Chapter 2

Overview of the Field

Control is a field with broad relevance to a number of engineering applications. Its impact on modern society is both profound and often poorly understood. In this chapter, we provide an overview of the field, illustrated with examples and vignettes, and describe the new environment for control.

2.1 What is Control?

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

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

A typical example of a modern control system is shown in Figure 2.1. The basic elements of of sensing, computation, and actuation are clearly seen. In modern control systems, computation is typically implemented on a digital computer,

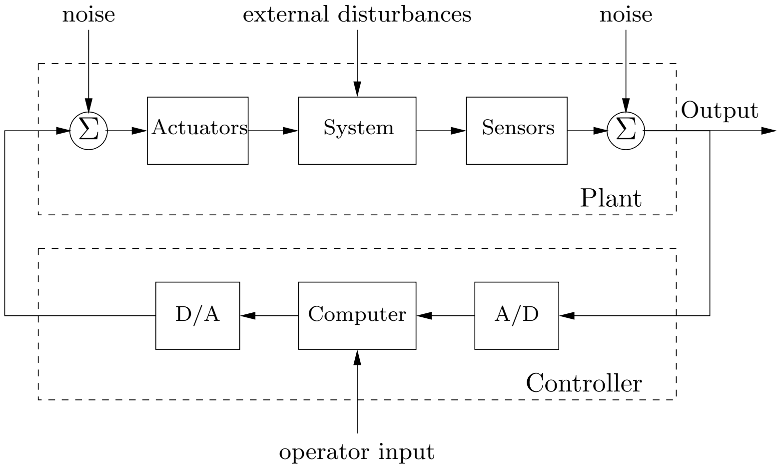

Figure 2.1. *Components of a modern control system.*

requiring the use of analog-to-digital (A/D) and digital-to-analog (D/A) converters. Uncertainty enters the system through noise in sensing and actuation subsystems, external disturbances that affect the underlying system physics, and uncertain dynamics in the physical system (parameter errors, unmodeled effects, etc).

The basic feedback loop of control is often combined with *feedforward* control, where a commanded actuator input is computed to achieve a desired action based on a model of the system. While feedback operates in a closed loop, with actions based on the deviation between measured and desired performance, feedforward operates in open loop, with actions taken based on plans. It is often advantageous to use feedback with feedforward to achieve both high performance and robustness.

It is important to note that while feedback is a central element of control, feedback as a phenomenon is ubiquitous in science and nature. Homeostasis in biological systems maintains thermal, chemical, and biological conditions through feedback. Global climate dynamics depend on the feedback interactions between the atmosphere, oceans, land, and the sun. Ecologies are filled with examples of feedback, resulting in complex interactions between animal and plant life. The dynamics of economies are based on the feedback between individuals and corporations through markets and the exchange of goods and services.

While ideas and tools from control can be applied to these systems, we focus our attention in this report on the application of feedback to engineering systems. We also limit ourselves to a small subset of the many aspects of control, choosing to focus on those that are undergoing the most change and are most in need of new ideas and techniques.

Control Theory

Control *theory* refers to the mathematical framework used to analyze and synthesize control systems. Over the last 50 years, there has been careful attention by control theorists to the issues of completeness and correctness. This includes substantial efforts by mathematicians and engineers to develop a solid foundation for proving stability and robustness of feedback controlled systems, and the development of computational tools that provide guaranteed performance in the presence of uncertainty. This rigor in approach is a hallmark of modern control and is largely responsible for the success it has enjoyed across a variety of disciplines.

It is useful in this context to provide a brief history of the development of modern control theory.

Automatic control traces its roots to the beginning of the industrial revolution, when simple governors were used to automatically maintain steam engine speed despite changes in loads, steam supply, and equipment. In the early 20th Century, the same principles were applied in the emerging field of electronics, yielding feedback amplifiers that automatically maintained constant performance despite large variations in vacuum tube devices.

