F_{2,ss} = \frac{k_{p2}}{k_{p1} + k_{p2}} mg\theta.$$

The proportional gains k_{p1} and k_{p2} of the controllers thus determine how the *compensation for the disturbance is distributed among the motors*.

Next we will consider the case when each motor drive is provided with a PI controller. The closed loop system is then described by the equations

$$\begin{aligned} m \frac{dv}{dt} &= (k_{p1} + k_{p2})(r - v) + k_{i1}I_1 + k_{i2}I_2 - mg\theta, \\ \frac{dI_1}{dt} &= r - v, \quad \frac{dI_2}{dt} = r - v. \end{aligned} \quad (15.13)$$

Assuming that r and θ are constant, the equilibrium is given by $v = r$ and constant I_1 and I_2 . The variables I_1 and I_2 are not unique, however, and they can have any values as long as

$$k_{p1}I_1 + k_{i2}I_2 = mg\theta.$$




Figure 15.20: Transient behavior of car with two controlled wheels. The system in (a) has a PI controller for each wheel drive. The system in (b) has one PI controller whose output is distributed to the wheel drives. The top plots show the velocity error $e = r - v$, the middle plots show the forces F_1 and F_2 generated by the wheels and the bottom plots show the total force $F_1 + F_2$.

The fact that the closed loop system has infinitely many equilibria is an indication that there may be difficulties. To explore this we will analyze the closed loop system. The transfer functions from r and θ to v are

$$G_{vr} = \frac{k_p s + k_i}{m s^2 + k_p s + k_i}, \quad G_{v\theta} = \frac{m g s}{m s^2 + k_p s + k_i}, \quad (15.14)$$

where $k_i = k_{i1} + k_{i2}$. The velocity v follows reference signals without steady state error since $G_{vr}(0) = 1$ and there will be no steady state error when encountering a slope since $G_{v\theta}(0) = 0$.

The system (15.13) is of third order but the transfer functions (15.14) are of second order, which means that there is a pole/zero cancellation. The canceled mode is governed by the equation $d(I_1 - I_2)/dt = 0$, which has an eigenvalue at the origin. The system has a Kalman decomposition (Figure 8.9a) with a first order subsystem $\Sigma_{\bar{r}\bar{o}}$ with integrator dynamics that is neither reachable from the forces F_1 and F_2 nor observable from v . Recall that a system with two parallel systems having the same mode is a prototype for a system that is neither reachable nor observable.

The fact that the closed loop system has an unstable mode that is neither reachable nor observable is seen in the simulation in Figure 15.20a. The simulation shows what happens when a pulse like perturbation is introduced at the input of the PI controller for wheel 1. The velocity error $e = r - v$ is quite well behaved and the error settles in about 10 time units. The force F_1 (solid curves) reacts quickly to reduce the effect of the disturbance but the force F_2 (dashed curves) goes in the

opposite direction. In steady state the force F_1 settles to a constant value as does F_2 . The steady state forces have different signs, however, which means that one wheel is driving and the other is braking: clearly not a satisfactory behavior.

The irregular behavior is caused by an unreachable and uncontrollable mode in the closed loop system, which caused by two PI controllers in parallel. Having understood what happens, it is straightforward to find a remedy: use one PI controller and distribute its output to the two wheels. The controller is then

$$F_1 = \alpha(k_p(r - v) + k_i I), \quad F_2 = (1 - \alpha)(k_p(r - v) + k_i I), \quad \frac{dI}{dt} = r - v, \quad (15.15)$$

where the parameter $0 < \alpha < 1$ tells how the forces are distributed between the wheels. The closed loop system is described by

$$\frac{dv}{dt} = k_p(r - v) + k_i I - mg\theta, \quad \frac{dI}{dt} = r - v, \quad (15.16)$$

and the forces are given by equation (15.15). This system is of second order and the transfer functions G_{vr} and $G_{v\theta}$ from reference r and slope θ to velocity v are given by equation (15.14).

A simulation of the system is shown in Figure 15.20b. The behavior of the forces are now much more reasonable because they both collaborate to reduce the disturbance. The behavior of the error e and the total force F are, however, the same as for the system with two PI controllers. ∇

The conclusion from the example can be generalized. If parallel systems are controlled by proportional controllers, then the controller gains determine how disturbance attenuation is divided among the subsystems. However, integral control cannot be used in the individual subsystems. Instead we can use a central integrator and distribute its output to the separate controllers for the subsystems.

15.6 Adaptation and Learning

In this section we will briefly discuss control systems with abilities to adapt, learn, and reason. Adaptation is used to adjust to a specified use or situation, learning is used to acquire knowledge or skill by study, instruction, or experience, and reasoning is the intellectual process of seeking truth or knowledge by inferring from either fact or logic. Cars and air vehicles with autonomy are areas where adaptation, learning, and reasoning are essential, but it should be recognized that the abilities of humans are still far ahead of what is achievable using human-engineered systems.

Adaptive Control

Adaptive control is a technique that can be used when there are significant variations in the process and its environment and where neither robust control nor gain scheduling is applicable. Model reference control and the self-tuning regulator are




Figure 15.21: Block diagrams of systems with (a) a model reference adaptive controller (MRAC) and (b) a self-tuning regulator (STR). The block P is the process, C is a controller with adjustable parameters. For the model reference system the requirements R are given in terms of the model F_m , which gives the ideal response to reference signals. The controller parameters θ are adjusted by the parameter adjustment mechanism PA . In the self-tuning regulator the controller parameters θ are adjusted indirectly based on a control design calculation CDC , where the process model is obtained by a recursive parameter estimator RPE .

two common adaptive systems: their block diagrams are shown in Figure 15.21. Notice that in both cases there are two feedback loops: one conventional feedback loop involving the process P and the controller C and a slower loop that adjusts the parameters θ of the controller.

Model reference adaptive control (MRAC) is primarily used for command signal following. A block diagram of the controller is shown in Figure 15.21a. The controller consists of three blocks F_m , C , and PA . The desired response to command signals is given by the transfer function F_m , which is shaped to satisfy the requirements R . The controller C has adjustable parameters θ . The parameter adjustment mechanism PA , receives the process input u , the process output y , and the desired response y_m , and it generates the the process parameters θ . A simple parameter adjustment mechanism is given by the “MIT rule” [?, Section 5.2]:

$$\frac{d\theta}{dt} = -\gamma e \frac{\partial e}{\partial \theta}, \quad (15.17)$$

where γ is a parameter, $e = y_m - y$, and $\partial e / \partial \theta$ is a sensitivity derivative. The MIT rule is a very simple way to adjust the parameters. There are many other rules, some of them are derived from Lyapunov theory [?].

The self-tuning regulator (STR) is used both for command signal following and for regulation. The controller is based on the idea of developing a process model automatically and applying some design method to find a suitable controller. A block diagram of a system is shown in Figure 15.21b. The controller has three blocks: a controller C with adjustable parameters θ , a recursive parameter estimator RPE and a controller design calculation CDC . The parameter estimator RPE estimates the process parameters θ recursively from the process input u and output y . The controller design block CDC determines the controller parameters from the

process parameters using some design method. In this calculation it is common to treat the estimates as the true parameters, a principle from decision making under uncertainty called the *certainty equivalence principle* [?]. Uncertainties in the estimates can be taken into account because many estimation schemes provide estimates of parameter uncertainty. The self-tuning regulator is very flexible because many different methods can be used for parameter estimation and control design.

Recursive least squares is a common method for estimating parameters in the model

$$\begin{aligned} y_{t+1} &= -a_1 y_t - a_2 y_{t-1} + \cdots + b_1 u_t + \cdots + e_{t+1} = \boldsymbol{\varphi}_t \boldsymbol{\theta} + e_{t+1}, \\ \boldsymbol{\varphi}_t &= [-y_t \ -y_{t-1} \ \cdots \ u_t \ u_{t-1} \ \cdots], \quad \boldsymbol{\theta} = [a_1 \ a_2 \ \cdots \ b_1 \ b_2 \ \cdots]^T \end{aligned} \quad (15.18)$$

by minimizing the mean square error $\sum e_k^2$. The estimates are given by

$$\begin{aligned} \hat{\boldsymbol{\theta}}_t &= \hat{\boldsymbol{\theta}}_{t-1} + K_t (y_t - \boldsymbol{\varphi}_t^T \hat{\boldsymbol{\theta}}_{t-1}), \quad K_t = P_t \boldsymbol{\varphi}_t^T, \\ P_t &= P_{t-1} \boldsymbol{\varphi}_t^T (\lambda + \boldsymbol{\varphi}_t P_{t-1} \boldsymbol{\varphi}_t^T)^{-1}, \end{aligned} \quad (15.19)$$

which is a special case of the Kalman filter: see equation (8.8) in Section 8.2. The parameter λ controls how quickly old data is forgotten. There are many variations of the parameter estimator: directional forgetting and square root algorithms, where the square root of P is updated instead of P itself [?], are of particular interest.

Applications of adaptive control are found in flight control, process control, and wheel-slip control and other automotive applications. We illustrate by a ship steering application.

Example 15.5 Adaptive ship steering

A conventional autopilot for ship steering is typically based on PID control. The major disturbances are due to wind and waves, which can change significantly. An adaptive controller can model wave generated forces and counteract them efficiently. Ship dynamics and wave forces can be captured by a model of the form (15.18) with 4 a -parameters, constrained to contain an integrator, and 2 b -parameters. Extensive sea trials have shown that the adaptive autopilot has better performance than the conventional autopilot in normal weather conditions and substantially better performance in bad weather conditions. Figure 15.22 shows results from evaluation of the Steermaster autopilot developed by Kockums and now marketed by Northrop Grumman. In the experiments, the conventional and adaptive autopilots were run repeatedly for about an hour each during normal operation. The figure shows that the adaptive autopilot has significantly smaller variations in heading than the conventional autopilot. The difference corresponds to about 3% less fuel consumption. ∇

A difficulty with adaptive control is that parameter estimation is performed when the system is in closed loop. It is then important where the excitation signals occur. Consider for example the standard feedback loop in Figure 15.21a. If the




Figure 15.22: Heading deviation (top) and rudder motion (bottom) for ship steering for (a) a conventional autopilot and (b) an adaptive autopilot. The experiments were performed on a 225000 ton tanker Seascope, wind velocity around 10 knots, from [?].

only perturbation on the system is the signal v we have $r = 0$ and $n = 0$ and we get

$$y = \frac{P(s)}{1 + P(s)C(s)}v, \quad u = -\frac{C(s)P(s)}{1 + P(s)C(s)}v.$$

It then follows that $y = -\frac{1}{C(s)}u$ and any attempt to find a model relating u and y will thus result in the negative inverse of the controller transfer function. However, if $v = 0$ and $w = 0$ we have instead

$$y = \frac{P(s)C(s)}{1 + P(s)C(s)}F(s)r, \quad u = \frac{C(s)}{1 + P(s)C(s)}F(s)r.$$

Hence, $y = P(s)u$ and the process model can indeed be estimated.