The foundations of the theory of control are rooted in the 1940s, with the development of methods for single-input, single-output feedback loops, including transfer functions and Bode plots for modeling and analyzing frequency response and stability, and Nyquist plots and gain/phase margin for studying stability of feedback systems [9]. By designing feedback loops to avoid positive reinforcement of disturbances around a closed loop system, one can ensure that the system is stable and disturbances are attenuated. This first generation of techniques is known collectively as "classical control" and is still the standard introduction to controls for engineering students.

In the 1960s, the second generation of control theory, known as "modern control," was developed to provide methods for multi-variable systems where many strongly coupled loops must be designed simultaneously. These tools made use of state space representations of control systems and were coupled with advances in numerical optimization and optimal control. These early state space methods made use of linear ordinary differential equations to study the response of systems, and control was achieved by placing the eigenvalues of the closed loop system to ensure stability.

At around this same time, optimal control theory also made great advances, with the establishment of the maximum principle of Pontryagin and the dynamic programming results of Bellman. Optimal control theory gave precise conditions under which a controller minimized a given cost function, either as an open loop input (such as computing the thrust for optimal trajectory generation) or as a closed loop feedback law. Estimation theory also benefited from results in optimal control, and the Kalman filter was developed and quickly became a standard tool used in many fields to estimate the internal states of a system given a (small) set of measured signals.

Finally, in the 1980s the third generation of control theory, known as "robust multi-variable control," added powerful formal methods to guarantee desired closed loop properties in the face of uncertainties. In many ways, robust control brought back some of the key ideas from the early theory of control, where uncertainty was a dominant factor in the design methodology. Techniques from operator theory were extremely useful here and there was stronger interaction with mathematics, both in terms of using existing techniques and developing new mathematics.

Over the past two decades, many other branches of control have appeared, including adaptive, nonlinear, geometric, hybrid, fuzzy, and neural control frameworks. All of these have built on the tradition of linking applications, theory, and computation to develop practical techniques with rigorous mathematics. Control also built on other disciplines, especially applied mathematics, physics, and operations research.

Today, control theory provides a rich methodology and a supporting set of mathematical principles and tools for analysis and design of feedback systems. It links four important concepts that are central to both engineered and natural systems: dynamics, modeling, interconnection, and uncertainty.

The role of dynamics is central to all control systems and control theory has developed a strong set of tools for analyzing stability and performance of dynamical systems. Through feedback, we can alter the behavior of a system to meet the needs of an application: systems that are unstable can be stabilized, systems that are sluggish can be made responsive, and systems that have drifting operating points can be held constant. Control theory provides a rich collection of techniques to analyze the stability and dynamic response of complex systems and to place bounds on the behavior of such systems by analyzing the gains of linear and nonlinear operators that describe their components. These techniques are particularly useful in the presence of disturbances, parametric uncertainty, and unmodeled dynamics—concepts that are often not treated in detail in traditional dynamics and dynamical systems courses.

Control theory also provides new techniques for (control-oriented) system modeling and identification. Since models play an essential role in analysis and design of feedback systems, sophisticated tools have been developed to build such models. These include input/output representations of systems (how disturbances propagate through the system) and data-driven system identification techniques. The use of “forced response” experiments to build models of systems is well developed in the control field and these tools find application in many disciplines, independent of the use of feedback. A strong theory of modeling has also been developed, allowing rigorous definitions of model fidelity and comparisons to experimental data.

A third key concept in control theory is the role of interconnection between subsystems. Input/output representations of systems allow one to build models of very complex systems by linking component behaviors. The dynamics of the resulting system is determined not only by the dynamics of the components, but by the interconnection structure between these components. The tools of control provide a methodology for studying the characteristics of these interconnections and when they lead to stability, robustness, and desired performance.

Finally, one of the powerful features of modern control theory is that it provides an *explicit* framework for representing uncertainty. Thus, we can describe a “set” of systems that represent the possible instantiations of a system or the possible descriptions of the system as it changes over time. While this framework is important for all of engineering, the control community has developed one of the most powerful collection of tools for dealing with uncertainty. This was necessary

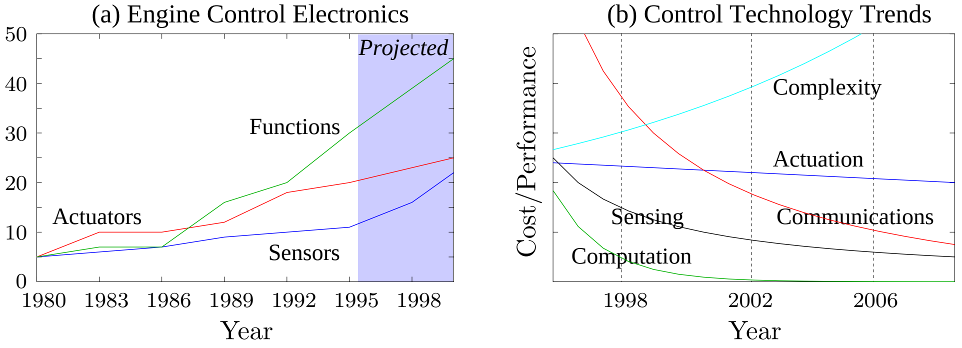

Figure 2.2. *Trends in control technology: (a) the number of sensors, actuators and control functions in engine controls [6] and (b) illustration of cost/performance trends for component technologies.*

because the use of feedback is not entirely benign. In fact, it can lead to catastrophic failure if the uncertainty is not properly managed (through positive feedback, for example).

Control Technology

Control *technology* includes sensing, actuation and computation, used together to produce a working system. Figure 2.2a shows some of the trends in sensing, actuation, and computation in automotive applications. As in many other application areas, the number of sensors, actuators, and microprocessors is increasing dramatically, as new features such as antilock brakes, adaptive cruise control, active restraint systems, and enhanced engine controls are brought to market. The cost/performance curves for these technologies, as illustrated in Figure 2.2b, is also insightful. The costs of electronics technologies, such as sensing, computation, and communications, is decreasing dramatically, enabling more information processing. Perhaps the most important is the role of communications, which is now inexpensive enough to offer many new possibilities.

Control is also closely related to the integration of software into physical systems. Virtually all modern control systems are implemented using digital computers. Often they are just a small part of much larger computing systems performing various other system management tasks. Because of this, control software becomes an integral part of the system design and is an enabler for many new features in products and processes. Online reconfiguration is a fundamental feature of computer controlled systems and this is, at its heart, a control issue.

This trend toward increased use of software in systems is both an opportunity and a challenge for control. As embedded systems become ubiquitous and communication between these systems becomes commonplace, it is possible to design systems that are not only reconfigurable, but also aware of their condition and environment, and interactive with owners, users, and maintainers. These “smart systems" provide improved performance, reduced downtime, and new functionality that was unimaginable before the advent of inexpensive computation, communications, and sensing. However, they also require increasingly sophisticated algorithms to guarantee performance in the face of uncertainty and component failures, and require new paradigms for verifying the software in a timely fashion. Our everyday experience with commercial word processors shows the difficulty involved in getting this right.

One of the emerging areas in control technology is the generation of such real-time embedded software [32]. While often considered within the domain of computer science, the role of dynamics, modeling, interconnection, and uncertainty is increasingly making embedded systems synonymous with control systems. Thus control must embrace software as a key element of control technology and integrate computer science principles and paradigms into the discipline. This has already started in many areas, such as hybrid systems and robotics, where the continuous mathematics of dynamics and control are intersecting with the discrete mathematics of logic and computer science.

Comparison with Other Disciplines

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

A key difference with many scientific disciplines is that control is fundamentally an engineering science. Unlike natural science, whose goal is to understand nature, the goal of engineering science is to understand and develop new systems that can benefit mankind. Typical examples are systems for transportation, electricity, communication and entertainment that have contributed dramatically to the comfort of life. While engineering originally emerged as traditional disciplines such as mining, civil, mechanical, electrical and computing, control emerged as a *systems* discipline around 1950 and cut across these traditional disciplines. The pinnacle of achievement in engineering science is to find new systems principles that are essential for dealing with complex man-made systems. Feedback is such a principle and it has had a profound impact on engineering systems.