To obtain reliable estimates of the process parameters, there must be sufficient variations in the control signal. This can be captured by the notion of *persistent excitation*. A signal $u(t)$ is persistently exciting of order n if

$$U = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T (A(p)u(k))^2 > 0$$

for all nonzero polynomials A of the differential operator $p = d/dt$ with degree less than or equal to $n - 1$. Persistency of excitation determines the number of parameters that can be estimated reliably. A constant is persistently exciting of order 1 and permits estimation of one parameter. A sinusoid is persistently exciting of order two.

To have a reliable parameter estimation it is important to be aware of where disturbances enter and to monitor the excitation. Load disturbances of the process are particularly harmful. A scheme for detecting harmful load disturbances is pre-

sented in [?]. To obtain reliable estimates it is necessary to monitor the excitation of the process and only update parameters when there is sufficient excitation of the process.

An interesting approach to control of uncertain systems was proposed by Feldbaum [?], who emphasized that control should be *investigating* as well as *directing*, and he coined the term dual control for this property. Feldbaum used optimal stochastic control to obtain a controller that was actively introducing perturbations in the process when the process was not properly excited by natural disturbances. The *hyper state* of a dual controller is the conditional probability distribution of the regular states of the process and the parameters. The computations of a dual controller can only be performed in simple cases because the state of the system is a conditional probability distribution over states and parameters [?]. Many heuristic schemes to monitor excitation and to introduce perturbations when needed have therefore been developed. In the field of machine learning this approach is called reinforcement learning.

Learning

A nonlinear function with a learning mechanism is a simple example of a learning system. Learning can be done in two different ways. In supervised learning, the function is created automatically by providing it with a large number of arguments and corresponding function values. In reinforcement learning a criterion for good fit is provided and learning is executed by selecting random arguments and changing the function until a good fit is provided. Representation of the learning mechanism is a central issue. A simple way to represent a function of several variables is to quantize the variables, which we illustrate by an example of unsupervised learning.

Example 15.6 Michie's BOXES method

Michie and Chambers developed a simple learning program called BOXES [?] for game playing. An early connection between learning and control was established when the program was applied to the classical control problem of stabilizing an inverted pendulum. Consider a cart-pendulum system such as that in Figure 7.2b on page 219, where cart position and velocity and pendulum angle and angular velocity are measured. The control signal is a function: $f : \mathbb{R}^4 \rightarrow \mathbb{R}$. BOXES was used to learn an approximation of this function. To implement the system the states were quantized crudely: 5 levels for position and angle, 3 levels for velocities, and 2 levels left (L) or right (R) for the control signal. The control law is then represented by a table with $5 \times 5 \times 3 \times 3 = 225$ entries. Each entry in the table has five numbers: LL (left life), RL (right life), LU (left usage), RU (right usage), which capture features of past experiments, and T (target), which is related to the current mean time for stabilization. The life (L) is a weighted average of the number of times the entry was used before failure. The usage (U) is a weighted number of times the entry was used before failure. The control actions taken upon entering a box is a heuristic function of the values of LL, RL, LU, RU, and T. The

system is initialized by introducing random numbers in the table. Experiments are run and the table is updated. In a typical experiment the system was able to stabilize the pendulum in 25 minutes after a 60 hour training period. ∇

Pendulum stabilization is perhaps not the best case to demonstrate learning, since a student could design a stabilizing controller in less than an hour. Control performance will also be better because a conventional design can avoid the crude quantization used in BOXES. There have, however, been significant developments in machine learning since the late 1960s when BOXES was developed.

Neural Networks

A severe drawback of schemes like BOXES is that the nonlinear function is represented by gridding the state variables, which requires very large tables. To have efficient learning schemes it is necessary to find more efficient ways to represent nonlinear functions and to find efficient learning mechanisms. An artificial neural network is such a representation.

Artificial neural networks were inspired by neuroscience although their current implementation is far from their biological origin. A neuron has many synapses which receives inhibitory or excitatory signals from other neurons. The neuron emits a pulse to other neurons if the net excitation over a short time interval is above a certain level. An *artificial neuron* mimics a real neuron but it operates continuously. A very simple model is

$$y = g\left(\sum_{k=1}^n w_k u_k\right), \quad (15.20)$$

where the parameters w_i are weights. The function g was originally a sigmoid shaped function, for example $g(x) = \tanh x$ or $g(x) = (1 + e^{-x})^{-1}$, but many other functions are currently used, such as $\max(0, x)$, $e^{-x^T Q x}$, and $(1 + x^T Q x)^{-1}$.

An *artificial neural network (ANN)* is a combination of neurons in a layered network as shown in Figure 15.23, which represents the function $f: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ as

$$f(u) = g\left(W^{(3)} g\left(W^{(2)} g\left(W^{(1)} u\right)\right)\right), \quad (15.21)$$

where $W^{(k)}$ is a matrix of the weights to the layer k , and g is a monotone function of the type discussed above (when x is a vector, $g(x)$ is a vector obtained when the function is applied to each element of the vector). The neural network may appear very restricted, but Kolmogorov proved that any continuous function on a finite cube can be represented by a neural network with only one hidden layer [?]. The strong advantage of having many layers was clarified by Håstad [?].

Neural networks have weights $w_{ij}^{(k)}$, which are determined experimentally by matching a large number of arguments and corresponding function values. A useful feature is that both the function and its inverse can be generated from data. The weights in a simple neuron can be determined by the gradient algorithm

$$w_i(k+1) = w_i(k) + \gamma u_i^0(k) (y^0(k) - y(k)), \quad (15.22)$$




Figure 15.23: Schematic diagram of a feedforward neural network with an input layer, a hidden layer, and an output layer.

where u_i^0 and y^0 are training data and y is computed from equation (15.20) with $u = u^0$. This rule can be interpreted as an approximation of a gradient scheme for minimizing the mean square error. It is similar to the MIT rule (15.17) for model reference adaptive control. In neurophysiology the algorithm (15.22) is known as Hebb's rule [?]. Parameters of multi-layer neural networks can be updated by gradient descent, where the gradient is computed by back propagation..

Deep Learning

Neural networks with many layers, so-called *deep learning*, have proven very useful in many fields. Preprocessing of data from information-rich sensors, such as spectrometers and cameras, is particularly useful for control. For example, in autonomous driving it is useful to recognize objects such as houses, road markings, traffic lights, cars, bicycles, and pedestrians, and to track these objects in real time.

A typical scene is illustrated in Figure 15.24, where object recognition is made by a *convolutional neural network (CNN)*. An image is represented as an $m \times n$ array and an additional index k is used to represent images of different colors (red, green, and blue) or different features (lines, corners, wheels, cars, bicycles, etc.). Artificial neural networks have only two operations—a weighted summation and a monotone function of a scalar variable—while the convolutional neural networks have more operations.

The *convolution (C)* function acts on a three-dimensional array and generates a three-dimensional array. The function can be written $y = C(x)$, where

$$y(i, j, k) = w_o(k) + \sum_u \sum_v \sum_l x(i-u, j-v, l)w(u, v, l, k).$$

The function depends on the bias term w_o and the weights w . Convolution is used to detect features, such as edges, corners, wheels, cars, bicycles, etc.

The *Rectified Linear Unit ReLU (R)* operates on an array and generates an array of the same size. The function can be written $y = R(x)$, where

$$y(i, j, k) = \max(x(i, j, k), 0).$$



Figure 15.24: Illustration of feature detection based on deep learning. Objects are detected and classified as car, bicycle, or person, their position and spatial extent are also determined. The images typically have low-resolution to prevent identification of license plates and humans. The particular network which generated this figures has 13 layers and more than 25 million parameters [?].

It operates element-wise on the image. The function is analogous to the monotone function used in conventional neural networks.

The *Max-pooling* function (P) operates on an array x of size $m \times n \times k$ and generates an array y of size $m/2 \times n/2 \times k$. The function can be written $y = P(x)$, where

$$y(i, j, k) = \max(x(2i, 2j, k), x(2i - 1, 2j, k), x(2i, 2j - 1, k), x(2i - 1, 2j - 1, k)).$$

Max-pooling is also called down sampling or subsampling, reduces the dimension of an array but retains the most important features of the image. There are also other pooling operators such as minimum and average pooling.

The function *Softmax* (S) operates on an array and generates an array of the same size. The function can be written $y = S(x)$, where

$$y(i, j, k) = \frac{e^{x(i, j, k)}}{\sum_l e^{x(i, j, l)}}, \quad \sum_k y(i, j, k) = 1.$$

The entries of the output are in the range $[0, 1]$, they indicate the subjective probability of finding the object k in position ij . The function is typically used in the final layer to assign a subjective probability to a detected feature.

A convolutional neural network can be large with many parameters, but parameters appear only in the convolution layer. The parameters are determined by optimization based on large training sets.

Determination of the sizes of the arrays, the number of layers and the order of the operators is an art that requires experience. In computer vision it is inspired by




Figure 15.25: The figure is a simplified version of the convolutional neural network used for the object detection in Figure 15.24. The input is a color image. The output is three layers of images representing the subjective probability of finding car, bicycle, or person at a given position in the image.

earlier efforts based on detection of lines, objects, corners, and blobs.