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

Control share many tools with the field of operations research. Optimization and differential games play central roles in each, and both solve problems of asset allocation in the face of uncertainty. The role of dynamics and interconnection (feedback) is much more ingrained within control, as well as the concepts of stability and dynamic performance.

Control is also closely associated with computer science, since virtually all modern control algorithms are implemented in software. However, control algorithms and software are very different from traditional computer software. The physics (dynamics) of the system are paramount in analyzing and designing them and their (hard) real-time nature dominates issues of their implementation. From a software-centric perspective, an F-16 is simply another peripheral, while from a control-centric perspective, the computer is just another implementation medium for the feedback law. Neither of these are adequate abstractions, and this is one of the key areas identified in this report as both an opportunity and a need.

2.2 Control System Examples

Control systems are all around us in the modern technological world. They maintain the environment, lighting, and power in our buildings and factories, they regulate the operation of our cars, consumer electronics, and manufacturing processes, they enable our transportation and communications systems, and they are critical elements in our military and space systems. For the most part, they are hidden from view, buried within the code of processors, executing their functions accurately and reliably. Nevertheless, their existence is a major intellectual and engineering accomplishment that is still evolving and growing, promising ever more important consequences to society.

Early Examples

The proliferation of control in engineered systems has occurred primarily in the latter half of the 20th Century. There are some familiar exceptions, such as the Watt governor described earlier and the thermostat (Figure 2.3a), designed at the turn of the century to regulate temperature of buildings.

The thermostat, in particular, is often cited as a simple example of feedback control that everyone can understand. Namely, the device measures the temperature in a building, compares that temperature to a desired set point, and uses the "feedback error" between these two to operate the heating plant, e.g., to turn heating on when the temperature is too low and to turn if off when temperature is too high. This explanation captures the essence of feedback, but it is a bit too simple even for a basic device such as the thermostat. Actually, because lags and delays exist in the heating plant and sensor, a good thermostat does a bit of anticipation, turning the plant off before the error actually changes sign. This avoids excessive temperature swings and cycling of the heating plant.

This modification illustrates that, even in simple cases, good control system design it not entirely trivial. It must take into account the dynamic behavior of the object being controlled in order to do a good job. The more complex the dynamic behavior, the more elaborate the modifications. In fact, the development of

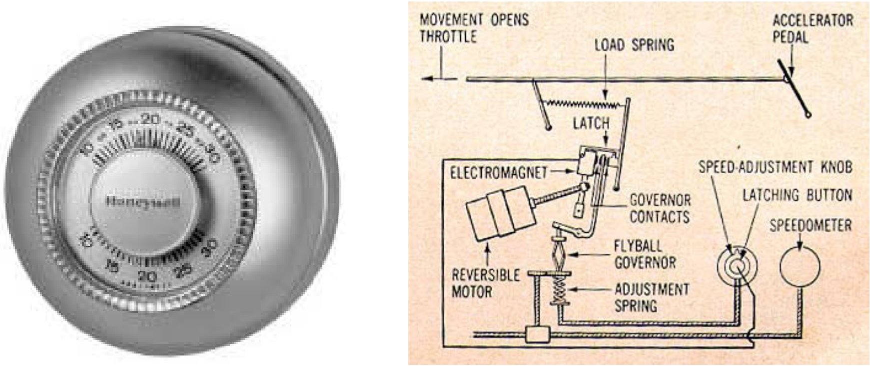

(a) (b)

Figure 2.3. *Early control devices: (a) Honeywell T86 thermostat, originally introduced in 1953, (b) Chrysler cruise control system, introduced in the 1958 Chrysler Imperial (note the centrifugal governor) [21].*

a thorough theoretical understanding of the relationship between dynamic behavior and good controllers constitutes the most significant intellectual accomplishment of the control community, and the codification of this understanding into powerful computer aided engineering design tools makes all modern control systems possible.