Figure 15.25 shows a simplified representation of the network used to generate Figure 15.24. The input to the network is a box with 3 RGB planes with 300×300 pixels and the output is a box of three images of size 20×20 pixels each representing a recognized object. The network has 25 million parameters which have to be estimated. The network can be described by the function

$$Q = S \circ C \circ R \circ C \circ (P \circ R \circ C)^3,$$

where $f \circ g$ denotes function composition, i.e. $f \circ g(x) = f(g(x))$. The function is a composition of 13 functions. The first operations creates a wide range of patterns using a combination of convolution (C), ReLU (R), and max-pooling (P). Recognition of the final objects car, bicycle, and pedestrian is done by $S \circ C \circ R \circ C$. The final result three images with numbers that indicate the subjective probability of finding an object at a particular location.

The arrays are illustrated by boxes in the figure. The first layer consists of three arrays of red, blue, and green pixels of the object. The convolution operates sequentially on part of the picture, the kernel is typically a fraction of the image size. Different convolution kernels are applied to detection edges, corners, and other features. The result of for each kernel is stored in a separate layer. Convolution is followed by a rectified linear operator. The size of the image is then reduced using max-pooling (P). The operations C , R , and P are repeated several times to create more features. The final part consists of three layers generated by C , R , and C . The final result, a classification as car, bicycle, pedestrian, or no object, is obtained by applying the softmax operator. The final result can be mapped on the original image as shown in Figure 15.25. There are several ways to improve the position of the objects and their sizes.

A camera with a convolutional neural network can be regarded as a trainable sensor that will detect, classify, and position objects such as cars, bicycles, and people in real time. It is clearly a useful component for autonomous driving and for other applications in which vision-based sensors are used. From the data it is easy to generate warnings, zones where a vehicle can safely enter, and other useful information. The computations are fast, since they only require evaluation

of simple functions; computations can also be parallelized. The network used to generate Figure 15.24 has more than 12 million parameters and it estimates 11 objects. The final trained network can be executed in 60 Hz on a standard PC with an NVIDIA Titan X graphics card.

15.7 Control Design in Common Application Fields

Control is sometimes called the *hidden technology* because it is successfully used practically everywhere without being noticed [?]. In this section we will present the role of control in five different applications fields, which provides a flavor of the commercial landscape of controls, the systems, the controllers, and the users.

Aerospace – High Performance Systems and Highly Skilled Users

The aerospace industry was an early user of control. The Wright brothers flew in 1905 because they had a good insight into dynamics and control. The first autopilot was designed by Sperry in 1914 and autonomous flight was demonstrated in 1947. Today aerospace is a flourishing application area for control, dominated by large companies for civil and military markets. The industry produces airplanes, helicopters, drones, satellites, rockets, missiles, and quadrocopters as well as infrastructure for flight control, which includes air traffic control and automatic landing systems. An indication of the size of the business is that about 2500 aircraft were produced in 2015, generating revenues over 20 billion dollars, and that is only one part of the industry. Aerospace companies typically have large central groups for systems and control.

Typical aerospace systems are operated by highly selected, skilled pilots and astronauts, who are well-trained in simulators and interact with the system by direct manipulation of actuators and reference signals or by changing operating modes. A consequence is that the systems are designed so that the user can directly influence the system in many different ways. Sometimes users have taken over operation of the system and saved the mission, as was done in Apollo 13 [168].

The aerospace industry has been a technology driver with new hardware and control techniques emerging from the industry. There are extreme requirements on safety, which has led to the practice of redundant systems and components, since adopted by the automotive industry. The industry pioneered the use of simulation and model-based systems engineering, development of high precision accelerometers and gyroscopes, and anti-lock braking (used on aircraft as early as 1929). Extremal control was first applied to control of aircraft engines already in 1951. Wide variations of operating conditions stimulated the development of gain scheduling and adaptive control. Optimal control and Kalman filtering were used in the early space efforts. Nonlinear control was used extensively in control of satellites. Unmanned air vehicles were used operationally already in 1970.

Automotive – Complex Systems Used by Ordinary People

The automotive industry is a multi-trillion dollar business. It is dominated by six large companies and many subcontractors. Control is used extensively both in the cars themselves and in the manufacturing of cars. Automobiles are used by ordinary people who interact with the system by changing modes and set points and by direct actuation. Control is executed using electronic control units (ECUs), microprocessors with input/output interfaces. Modern cars have more than 100 electronic control units.

Servo-assisted power braking was used in racing cars in 1914 and became commonly used in the 1920s. Computer control was introduced in the late 1970s to cope with the stringent emission requirements. Once computers were introduced they were applied to more functions: suspension control, anti-lock braking systems (ABS), electronic braking systems (EBS) and electronic stability control (ESC). These systems used accelerometers and gyroscopes to control the brakes individually to improve stability and steering. Adaptive cruise control, based on radar sensors, maintains a constant distance to the car in front. The excellent experience with these systems inspired car manufacturers to introduce more sophisticated systems such as collision avoidance and parking assist. Autonomous driving is well on its way. Control is a key element both on its own but also in combination with computer vision.

Model-based systems engineering is used extensively to improve the efficiency of engineering. The design of the Toyota Prius is an example where modeling and simulation replaced much of the traditional testing using hardware prototypes. Another example is design of climate control systems. The major European car manufacturers and their component suppliers have created an infrastructure for model-based design where the suppliers deliver components with validated dynamical models enabling the car manufacturers to simulate complete systems [?].

A worldwide development partnership of vehicle manufacturers, suppliers, and software companies called AUTOSAR (AUTomotive Open System Architecture) was formed in 2003 (see <http://www.autosar.org>). Standards that enable modularity, scalability, transferability, and reusability of functions have been created, providing a standardized platform for automotive software systems. The standard enables system-wide configuration and optimization to meet run-time requirements of automotive devices.

The large size of the automotive industry provides a mass market for a wide range of industries to develop components and subsystems. The industry stimulated the development of inexpensive emission sensors, accelerometers, and gyroscopes, and even more importantly the microcontroller and the programmable logic controller (PLC).

Early manufacturing systems were automation systems controlled by relays for logic and sequencing. General Motors challenged the electronics industry with requirements for a standard machine controller that could replace the relays resulting in the PLC. The system architecture is based on round robin schedulers with different cycle rates. PLCs were originally programmed in a graphical language called

ladder diagrams (LD), which emulated the ladder logic used to describe relay circuits. Several different programming styles were later standardized: function block diagrams (FBD), sequential function charts (SFC), and structured text (ST). PLCs developed rapidly and became a standard tool for automation in many industries.

Process Industry – Complex Systems with Many Different Users

Process control provides automation for a variety of industries such as chemicals, oil refining, pulp and paper, pharmaceuticals power plants and many others. A characteristic feature of the industry is that control and automation is typically delivered by special companies. This began with instrument companies that developed sensors, recorders, and controllers, including Taylor Instruments, founded in 1851, and Foxboro, founded in 1908. By the mid 1930s there were many companies who supplied sensors, actuators, and controllers to the process industry.

It is normally difficult to find out how much of the turnover of a business is related to control. In process control this data is available, since automation is done by special companies. Control and automation is a 100 billion dollar industry. Distributed control systems account for about 20% of the market, the rest is sensors, actuators, software, and other components. Five dominating suppliers have more than 50% of the market.

Functions of control and logic and sequencing are essential for process operation. Early process control systems had cabinets with analog controllers for regulation and cabinets with relays for logic and sequencing governing startup, shut down, and equipment protection. As technology developed the relays were replaced by programmable logic controllers (PLCs), originating in the automotive industry, and the analog controllers were replaced by distributed control systems (DCS).

The DCS is now the standard tool to provide control in the process industry, as illustrated in Figure 15.5. It has facilities for connecting sensors, actuators, and algorithms and can be viewed as a toolbox for implementing control systems. It is interesting to note that ExxonMobil has recently contracted Lockheed Martin to specify the next generation of distributed control system for process control. The system will be open, secure, and based on standards, leveraging experiences from the Future Airborne Capability Environment (FACE) consortium in the aerospace industry [?].

Process control systems typically have thousands of sensors and actuators, and the systems are also widely distributed geographically. Sensors and actuators are connected to the DCS system by standardized networks (IEC 61784).

Valves are commonly used for actuation in process control. It is customary to have cascade control with inner analog loops with valve positioners to reduce effects of friction and nonlinearities at the lowest level of the hierarchy, and feedback loops for control of pressure and temperature and quality variables at the higher levels.

There are several different types of users of the DCS: the plant managers who set production schedules and directs equipment maintenance, the process engi-

neers who select, configure, and modify the system, the instrument engineers who tune controllers and maintain sensors and actuators and the operators who supervise the operation of the system (see Figure 15.5).

Distributed control systems have many control algorithms that easily can be configured using graphical interfaces. The control algorithms are implemented by process and instrument engineers both by company personnel and by consultants. Controllers are tuned during operation and the system is occasionally reconfigured. Algorithms and languages are standardized by international committees. A wide range of standards for control and automation are set by the International Society of Automation (ISA) and the International Electrical Commission (IEC). There are also some standards for communication organized by special groups.

Although PID control was used in many fields, the major development of the controller and its tuning procedure occurred in the process industries. Most of the controllers (typically 97% [?]) are PID controllers, only a small fraction of them using derivative action. A recent investigation of 100 boiler-turbine units in the Guangdong Province in China showed 94.4% PI, 3.7% PID, and 1.9% advanced controllers [?] .

Control paradigms such as cascade, selector, and midrange control are common, as are gain scheduling, automatic tuning, and model predictive control. Model predictive control emerged from efforts at the Shell oil company to develop effective techniques for control of multi-variable processes. It was originally called dynamic matrix control [?, ?].

Telecommunication – Billions of Systems

Black's invention of the negative feedback amplifier was inspired by the needs to make phone calls over long distances. Intellectual giants like Bode, Nyquist, and Shannon developed theoretical foundations of control and communications.

Today the global telecommunication system is said to be the world's largest man-made artifact, with the total number of mobile subscriptions in the beginning of 2016 at around 7.4 billion. In many countries the number of mobile subscriptions exceeds the population. The *Internet of Things (IoT)* is of particular interest because it enables simple ways of using feedback, and combined with the cloud it offers many interesting opportunities for novel control applications. It is expected that the number of IoT devices will surpass mobile subscriptions by 2018. Telecommunications is a high pace industry where the consumer preferences change quickly, making it hard to predict what products will be like 2–3 years from now and the rules of the game change continuously.

The development of camera modules, GPS modules, accelerometers, and gyros for the mobile phone industry has decreased the cost for such sensors by several orders of magnitude because of the large volumes involved. Reduction in size and improvements in power efficiency have been required to fit sensors into hand held devices with reasonable battery life times. The inexpensive components have then found several uses in other fields, for example virtual reality.

Cost efficiency is vital for production in large volumes, making it economical

to put large engineering efforts into cost optimization even for minute details. The requirements from the communication field lead the development of smaller and more energy efficient solutions.

There is extensive standardization in the telecommunication markets, forced by the fact that the technologies share the common radio frequency spectrum and a carefully controlled use of this limited resource has been needed. It is also highly beneficial for the consumer if devices from different manufacturers can function together. The development of technology is coordinated in different groups. The 3GPP consortium develops the standards for mobile communication using GSM (2G), WCDMA (3G), LTE (4G) as well as future communication standards such as 5G and beyond. Similarly, IEEE working groups develop standards, for example 802.11 for wireless routers. Competing operators, vendors, and mobile equipment manufacturers meet in standardization meetings to discuss and decide on new functionality and performance requirements for future devices. New inventions and intellectual property rights are very important for competitive reasons which leads to both portfolio agreements and patent battles.

Control enters on many levels, from analog electronics where it is used to improve performance such as linearity and power efficiency of power amplifiers, to higher functional levels where control is used to continuously choose suitable system parameters such as transmission power and coding schemes depending on existing communication conditions. PID control is often used together with and gain scheduling and simple adaptive schemes.

15.8 Further Reading

A comprehensive treatment of architecture and design of complex systems is given in [?]. Much of the development of PID controllers occurred in the power and process industries, where also many of the associated bottom-up control architectures like cascade, selector, and midranging control appeared. A detailed treatment is given in [18]. Complementary filtering is well described in []. The internal model control architecture is described in [?]. The internal model principle which says that a good controller should contain a model of the process is formulated in [?]. The Smith predictor [?, ?] is closely related to internal model control because it also uses a parallel model. It was invented by Otto J. M. Smith, a legendary professor from University of California, Berkeley who also invented posicast control [?], a scheme for controlling highly oscillatory systems.

Decoupling has long been used in many areas of control, the architecture with inverted decoupling is introduced in [?], a state space version is given in [?]. The relative gain array, a simple way to quickly estimate the interaction, was introduced in [?]. A detailed treatment is found in [?].

Gain scheduling is widely used in practice, overviews are presented in [?, ?, ?]. Adaptive control emerged from the desire to avoid gain scheduling in flight control and the dream of having a controller that can automatically adjust itself to good performance in the process industry. Early pioneering work is found in [?]. Current

knowledge is found in [21, ?, ?]. System identification is a key element of adaptive control; the book [?] is an excellent reference. In spite of advances in adaptive control, gain scheduling is still the dominant control scheme for flight control [?], with one reason being the significant developments in air-data sensors. The process engineers dream of having a simple universal controller that self-adjusts to provide good robust performance; the relay auto-tuner for PID control is a partial answer.

Rosenblatt's perceptron [?] was the first neural network, which was used to separate hyperplanes in pictures. An analog version, the Addaline [?], invented by Widrow, was used in simple adaptive systems and for noise cancellation. A severe and partially unfair criticism was given by Minsky and Papert [?] who did not realize the advantage with many layers. There was a revival when backpropagation was introduced to find parameters in networks with many layers [?]. Backpropagation is closely related to dynamic programming [?]. The disadvantage of networks with few layers was clarified in [?]. An early application to picture classification is given in [?], using deep structures and convolutional networks. There have lately been significant advances in object recognition, driven by improved algorithms and good test bases. The ImageNet Large Scale Visual Recognition Challenge is a benchmark with millions of images and hundreds of objects. Competitions have been run since 2010. The classification error was brought down from around 30% in 2010 to 3% in 2016 [?, ?, ?]. The usefulness of deep convolutional networks has been proven in many domains: character recognition, computer vision, and game-playing. AlphaGo's algorithm uses a Monte Carlo tree search to find its moves based on knowledge previously obtained by machine learning based on extensive training, both from playing with humans and computers [?].

Bibliography

- [Abk69] M. A. Abkowitz. *Stability and Motion Control of Ocean Vehicles*. MIT Press, Cambridge, MA, 1969.
- [Ack72] J. Ackermann. Der Entwurf linearer Regelungssysteme im Zustandsraum. *Regelungstechnik und Prozessdatenverarbeitung*, 7:297–300, 1972.
- [Ack85] J. Ackermann. *Sampled-Data Control Systems*. Springer, Berlin, 1985.
- [AF66] Michael Athans and Peter Falb. *Optimal Control*. McGraw-Hill, New York, NY, 1966. Dover Reprint 2007.
- [Agn76] C. E. Agnew. Dynamic modeling and control of congestion-prone systems. *Operations Research*, 24(3):400–419, 1976.
- [ÅH86] Karl Johan Åström and A. Helmersson. Dual control of an integrator with unknown gain. *Comp. & Maths. with Appls.*, 12A(6):653–662, 1986.
- [ÅH05] K. J. Åström and T. Hägglund. *Advanced PID Control*. ISA—The Instrumentation, Systems, and Automation Society, Research Triangle Park, NC, 2005.
- [Ahl66] L. V. Ahlfors. *Complex Analysis*. McGraw-Hill, New York, 1966.
- [ÅHL94] K. J. Åström, C. C. Hang, and B. C. Lim. A new smith predictor for controlling a process with an integrator and long dead-time. *IEEE Transactions on Automatic Control*, 39:343–345, 1994.
- [ÅKL05] K. J. Åström, R. E. Klein, and A. Lennartsson. Bicycle dynamics and control. *IEEE Control Systems Magazine*, 25(4):26–47, 2005.
- [Alu15] R. Alur. *Principles of Cyber-Physical Systems*. MIT Press, 2015.
- [AM90] B. D. O. Anderson and J. B. Moore. *Optimal Control Linear Quadratic Methods*. Prentice Hall, Englewood Cliffs, NJ, 1990. Republished by Dover Publications, 2007.
- [AM10] P. Albertos and I. Mareels. *Feedback and Control for Everyone*. Springer, 2010.
- [Ano92] Anon. V-model development standard for it-systems of the federal republic of germany. Technical report, Federal German Armed Forces, 1992.
- [Apo67] T. M. Apostol. *Calculus*, Vol. II: *Multi-Variable Calculus and Linear Algebra with Applications*. Wiley, New York, 1967.
- [Apo69] T. M. Apostol. *Calculus*, Vol. I: *One-Variable Calculus with an Introduction to Linear Algebra*. Wiley, New York, 1969.
- [Ari94] R. Aris. *Mathematical Modeling Techniques*. Dover, New York, 1994. Originally published by Pitman, 1978.
- [Arn78] V. I. Arnold. *Mathematical Methods in Classical Mechanics*. Springer, New York, 1978.
- [Arn87] V. I. Arnold. *Ordinary Differential Equations*. MIT Press, Cambridge, MA, 1987. 10th printing 1998.
- [AS82] R. H. Abraham and C. D. Shaw. *Dynamics—The Geometry of Behavior*, Part 1: *Periodic Behavior*. Aerial Press, Santa Cruz, CA, 1982.

- [ASMN03] M. Atkinson, M. Savageau, J. Myers, and A. Ninfa. Development of genetic circuitry exhibiting toggle switch or oscillatory behavior in *Escherichia coli*. *Cell*, 113(5):597–607, 2003.
- [Åst99] Karl Johan Åström. Automatic control—The hidden technology. In P. M. Frank, editor, *Advances in Control—Highlights of the ECC '99*, pages 1–29, London, UK, 1999. Springer.
- [Åst00] K. J. Åström. Limitations on control system performance. *European Journal on Control*, 6(1):2–20, 2000.
- [Åst06] K. J. Åström. *Introduction to Stochastic Control Theory*. Dover, New York, 2006. Originally published by Academic Press, New York, 1970.
- [Ath75] D. P. Atherton. *Nonlinear Control Engineering*. Van Nostrand, New York, 1975.
- [AVK87] A. A. Andronov, A. A. Vitt, and S. E. Khaikin. *Theory of Oscillators*. Dover, New York, 1987.
- [ÅW97] K. J. Åström and B. Wittenmark. *Computer-Control Systems: Theory and Design*. Prentice Hall, Englewood Cliffs, NJ, 3rd edition, 1997.
- [ÅW08a] K. J. Åström and B. Wittenmark. *Adaptive Control*. Dover, New York, 2nd edition, 2008. Originally published by Addison Wesley, 1995.
- [ÅW08b] K. J. Åström and B. Wittenmark. *Adaptive Control*. Dover, New York, 2nd edition, 2008. Originally published by Addison Wesley, 1995.
- [BÅ70] R. Bellman and K. J. Åström. On structural identifiability. *Mathematical Biosciences*, 7:329–339, 1970.
- [Bas01] T. Basar, editor. *Control Theory: Twenty-five Seminal Papers (1932–1981)*. IEEE Press, New York, 2001.
- [BB91] T. Basar and P. Bernhard. *H^∞ -Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Birkhauser, Boston, 1991.
- [BBvB⁺01] K. Beck, M. Beedle, A. van Bennekum, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, R. Jeffries, Jon Kern, Brian Marick, Robert C. Martin, Steve Mallor, Ken Shwaber, and Jeff Sutherland. The Agile Manifesto. Technical report, The Agile Alliance, 2001. Available at <http://agilemanifesto.org>.
- [BC48] G. S. Brown and D. P. Campbell. *Principles of Servomechanisms*. Wiley, New York, 1948.
- [BD04] W. E. Boyce and R. C. DiPrima. *Elementary Differential Equations*. Wiley, New York, 2004.
- [Bec05] J. Bechhoefer. Feedback for physicists: A tutorial essay on control. *Reviews of Modern Physics*, 77:783–836, 2005.
- [Bel57] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton NJ, 1957.
- [Ben79] S. Bennett. *A History of Control Engineering: 1800–1930*. Peter Peregrinus, Stevenage, 1979.
- [Ben93] S. Bennett. *A History of Control Engineering: 1930–1955*. Peter Peregrinus, Stevenage, 1993.
- [Ber54] L. L. Beranek. *Acoustics*. McGraw-Hill, New York, 1954.
- [Ber89] R. N. Bergman. Toward physiological understanding of glucose tolerance: Minimal model approach. *Diabetes*, 38:1512–1527, 1989.
- [BG68] B. Brawn and F. Gustavson. Program behavior in a paging environment. *Proceedings of the AFIPS Fall Joint Computer Conference*, pages 1019–1032, 1968.