There are many other control system examples, of course, that have developed over the years with progressively increasing levels of sophistication and impact. An early system with broad public exposure was the "cruise control" option introduced on automobiles in 1958 (see Figure 2.3b). With cruise control, ordinary people experienced the dynamic behavior of closed loop feedback systems in action—the slowdown error as the system climbs a grade, the gradual reduction of that error due to integral action in the controller, the small (but unavoidable) overshoot at the top of the climb, etc. More importantly, by experiencing these systems operating reliably and robustly, the public learned to trust and accept feedback systems, permitting their increasing proliferation all around us. Later control systems on automobiles have had more concrete impact, such as emission controls and fuel metering systems that have achieved major reductions of pollutants and increases in fuel economy.

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

Aerospace Applications

Similarly, control systems have been critical enablers in the aerospace and military world. We are familiar, for example, with the saturation bombing campaigns of

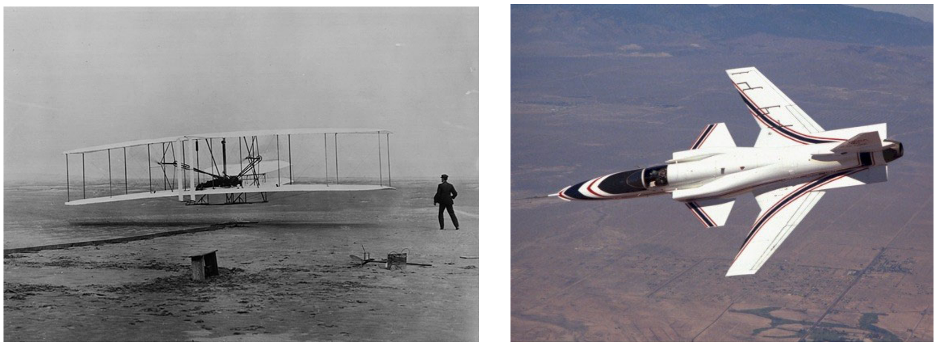
(a) (b)

Figure 2.4. *Flight systems: (a) 1903 Wright Flyer, (b) X-29 forward swept wing aircraft, in 1987. X-29 photograph courtesy of NASA Dryden Flight Research Center.*

World War II, which dropped unguided explosives almost indiscriminately on population centers in order to destroy selected industrial or military targets. These have been replaced with precision guided weapons with uncanny accuracy, a single round for a single target. This is enabled by advanced control systems, combining inertial guidance sensors , radar and infrared homing seekers, satellite navigation updates from the global positioning system, and sophisticated processing of the "feedback error," all combined in an affordably disposable package.

We are also familiar with early space launches. Slender rockets balanced precariously on the launch pad, failing too often in out-of-control tumbles or fireballs shortly after ignition. Robust, reliable, and well-designed control systems are not optional here, because boosters themselves are unstable. And control systems have lived up to this challenge. We now take routine launch operations for granted, supporting manned space stations, probes to the outer planets, and a host of satellites for communications, navigation, surveillance, and earth observation missions. Of course, these payloads are themselves critically dependent on robust, reliable and well-designed control systems for everything from attitude control, to on-orbit station-keeping, thermal management, momentum management, communications, etc.

Flight Control

Another notable success story for control in the aerospace world comes from the control of flight. This example illustrates just how significant the intellectual and technological accomplishments of control have been and how important their continued evolution will be in the future.

Control has played a key role in the development of aircraft from the very beginning. Indeed, the Wright brother's first powered flight was successful only because the aircraft included control surfaces (warpable wings and forward-mounted vertical and horizontal fins) that were adjusted continuously by the pilot to stabilize the flight [19] (see Figure 2.4a). These adjustments were critical because the Wright Flyer itself was unstable, and could not maintain steady flight on its own.