- [BG87] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Englewood Cliffs, 1987.
- [BH75] A. E. Bryson, Jr. and Y.-C. Ho. *Applied Optimal Control: Optimization, Estimation, and Control*. Wiley, New York, 1975.
- [Bia95] B. Bialkowski. Process control sample problems. In N. J. Sell, editor, *Process Control Fundamentals for the Pulp & Paper Industry*. Tappi Press, Norcross, GA, 1995.
- [BK64] R. E. Bellman and R. Kalaba. *Selected Papers on Mathematical Trends in Control Theory*. Dover, New York, 1964.
- [BKJS16] M. Blanke, M. Kinnaert, J.Lunze, and M. Staroswiecki. *Diagnosis and Fault-Tolerant Control*. Springer, 2016. Third Edition.
- [Bla34] H. S. Black. Stabilized feedback amplifiers. *Bell System Technical Journal*, 13:1–2, 1934.
- [Bla77] H. S. Black. Inventing the negative feedback amplifier. *IEEE Spectrum*, pages 55–60, 1977.
- [Bla91] J. H. Blakelock. *Automatic Control of Aircraft and Missiles*. Addison-Wesley, Cambridge, MA, 2nd edition, 1991.
- [Bli90] G. Blickley. Modern control started with Ziegler-Nichols tuning. *Control Engineering*, 37:72–75, 1990.
- [BM99] Alberto Bemporad and Manfred Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407 – 427, 1999.
- [Boa16] BKCASE Editorial Board, editor. *The Guide to the Systems Engineering Body of Knowledge (SEBoK)*. Stevens Institute of Technology, v. 1.6 edition, 2016. Accessed 27 August 2016, <http://sebokwiki.org>.
- [Bod45] H. W. Bode. *Network Analysis and Feedback Amplifier Design*. Van Nostrand, New York, 1945.
- [Bod60] H. W. Bode. Feedback—The history of an idea. In *Symposium on Active Networks and Feedback Systems*. Polytechnic Institute of Brooklyn, New York, 1960. Reprinted in [30].
- [BP96] 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.
- [Bri66] E. H. Bristol. On a new measure of interactions for multivariable process control. *IEEE Transactions on Automatic Control*, 11(1):133–134, 1966.
- [Bro70] R. W. Brockett. *Finite Dimensional Linear Systems*. Wiley, New York, 1970.
- [Bro00] 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, Berlin, 2000.
- [BRS60] J. F. Blackburn, G. Reethof, and J. L. Shearer. *Fluid Power Control*. MIT Press, Cambridge, MA, 1960.
- [BT78] C. Brosilov and M. Tong. Inferential control systems. *AIChE Journal*, 24(3):457–465, 1978.
- [Can03] R. H. Cannon. *Dynamics of Physical Systems*. Dover, New York, 2003. Originally published by McGraw-Hill, 1967.
- [CCG⁺02] A. Cimatti, E. M. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella. Nusmv 2: An opensource tool for symbolic model checking. In *Proceedings of the 14th International Conference on Computer Aided Verification*, pages 359–364. Springer-Verlag, 2002.

- [CCS15] E. Crawley, B. Cameron, and D. Selva. *Architecture: Strategy and Product Development of Complex Systems*. Pearson, New York, NJ, 2015.
- [CD75] R. F. Coughlin and F. F. Driscoll. *Operational Amplifiers and Linear Integrated Circuits*. Prentice Hall, Englewood Cliffs, NJ, 6th edition, 1975.
- [CD91] F. M. Callier and C. A. Desoer. *Linear System Theory*. Springer-Verlag, London, 1991.
- [CFG⁺06] L. B. Cremean, T. B. Foote, J. H. Gillula, G. H. Hines, D. Kogan, K. L. Kriechbaum, J. C. Lamb, J. Leibs, L. Lindzey, C. E. Rasmussen, A. D. Stewart, J. W. Burdick, and R. M. Murray. Alice: An information-rich autonomous vehicle for high-speed desert navigation. *Journal of Field Robotics*, 23(9):777–810, 2006.
- [CJ59] H. S. Carslaw and J. C. Jaeger. *Conduction of Heat in Solids*. Clarendon Press, Oxford, UK, 2nd edition, 1959.
- [CM51] H. Chestnut and R. W. Mayer. *Servomechanisms and Regulating System Design*, Vol. 1. Wiley, New York, 1951.
- [CR80] C. R. Cutler and B. C. Ramaker. Dynamic matrix control—A computer control algorithm. In *Proceedings Joint Automatic Control Conference*, San Francisco, CA, 1980.
- [Cro75] Crocus. *Systemes d'Exploitation des Ordinateurs*. Dunod, Paris, 1975.
- [CS08] G. C. and K. S. Sin. *Adaptive Filtering Prediction and Control*. Dover, New York, 2008. Originally published by Prentice Hall, 1984.
- [CT84] C. Cobelli and G. Toffolo. Model of glucose kinetics and their control by insulin, compartmental and non-compartmental approaches. *Mathematical Biosciences*, 72(2):291–316, 1984.
- [Cul66] W. J. Culver. On the existence and uniqueness of the real logarithm of a matrix. *Proc. American Mathematical Society*, 17(5):1146–1151, 1966.
- [DB04] R. C. Dorf and R. H. Bishop. *Modern Control Systems*. Prentice Hall, Upper Saddle River, NJ, 10th edition, 2004.
- [DFT92] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum. *Feedback Control Theory*. Macmillan, New York, 1992.
- [DGH⁺02] Y. Diao, N. Gandhi, J. L. Hellerstein, S. Parekh, and D. M. Tilbury. Using MIMO feedback control to enforce policies for interrelated metrics with application to the Apache web server. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium*, pages 219–234, 2002.
- [DGKF89] J. C. Doyle, K. Glover, P. P. Khargonekar, and B. A. Francis. State-space solutions to standard H_2 and H_∞ control problems. *IEEE Transactions on Automatic Control*, 34(8):831–847, 1989.
- [DH85] J. P. Den Hartog. *Mechanical Vibrations*. Dover, New York, 1985. Reprint of 4th ed. from 1956; 1st ed. published in 1934.
- [dJ02] H. de Jong. Modeling and simulation of genetic regulatory systems: A literature review. *Journal of Computational Biology*, 9:67–103, 2002.
- [DM02a] L. Desborough and R. Miller. Increasing customer value of industrial control performance monitoring—Honeywell's experience. In *Sixth International Conference on Chemical Process Control*. AIChE Symposium Series Number 326 (Vol. 98), 2002.
- [DM02b] L. Desborough and R. Miller. Increasing customer value of industrial control performance monitoring - Honeywell's experience. In *Sixth International Conference on Chemical Process Control*. AIChE Symposium Series Number 326 (Volume 98), 2002.
- [DM14] D. Del Vecchio and R. M. Murray. *Biomolecular Feedback Systems*. Princeton University Press, 2014.

- [DOD⁺15] G. Drion, T. O’Leary, J. Dethier, A. Franci, and R. Sepulchre. Neuronal behaviors: A control perspective. *Proc. IEEE Control and Decision Conference*, pages 1923–1944, 2015.
- [Dos68] F. H. Dost. *Grundlagen der Pharmakokinetik*. Thieme Verlag, Stuttgart, 1968.
- [Doy78] J. C. Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on Automatic Control*, 23(4):756–757, 1978.
- [Dra55] C. S. Draper. Flight control. *Jornal Royal Aeronautical Society*, 59(July):451–477, 1955. 45th Wilber Wright Memorial Lecture.
- [Dub57] L. E. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79:497–516, 1957.
- [Dys04] F. Dyson. A meeting with Enrico Fermi. *Nature*, 247(6972):297, 2004.
- [EG05] S. P. Ellner and J. Guckenheimer. *Dynamic Models in Biology*. Princeton University Press, Princeton, NJ, 2005.
- [EHBM08] B. Eisenhower, G. Hagen, A. Banaszuk, and I. Mezić. Passive control of limit cycle oscillations in a thermoacoustic system using asymmetry. *Journal of Applied Mechanics*, 75:011021–1–011021–9, 2008.
- [EKR03] E. N. Elnozahy, M. Kistler, and R. Rajamony. Energy-efficient server clusters. In *Power-Aware Computer Systems*, pages 179–197. Springer, 2003.
- [EL00] M. B. Elowitz and S. Leibler. A synthetic oscillatory network of transcriptional regulators. *Nature*, 403(6767):335–338, 2000.
- [Ell94] J. R. Ellis. *Vehicle Handling Dynamics*. Mechanical Engineering Publications, London, 1994.
- [ESGK02] H. El-Samad, J. P. Goff, and M. Khammash. Calcium homeostasis and parturient hypocalcemia: An integral feedback perspective. *Journal of Theoretical Biology*, 214:17–29, 2002.
- [FCF⁺06] P. G. Fabietti, V. Canonico, M. O. Federici, M. Benedetti, and E. Sarti. Control oriented model of insulin and glucose dynamics in type 1 diabetes. *Medical and Biological Engineering and Computing*, 44:66–78, 2006.
- [Fel65] A. A. Feldbaum. *Optimal Control Theory*. Academic Press, New York, 1965.
- [FLMR92] M. Fliess, J. Levine, P. Martin, and P. Rouchon. On differentially flat nonlinear systems. *Comptes Rendus des Séances de l’Académie des Sciences, Serie I*, 315:619–624, 1992.
- [For61] J. W. Forrester. *Industrial Dynamics*. MIT Press, Cambridge, MA, 1961.
- [Fou07] J. B. J. Fourier. On the propagation of heat in solid bodies. Memoir, read before the Class of the Instut de France, 1807.
- [FPEN05] G. F. Franklin, J. D. Powell, and A. Emami-Naeini. *Feedback Control of Dynamic Systems*. Prentice Hall, Upper Saddle River, NJ, 5th edition, 2005.
- [Fra87] B. A. Francis. *A Course in \mathcal{H}_∞ Control*. Springer-Verlag, Berlin, 1987.
- [Fra07] A. Fradkov. *Cybernetical Physics: From Control of Chaos to Quantum Control*. Springer, Berlin, 2007.
- [Fri04] B. Friedland. *Control System Design: An Introduction to State Space Methods*. Dover, New York, 2004.
- [FW76] B. A. Francis and M. W. Wonham. The internal model principle of control theory. *Automatica*, 12(5):492–500, 1976.

- [Gan60] F. R. Gantmacher. *The Theory of Matrices, Vol. 1 and 2*. Chelsea Publishing Company, 1960.
- [GB42] M. A. Gardner and J. L. Barnes. *Transients in Linear Systems*. Wiley, New York, 1942.
- [GF71] L. Gunkel and G. F. Franklin. A general solution for linear sampled data systems. *IEEE Transactions on Automatic Control*, AC-16:767–775, 1971.
- [GGS01] G. C. Goodwin, S. F. Graebe, and M. E. Salgado. *Control System Design*. Prentice Hall, Upper Saddle River, NJ, 2001.
- [GH83] J. Guckenheimer and P. Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, Berlin, 1983.
- [Gil63] E. Gilbert. Controllability and observability in multivariable control systems. *SIAM Journal of Control*, 1(1):128–151, 1963.
- [GL95] M. Green and D. J. N. Limebeer. *Linear Robust Control*. Prentice Hall, Englewood Cliffs, NJ, 1995.
- [GM61] D. Graham and D. McRuer. *Analysis of Nonlinear Control Systems*. Wiley, New York, 1961.
- [God83] K. Godfrey. *Compartment Models and Their Application*. Academic Press, New York, 1983.
- [Gol53] H. Goldstein. *Classical Mechanics*. Addison-Wesley, Cambridge, MA, 1953.
- [Gol70] S. W. Golomb. Mathematical models—Uses and limitations. *Simulation*, 4(14):197–198, 1970.
- [GP82] M. Giobaldi and D. Perrier. *Pharmacokinetics*. Marcel Dekker, New York, 2nd edition, 1982.
- [GPD59] J. C. Gille, M. J. Pelegrin, and P. Decaulne. *Feedback Control Systems; Analysis, Synthesis, and Design*. McGraw-Hill, New York, 1959.
- [GQH16] R. Gonzalez, F. Qi, and B. Huang. *Process Control System Fault Diagnosis: A Bayesian Approach*. Wiley, 2016.
- [Gre59] P. C. Gregory, editor. *Proc. Self Adaptive Flight Control Symposium*. Wright Air Development Center, Wright-Patterson Air Force Base, Ohio, 1959.
- [Gui63] E. A. Guillemin. *Theory of Linear Physical Systems*. MIT Press, Cambridge, MA, 1963.
- [Hah67] W. Hahn. *Stability of Motion*. Springer, Berlin, 1967.
- [Hås87] J. T. Håstad. *Computational Limitations of Small Depth Circuits*. MIT Press, Cambridge, Massachusetts, 1987.
- [HB90] J. K. Hedrick and T. Batsuen. Invariant properties of automobile suspensions. In *Proceedings of the Institution of Mechanical Engineers*, volume 204, pages 21–27, London, 1990.
- [HC03] C. Hajiyev and F. Caliskan. *Fault Diagnosis and Reconfiguration in Flight Control Systems*. Kluwer, 2003.
- [HD95] M. B. Hoagland and B. Dodson. *The Way Life Works*. Times Books, New York, 1995.
- [HDPT04] J. L. Hellerstein, Y. Diao, S. Parekh, and D. M. Tilbury. *Feedback Control of Computing Systems*. Wiley, New York, 2004.
- [Heb49] D. O. Hebb. *The Organization of Behavior*. Wiley & Sons, New York, 1949.
- [Her04] D. V. Herlihy. *Bicycle—The History*. Yale University Press, New Haven, CT, 2004.

- [HH52] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117(500–544), 1952.
- [HK83] J. Hammer and P.P. Khargonekar. Decoupling of linear systems by dynamic output feedback. *Mathematical Systems Theory*, 17:135–157, 1983.
- [HMTG00] C. V. Hollot, V. Misra, D. Towsley, and W-B. Gong. A control theoretic analysis of RED. In *Proceedings of IEEE Infocom*, pages 1510–1519, 2000.
- [Hol03] G. J. Holzmann. *The SPIN Model Checker*. Addison-Wesley, 2003.
- [Hor63] I. M. Horowitz. *Synthesis of Feedback Systems*. Academic Press, New York, 1963.
- [Hor75] I. M. Horowitz. Superiority of transfer function over state-variable methods in linear, time-invariant feedback system design. *IEEE Transactions on Automatic Control*, AC-20(1):84–97, 1975.
- [Hor91] I. M. Horowitz. Survey of quantitative feedback theory. *International Journal of Control*, 53:255291, 1991.
- [HP01] HP. Model 200a audio oscillator, 1939, 2001. <http://www.hp.com/hpinfo/about/hp/histnfacts/museum/earlyinstruments/0002> [Accessed 2017-04-11].
- [HSH15] R. J. Hawkins, J. K. Speakes, and D. E. Hamilton. Monetary policy and PID control. *Journal of Economic Interaction and Coordination*, 10(1):183–197, 2015.
- [Hug93] T. P. Hughes. *Elmer Sperry: Inventor and Engineer*. John Hopkins University Press, Baltimore, MD, 1993.
- [Hun07] J. Y. Hung. Posicast control past and present. *IEEE Multidisciplinary Engineering Education Magazine*, 3:7–11, 2007.
- [HW00] D. Hanahan and R. A. Weinberg. The hallmarks of cancer. *Cell*, 100:57–70, 2000.
- [Isi95] A. Isidori. *Nonlinear Control Systems*. Springer-Verlag, Berlin, 3rd edition, 1995.
- [Ito70] M. Ito. Neurophysiological aspects of the cerebellar motor system. *International Journal of Neurology*, 7:162–178, 1970.
- [Jac65] H. A. Jacobs. *Frank Lloyd Wright: America’s Greatest Architect*. Harcourt, Brace & World., 1965.
- [Jac72] J. A. Jacquez. *Compartment Analysis in Biology and Medicine*. Elsevier, Amsterdam, 1972.
- [Jac95] V. Jacobson. Congestion avoidance and control. *ACM SIGCOMM Computer Communication Review*, 25:157–173, 1995.
- [Jan14] P. K. Janert. *Feedback Control for Computer Systems*. O’Reilly Media, 2014.
- [JBRM99] J. B. J. B Rawlings and D. Q. Mayne. *Model Predictive Control Theory and Design*. Nob Hill Publishing, Madison, WI, 1999.
- [JNP47] H. James, N. Nichols, and R. Phillips. *Theory of Servomechanisms*. McGraw-Hill, New York, 1947.
- [JT61] P. D. Joseph and J. T. Tou. On linear control theory. *Transactions of the AIEE*, 80(18), 1961.
- [Jun02] W. G. Jung, editor. *Op Amp Applications*. Analog Devices, Norwood, MA, 2002.
- [Kal60] R. E. Kalman. Contributions to the theory of optimal control. *Boletín de la Sociedad Matemática Mexicana*, 5:102–119, 1960.

- [Kal61a] R. E. Kalman. New methods and results in linear prediction and filtering theory. Technical Report 61-1, Research Institute for Advanced Studies (RIAS), Baltimore, MD, February 1961.
- [Kal61b] R. E. Kalman. On the general theory of control systems. In *Proceedings of the First IFAC Congress on Automatic Control, Moscow, 1960*, volume 1, pages 481–492. Butterworths, London, 1961.
- [KÅT⁺79] C. G. Källström, Karl Johan Åström, N. E. Thorell, J. Eriksson, and L. Sten. Adaptive autopilots for tankers. *Automatica*, 15:241–254, 1979.
- [KB61] R. E. Kalman and R. S. Bucy. New results in linear filtering and prediction theory. *Transactions of the ASME (Journal of Basic Engineering)*, 83 D:95–108, 1961.
- [Kel85] F. P. Kelly. Stochastic models of computer communication. *Journal of the Royal Statistical Society*, B47(3):379–395, 1985.
- [Kel94] K. Kelly. *Out of Control*. Addison-Wesley, Reading, MA, 1994. Available at <http://www.kk.org/outofcontrol>.
- [KFA69] R. E. Kalman, P. L. Falb, and M. A. Arbib. *Topics in Mathematical System Theory*. McGraw-Hill, New York, 1969.
- [KG02] B. C. Kuo and F. Golnaraghi. *Automatic Control Systems*. Wiley, New York, 8th edition, 2002.
- [Kha01] H. K. Khalil. *Nonlinear Systems*. Macmillan, New York, 3rd edition, 2001.
- [KHN63] R. E. Kalman, Y. Ho, and K. S. Narendra. *Controllability of Linear Dynamical Systems*, volume 1 of *Contributions to Differential Equations*. Wiley, New York, 1963.
- [Kit95] C. Kittel. *Introduction to Solid State Physics*. Wiley, New York, 1995.
- [Kit04] Hiroaki Kitano. Biological robustness. *Nature Reviews Genetics*, 5(11):826–837, 2004.
- [KKK95] M. Krstić, I. Kanellakopoulos, and P. Kokotović. *Nonlinear and Adaptive Control Design*. Wiley, 1995.
- [Kle75] L. Kleinrock. *Queuing Systems*, Vols. I and II. Wiley-Interscience, New York, 2nd edition, 1975.
- [Kle89] R. E. Klein. Using bicycles to teach dynamics. *Control Systems Magazine*, 9(3):4–8, 1989.
- [KN00] U. Kiencke and L. Nielsen. *Automotive Control Systems: For Engine, Driveline, and Vehicle*. Springer, Berlin, 2000.
- [KNP11] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In G. Gopalakrishnan and S. Qadeer, editors, *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*, pages 585–591. Springer, 2011.
- [Kol57] A.N. Kolmogorov. On the representations of continuous functions of many variables by superpositions of continuous functions of one variable and addition. *Doklady Akademii Nauk USSR*, 114(5):953–956, 1957.
- [Kra63] N. N. Krasovski. *Stability of Motion*. Stanford University Press, Stanford, CA, 1963.
- [Krs03] K. B. Ariyur and M. Krstić. *Real-Time Optimization by Extremum-Seeking Control*. Wiley, 2003.
- [Kru09] Paul Krugman. *The Return of Depression Economics and the Crisis of 2008*. W. W. Norton & Company, New York, 2009.
- [KS01] J. Keener and J. Sneyd. *Mathematical Physiology*. Springer, New York, 2001.