Because pilot workload is high when flying unstable aircraft, most early aircraft that followed the Wright Flyer were designed to be statically stable. Still, as the size and performance capabilities of aircraft grew, their handling characteristics deteriorated. Designers then installed so-called "stability augmentation systems"—automatic control systems designed to modify dynamic behavior of aircraft slightly in order to make them easier to fly. These systems first appeared during the World War II years. They used early inertial sensors to measure flight motions, analog electronic systems to construct and process feedback errors, and hydraulic systems to actuate the linkages of selected control surfaces (vertical and horizontal tails, ailerons, etc).

Two issues surfaced immediately as these systems were being fielded: (1) how to design the control logic systematically (early systems were essentially developed by trial-and-error), and (2) how to build the systems such that they would operate reliably. Early systems proved to be quite unreliable. Hence, only a small fraction of the full authority of the control surfaces was typically allocated to the automatic system, with the bulk of authority reserved for manual control, so the pilot could always override the automation.

Control theorists provided the solution for the first issue. They developed modeling and simulation methods (based on differential equations and transfer functions) that accurately describe aircraft dynamics, and they developed increasingly powerful generations of control analysis and design methods to design control laws. Classical control methods enabled the systematic design of early stability augmentation systems, while modern control and robust multi-variable control are critical in all of today's modern flight systems.

But analysis and design methods alone could not address the second issue of early stability augmentation systems, namely the need for highly reliable control implementations. That issue was resolved with the development of airborne digital computers and redundant architectures. These are now routinely used on all commercial and military aircraft. They have become so highly reliable that the old solution of granting only partial authority to automation has long been abandoned. In fact, most modern flight control implementations do not even include mechanical linkages between pilots and control surfaces. All sensed signals and control commands go through the digital implementation (e.g., fly-by-wire).

Today, we even entrust the very survival of aircraft to automation. Examples include the all weather auto-land functions of commercial transports, in which safe go-around maneuvers are not available if failures were to occur at certain critical flight phases. Other examples include the F-16, B-2, and X-29 military aircraft (see Figure 2.4), whose basic dynamics are unstable like the Wright Flyer, but so much more violently that manual stabilization is not possible. Finally, in modern flight systems there is a growing trend to automate more and more functions—all the way to removing the pilot entirely from the cockpit. This is already commonplace in certain military reconnaissance and surveillance missions and will soon be extended to more lethal ones, such as suppressing enemy air defenses with unmanned aerial vehicles (UAVs).

The following vignette describes some of these advances, from the perspective of one of its successful practitioners.

Vignette: Fighter Aircraft and Missiles (Kevin A. Wise, The Boeing Company)

The 1990s has been a decade of significant accomplishments and change for the aerospace community. New systems such as unstable, tailless aircraft, propulsion controlled ejection seats, and low-cost, accurate, GPS guided munitions were developed. Fly-by-wire flight control systems have become the standard, making control system design and analysis central to military aircraft and missile system development. Improving pilot safety and reducing costs were key focus areas in industry.

Flight control system design methods using feedback linearization paved the way for new gain scheduled flight control systems for aircraft. This method, applied to the X-36 Tailless Agility Research aircraft and the F-15 ACTIVE, uniquely allows engineers to better design flying qualities into the aircraft, reducing design and development costs and improving pilot acceptance. Advances in robustness theory improved analysis tools allowing engineers to accurately predict and thus expand departure boundaries for these highly unstable aircraft. To further improve safety, these control laws were augmented with neural networks for reconfigurable and damage adaptive flight control.

Missile systems, such as the Joint Direct Attack Munition (JDAM) and the Miniaturized Munition Technology Demonstrator (MMTD) developed their flight control designs using state feedback optimal control, and then projecting out those states not measured by sensors . This method eliminated sensor hardware, reducing weight and costs, and proved to be completely automatable. The Fourth Generation Escape System (GEN4) ejection seat also used this approach for its control laws. In addition to needing optimal performance, advances in robustness theory were used to characterize the seat's control system performance to uncertain crew member size and weight (95% male to 5% female). Autocode software tools for implementing controls systems also emerged in the 1990s. These computer aided design tools provide a single environment for control design and analysis as well as software design and test. They have greatly reduced the implementation and testing costs of flight control systems.