- [KSH12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [Kum01] P. R. Kumar. New technological vistas for systems and control: The example of wireless networks. *Control Systems Magazine*, 21(1):24–37, 2001.
- [Kun93] P. Kundur. *Power System Stability and Control*. McGraw-Hill, New York, 1993.
- [KV86] P. R. Kumar and P. Varaiya. *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Prentice Hall, Englewood Cliffs, NJ, 1986.
- [Lam03] L. Lamport. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Pearson Education, 2003.
- [LaS60] J. P. LaSalle. Some extensions of Lyapunov’s second method. *IRE Transactions on Circuit Theory*, CT-7(4):520–527, 1960.
- [LBD⁺89] Y. LeCun, B. Boser, J.S. Denker, D. Henderson, R. E. Howard, W. Hubbrd, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.
- [Lew03] A. D. Lewis. A mathematical approach to classical control. Technical report, Queens University, Kingston, Ontario, 2003.
- [Lju87] L. Ljung. *System Identification—Theory for the User*. Prentice Hall, Upper Saddle River, NJ, 1987. Second Edition 1999.
- [Lju99] L. Ljung. *System Identification – Theory for the User*. Prentice Hall, Upper Saddle River, NJ, 2nd edition, 1999.
- [LLAE⁺17] F. Lamnabhi-Lagarrigue, A. Annaswamy, S. Engell, A. Isaksson, P. Khargonekar, R. M. Murray, H. Nijmeijer, T. Samad, D. Tilbury, and P. Van den Hof. Systems & control for the future of humanity, research agenda: Current and future roles, impact and grand challenges. *Annual Reviews in Control*, 43:1–64, 2017.
- [LPD02] S. H. Low, F. Paganini, and J. C. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, pages 28–43, February 2002.
- [LPW⁺02] S. H. Low, F. Paganini, J. Wang, S. Adlakha, and J. C. Doyle. Dynamics of TCP/RED and a scalable control. In *Proceedings of IEEE Infocom*, pages 239–248, 2002.
- [LS01] D. Levy and J. Sanford. Technology pioneer william r. hewlett dead at 87. *Stanford Report*, January 12 2001. <http://news.stanford.edu/news/2001/january17/hewlett-a.html> [Accessed 2017-04-11].
- [LS06] D. J. N. Limebeer and Sharp R. S. Bicycles, motorcycles and models. *Control Systems Magazine*, 26(5):34–61, 2006.
- [LS15] E. A. Lee and S. A. Seshia. *Introduction to Embedded Systems, A Cyber-Physical Systems Approach*. <http://LeeSeshia.org>, 2015. ISBN 978-1-312-42740-2.
- [Lun05] K. H. Lundberg. History of analog computing. *IEEE Control Systems Magazine*, pages 22–28, March 2005.
- [LV11] E. A. Lee and P. Varaiya. *Structure and Interpretation of Signals and Systems*. Lee-Varaiya.org, 2011. Available online at <http://leevaraiya.org>.
- [LW13] E. Lavretsky and K. A. Wise. *Robust and Adaptive Control with Aerospace Applications*. Springer, London, 2013.
- [LWE00] D. J. Leith and Leithead W. E. Survey of gain-scheduling analysis and design. *International Journal on Control*, 73(11):1001–1025, 2000.
- [MÅB⁺03] R. M. Murray, K. J. Åström, S. P. Boyd, R. W. Brockett, and G. Stein. Future directions in control in an information-rich world. *Control Systems Magazine*, April 2003.

- [Mac37] D. A. MacLulich. *Fluctuations in the Numbers of the Varying Hare (*Lepus americanus*)*. University of Toronto Press, 1937.
- [Mac45] L.A. MacColl. *Fundamental Theory of Servomechanisms*. Van Nostrand, Princeton, NJ, 1945. Dover reprint 1968.
- [Mac89] J. M. Maciejowski. *Multivariable Feedback Design*. Addison Wesley, Reading, MA, 1989.
- [Mal59] J. G. Malkin. *Theorie der Stabilität einer Bewegung*. Oldenbourg, München, 1959.
- [Man02] R. Mancini. *Op Amps for Everyone*. Texas Instruments, Houston, TX, 2002.
- [May70] O. Mayr. *The Origins of Feedback Control*. MIT Press, Cambridge, MA, 1970.
- [MB61] E. Mishkin and L. Braun. *Adaptive Control Systems*. McGraw-Hill, New York, 1961.
- [MC68] D. Michie and R. A. Chambers. Boxes: An experiment in adaptive control. In E. Dale and D. Michie (Eds.), editors, *Machine Intelligence 2*, pages 137–152. Oliver and Boyd, Edinburgh, 1968.
- [McA83] T. McAvoy. *Interaction Analysis*. ISA, Research Triangle Park, NC, 2083.
- [McF53] M. W. McFarland, editor. *The Papers of Wilbur and Orville Wright*. McGraw-Hill, New York, 1953.
- [MG90] D. C. McFarlane and K. Glover. *Robust Controller Design Using Normalized Coprime Factor Plant Descriptions*. Springer, New York, 1990.
- [MH98] J. E. Marsden and M. J. Hoffmann. *Basic Complex Analysis*. W. H. Freeman, New York, 1998.
- [Mil66] H. T. Milhorn. *The Application of Control Theory to Physiological Systems*. Saunders, Philadelphia, 1966.
- [Min02] D. A. Mindel. *Between Human and Machine: Feedback, Control, and Computing Before Cybernetics*. Johns Hopkins University Press, Baltimore, MD, 2002.
- [Min08] D. A. Mindel. *Digital Apollo: Human and Machine in Spaceflight*. The MIT Press, Cambridge, MA, 2008.
- [MLK06] A. Makroglou, J. Li, and Y. Kuang. Mathematical models and software tools for the glucose-insulin regulatory system and diabetes: An overview. *Applied Numerical Mathematics*, 56:559–573, 2006.
- [MLS94] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.
- [MP69] M. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, Massachusetts, 1969.
- [MR94] J. E. Marsden and T. S. Ratiu. *Introduction to Mechanics and Symmetry*. Springer-Verlag, New York, 1994.
- [MRRS00] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- [MS15] J. R. Marden and J. S. Shamma. Game Theory and Distributed Control. In *Handbook of Game Theory with Economic Applications*, chapter 16, pages 861–899. Elsevier, 2015.
- [Mur03] R. M. Murray, editor. *Control in an Information Rich World: Report of the Panel on Future Directions in Control, Dynamics and Systems*. SIAM, Philadelphia, 2003.
- [Mur04] J. D. Murray. *Mathematical Biology*, Vols. I and II. Springer-Verlag, New York, 3rd edition, 2004.

- [Mur07] R. M. Murray. Recent research in cooperative control of multi-vehicle systems. *ASME Journal of Dynamic Systems, Measurement and Control*, 129(5):571–583, 2007.
- [MW80] Zohar Manna and Richard Waldinger. A deductive approach to program synthesis. *ACM Transactions on Programming Languages and Systems*, 2(1):90–121, 1980.
- [Nah88] P. J. Nahin. *Oliver Heaviside: Sage in Solitude: The Life, Work and Times of an Electrical Genius of the Victorian Age*. IEEE Press, New York, 1988.
- [NS99] 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. University of Groningen, 1999.
- [Nyq32] H. Nyquist. Regeneration theory. *Bell System Technical Journal*, 11:126–147, 1932.
- [Nyq56] H. Nyquist. The regeneration theory. In R. Oldenburger, editor, *Frequency Response*, page 3. MacMillan, New York, 1956.
- [Oga01] K. Ogata. *Modern Control Engineering*. Prentice Hall, Upper Saddle River, NJ, 4th edition, 2001.
- [Old56] R. Oldenburger, editor. *Frequency Response*. MacMillan, New York, 1956.
- [OSFM07] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- [OWN96] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab. *Signals and Systems*. Prentice-Hall, second edition edition, 1996.
- [Par93] L. E. Parker. Designing control laws for cooperative agent teams. In *Proc. IEEE International Conference on Robotics and Automation*, pages 582–587, 1993.
- [PB86] G. Pacini and R. N. Bergman. A computer program to calculate insulin sensitivity and pancreatic responsivity from the frequently sampled intravenous glucose tolerance test. *Computer Methods and Programs in Biomedicine*, 23:113–122, 1986.
- [Phi48] G. A. Philbrick. Designing industrial controllers by analog. *Electronics*, 21(6):108–111, 1948.
- [PN00] W. F. Powers and P. R. Nicastrì. Automotive vehicle control challenges in the 21st century. *Control Engineering Practice*, 8:605–618, 2000.
- [PPP02] S. Prajna, A. Papachristodoulou, and P. A. Parrilo. SOSTOOLS: Sum of squares optimization toolbox for MATLAB, 2002. Available from <http://www.cds.caltech.edu/sostools>.
- [Pto14] Claudius Ptolemaeus, editor. *System Design, Modeling, and Simulation using Ptolemy II*. Ptolemy.org, 2014.
- [QB97] S. J. Qin and T. A. Badgwell. An overview of industrial model predictive control technology. In J.C. Kantor, C.E. Garcia, and B. Carnahan, editors, *Fifth International Conference on Chemical Process Control*, pages 232–256, 1997.
- [RÅ15] A Rantzer and K.J. Åström. Control theory. In N. J. Higham, editor, *The Princeton Companion to Applied Mathematics*. Princeton University Press, Princeton and Oxford, 2015.
- [RHW86] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating error. *Nature*, 323(6088):533–536, 1986.
- [Rig63] D. S. Riggs. *The Mathematical Approach to Physiological Problems*. MIT Press, Cambridge, MA, 1963.
- [RM71] H. H. Rosenbrock and P. D. Moran. Good, bad or optimal? *IEEE Transactions on Automatic Control*, AC-16(6):552–554, 1971.