The new challenge faced by the control community is the development of unmanned combat systems (munitions as well as aircraft) and concepts of operations for these systems to address the intelligent, increasingly hostile, rapidly changing environments faced by our war fighters. These systems must detect, identify, locate, prioritize, and employ ordinance to achieve permanent destruction of high value targets. New developments in intelligent control, vision based control, mission planning, path planning, decision aiding, communication architectures, logistics and support concepts, and last but not least, software development, validation, and verification are needed to support these systems and make them affordable.

2.3 The Increasing Role of Information-Based Systems

Early applications of control focused on the physics of the system being controlled, whether it was the thermal dynamics of buildings, the flight mechanics of an airplane, or the tracking properties of a disk drive head. The situation we now face is one in which pervasive computing, sensing, and communications are common and the way that we interact with machines and they interact with each other is changing rapidly. The consequences of this tremendous increase in information are also manifest in control, where we are now facing the challenges of controlling large-scale systems and networks that are well beyond the size and complexity of the traditional applications of control.

One indication of this shift is the role that embedded systems and software play in modern technology, described briefly above. Modern computer control systems are capable of enormous amounts of decision making and control logic. Increasingly, these software systems are interacting with physical processes and introducing feedback algorithms to improve performance and robustness. Already, the amount of logic-based code is overshadowing the traditional control algorithms in many applications. Much of this logic is interwoven with the closed loop performance of the system, but systematic methods for analysis, verification, and design have yet to be developed.

Another area where control of information-based systems will be increasingly important is in resource allocation systems. In this context, control can be described as the science and engineering of optimal dynamic resource allocation under uncertainty. We start with a mathematical model, of a system that describes how current actions or decisions can affect the future behavior of the system, including our uncertainty in that behavior. "Resource allocation" means that our decisions can be interpreted as managing a tradeoff between competing goals, or choosing from a limited set of possible actions. "Uncertainty" is critical: there is some possible variation in the system's behavior, so that decisions have to be made taking different possibilities into account. Sources of uncertainty include incomplete or corrupted information available to the decision maker, uncertainty in the mathematical model used to model the system, and unpredictability of commands due to noise and disturbance signals that affect the system. While often considered an operations research problem, the role of dynamics and instabilities points to a clear need for control theory as well.

One of the consequences of this shift toward information-based systems is that we are moving from an era where physics was the bottleneck to progress to one in which complexity is the bottleneck.

There are already many examples of this new class of systems that are being deployed. Congestion control in routers for the Internet, power control in wireless communications systems, and real-time use of information in service and supply chains are a few examples. In all of these systems, it is the interaction of information flow with the underlying physics that is responsible for the overall performance. Another example is the air traffic control network, where the density of flights, demand for efficiency, and intolerance for failure have created a situation that couples