- [Ros62] F. Rosenblatt. *Principles of Neurodynamics*. Spartan Books, New York, 1962.
- [RST12] G. Rafal, R. G. Sanfelice, and A. Teel. *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton University Press, 2012.
- [Rug91] W. J. Rugh. Analytic framework for gain scheduling. *Control Systems Magazine*, 11(1):79–84, 1991.
- [Rug95] W. J. Rugh. *Linear System Theory*. Prentice Hall, Englewood Cliffs, NJ, 2nd edition, 1995.
- [SA92] J. S. Shamma and M. Athans. Gain scheduling: potential hazards and possible remedies. *Control Systems Magazine*, 12(3):101–107, 1992.
- [SA14] T. Samad and A. M. Annaswamy, editors. *The Impact of Control Technology*. IEEE Control Systems Society, 2nd edition, 2014. Available at www.ieeeccs.org.
- [SÅD⁺07] G. Schitter, K. J. Åström, B. DeMartini, P. J. Thurner, K. L. Turner, and P. K. Hansma. Design and modeling of a high-speed AFM-scanner. *IEEE Transactions on Control System Technology*, 15(5):906–915, 2007.
- [Saf12] Michael G. Safonov. Origins of robust control: Early history and future speculations. *Annual Reviews in Control*, 36:173–181, 2012.
- [Sar91] D. Sarid. *Atomic Force Microscopy*. Oxford University Press, Oxford, UK, 1991.
- [Sas99] S. Sastry. *Nonlinear Systems*. Springer, New York, 1999.
- [SC92] R. Shishko and R. G. Chamberlain. NASA systems engineering handbook. Technical report, National Aeronautics and Space Administration, 1992.
- [Sch87] M. Schwartz. *Telecommunication Networks*. Addison Wesley, Reading, MA, 1987.
- [Sch01] G. Schitter. High performance feedback for fast scanning atomic force microscopes. *Review of Scientific Instruments*, 72(8):3320–3327, 2001.
- [SEM04] D. E. Seborg, T. F. Edgar, and D. A. Mellichamp. *Process Dynamics and Control*. Wiley, Hoboken, NJ, 2nd edition, 2004.
- [Sen01] S. D. Senturia. *Microsystem Design*. Kluwer, Boston, MA, 2001.
- [Sha01] L. Sha. Using simplicity to control complexity. *IEEE Software*, 18:20–28, 2001.
- [Shi96] F. G. Shinskey. *Process-Control Systems. Application, Design, and Tuning*. McGraw-Hill, New York, 4th edition, 1996.
- [SHM⁺16a] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [SHM⁺16b] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, and others. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [Sim56] H. A. Simon. Dynamic programming under uncertainty with a quadratic criterion function. *Econometrica*, 24:74–81, 1956.
- [SLL16] Li Sun, Donghai Li, and Kwang Y. Lee. Optimal disturbance rejection for pi controller with constraints on relative delay margin. *ISA Transactions*, 2016.
- [Smi52] J. M. Smith. The importance of the nervous system in the evolution of animal flight. *Evolution*, 6:127–129, 1952.
- [Smi57a] O. J. M. Smith. Closed control of loops with dead time. *Chemical Engineering Progress*, 53:217–219, 1957.

- [Smi57b] O. J. M. Smith. Posicast control of damped oscillatory systems. *Proc. IRE*, 45:1249–1255, 1957.
- [Son98] E. P. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer, New York, 2nd edition, 1998.
- [SP05] S. Skogestad and I. Postlethwaite. *Multivariable Feedback Control*. Wiley, Hoboken, NJ, 2nd edition, 2005.
- [SS02] E. B. Saf and A. D. Snider. *Fundamentals of Complex Analysis with Applications to Engineering, Science and Mathematics*. Prentice Hall, Englewood Cliffs, NJ, 2002.
- [SSS⁺04] Jörg Stelling, Uwe Sauer, Zoltan Szallasi, Francis J Doyle III, and John Doyle. Robustness of cellular functions. *Cell*, 118(6):675–685, 2004.
- [Sta68] L. Stark. *Neurological Control Systems—Studies in Bioengineering*. Plenum Press, New York, 1968.
- [Ste80] G. Stein. Adaptive flight control: A pragmatic view. In K. S. Narendra and R. V. Monopoli, editors, *Applications of Adaptive Control*. Academic Press, 1980.
- [Ste02] J. Stewart. *Calculus: Early Transcendentals*. Brooks Cole, Pacific Grove, CA, 2002.
- [Ste03] G. Stein. Respect the unstable. *Control Systems Magazine*, 23(4):12–25, 2003.
- [Str88] G. Strang. *Linear Algebra and Its Applications*. Harcourt Brace Jovanovich, San Diego, 3rd edition, 1988.
- [Str94] S. H. Strogatz. *Nonlinear Dynamics and Chaos, with Applications to Physics, Biology, Chemistry, and Engineering*. Addison-Wesley, Reading, MA, 1994.
- [SV89] M. W. Spong and M. Vidyasagar. *Dynamics and Control of Robot Manipulators*. John Wiley, 1989.
- [SZ14] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Tan96] A. S. Tannenbaum. *Computer Networks*. Prentice Hall, Upper Saddle River, NJ, 3rd edition, 1996.
- [Teo37] T. Teorell. Kinetics of distribution of substances administered to the body, I and II. *Archives Internationales de Pharmacodynamie et de Therapie*, 57:205–240, 1937.
- [Tha89] G. T. Thaler. *Automatic Control Systems*. West Publishing, St. Paul, MN, 1989.
- [THÅ00] Tore T. Häggglund and K. J. Åström. Supervision of adaptive control algorithms. *Automatica*, 36(8):1171–1180, 2000.
- [Til01] M. Tiller. *Introduction to Physical Modeling with Modelica*. Springer, Berlin, 2001.
- [Tru55] J. G. Truxal. *Automatic Feedback Control System Synthesis*. McGraw-Hill, New York, 1955.
- [TS90] D. Tipper and M. K. Sundareshan. Numerical methods for modeling computer networks under nonstationary conditions. *IEEE Journal of Selected Areas in Communications*, 8(9):1682–1695, 1990.
- [Tsi54] H. S. Tsien. *Engineering Cybernetics*. McGraw-Hill, New York, 1954.
- [Tus52] A. Tustin. Feedback. *Scientific American*, 48–54, 1952.
- [Vid85] M Vidyasagar. *Control Systems Synthesis*. MIT Press, Cambridge, MA, 1985.
- [Vin93] G. Vinnicombe. Frequency domain uncertainty and the graph topology. *IEEE Transactions on Automatic Control*, 38:1371–1383, 1993.
- [Vin01] G. Vinnicombe. *Uncertainty and Feedback: \mathcal{H}_∞ Loop-Shaping and the v -Gap Metric*. Imperial College Press, London, 2001.

- [Wad97] H.L. Wade. Inverted decoupling: a neglected technique. *ISA Transactions*, 36(1):3–10, 1997.
- [Whi99] F. J. W. Whipple. The stability of the motion of a bicycle. *Quarterly Journal of Pure and Applied Mathematics*, 30:312–348, 1899.
- [Wid41] D. V. Widder. *Laplace Transforms*. Princeton University Press, Princeton, NJ, 1941.
- [Wie48] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. Wiley, 1948.
- [Wig90] S. Wiggins. *Introduction to Applied Nonlinear Dynamical Systems and Chaos*. Springer-Verlag, Berlin, 1990.
- [Wil99] H. R. Wilson. *Spikes, Decisions, and Actions: The Dynamical Foundations of Neuroscience*. Oxford University Press, Oxford, UK, 1999.
- [Wil04] D. G. Wilson. *Bicycling Science*. MIT Press, Cambridge, MA, 3rd edition, 2004. With contributions by Jim Papadopoulos.
- [Wis07] K. A. Wise. Guidance and control for military systems: Future challenges. In *AIAA Conference on Guidance, Navigation, and Control*, 2007. AIAA Paper 2007-6867.
- [WT24] E. P. M. Widmark and J. Tandberg. Über die Bedingungen für die Akkumulation indifferenten Narkotika. *Biochemische Zeitung*, 148:358–389, 1924.
- [WW97] F. R. Whitt and D. G. Wilson. *Bicycling Science*. MIT Press, 1997.
- [YH91] S. Yamamoto and I. Hashimoto. Present status and future needs: The view from Japanese industry. In Y. Arkun and W. H. Ray, editors, *Chemical Process Control—CPC IV*, 1991.
- [YHSD00a] T M Yi, Y Huang, M I Simon, and J Doyle. Robust perfect adaptation in bacterial chemotaxis through integral feedback control. *Proceedings of the National Academy of Sciences*, 97(9):4649–4653, 2000.
- [YHSD00b] T. M. Yi, Y. Huang, M. I. Simon, and J. C. Doyle. Robust perfect adaptation in bacterial chemotaxis through integral feedback control. *Proceedings of the National Academy of Sciences*, 97(9):4649–4653, 2000.
- [Zam81] G. Zames. Feedback and optimal sensitivity: Model reference transformations, multiplicative seminorms, and approximative inverse. *IEEE Transactions on Automatic Control*, AC-26(2):301–320, 1981.
- [ZD63] L. A. Zadeh and C. A. Desoer. *Linear System Theory: the State Space Approach*. McGraw-Hill, New York, 1963.
- [ZDG96] J. C. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice Hall, Englewood Cliffs, NJ, 1996.
- [ZN42] J. G. Ziegler and N. B. Nichols. Optimum settings for automatic controllers. *Transactions of the ASME*, 64:759–768, 1942.