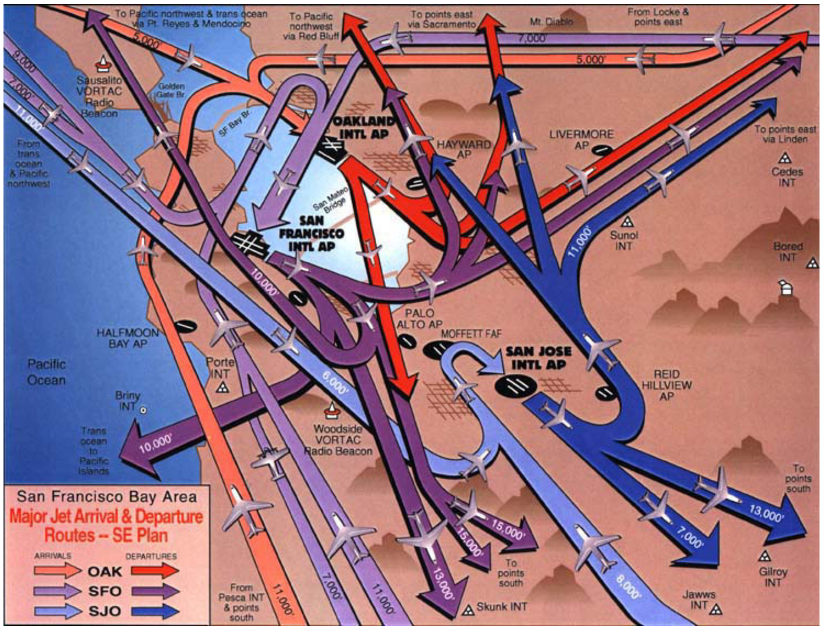

Figure 2.5. *San Francisco Bay area aircraft arrival and departure routes. Figure courtesy of Federal Aviation Authority.*

vast amounts of information—everything from the location of the planes to the individual customer itineraries—that must be managed to maintain high performance, robust, and reliable operation at all times. Figure 2.5 shows just one small part of this problem, the local departure and arrival routes in the San Francisco Bay area.

There is an important role for control in many of these applications. As in traditional application areas, control serves as a mechanism for getting both information and, more importantly, *action* out of data. Furthermore, the theory of control provides insights and tools for analyzing and designing interconnected systems with desirable stability and robustness properties.

One fundamental change in the use of control is the role of communications and networking. This will radically change the use of feedback in modern systems, through increased access to large amounts of information as well as the new environment in which control systems will have to operate. Control computations must increasingly be done in a distributed, partially asynchronous environment where tight control over the timing of data and computation is not available, due for example to the existence of packet-based communications networks between sensing, actuation, and computational nodes. Many traditional approaches may no longer work in this context and we anticipate the need to develop new paradigms for designing robust, high performance, feedback systems in this information rich environment.

The role of uncertainty in information rich systems is also critical (and largely unexplored) and concepts from control will play an important role in managing this uncertainty in the analysis, design, and operation of large-scale, interconnected systems. Uncertainty must be represented in order to build tractable models for answering questions that take into account the whole range of possible variations in the details of components and their interconnections. Control ideas will be increasingly important as a tool for managing both the complexity and uncertainty in these systems, and must be made available to the designers of such systems, through education and design software. One aspect of this that is likely to be particularly important is the exploration of fundamental limits of performance, robustness, and stability, since tradeoffs between these will be the primary design challenge in this space.

Examples of the need for increased development in this area can be seen in the applications discussed in the next chapter. Vehicle, mission, and airspace management systems for transportation; source, power, and router control for networks; and genetic, cellular, and integrative feedback networks in biological systems are just a few examples. The simplest of these problems lies at the boundaries of current tools and understanding, and future progress will require a much deeper understanding of the integration between control, communications, computing, and networks as well as modeling, analysis, and control of complex decision systems.

2.4 Opportunities and Challenges Facing the Field

Control has developed into a major field in which generations of engineers are able to solve problems of practical importance and enormous impact. Over the past few years, the opportunities for control have expanded enormously, but there are many challenges that must be addressed to realize the potential for impact. In this section we attempt to characterize some of the overarching themes that describe these opportunities and challenges, and recommend an approach for moving forward.

Characteristics of the New Environment

The future of control will be driven by a new environment that differs substantially from that of the past 40 years. Some of the features of this new environment are already apparent and provide insight into the new research directions that must be pursued.

Ubiquitous Computation, Communication and Sensing. The dominant change in the engineering environment is the presence of ever more powerful computation and cheaper communication. The new software and storage products that these developments have spawned have further changed the engineering landscape in many areas. In addition, microelectronics and MEMS have made available inexpensive sensors , such as those shown in Figure 2.6, and new actuator concepts that can be made available via communication networks, allowing increasingly sensor rich and actuator rich control.

It will require decades to take full advantage of these developments. Some innovation will involve standalone improvements to individual systems and some

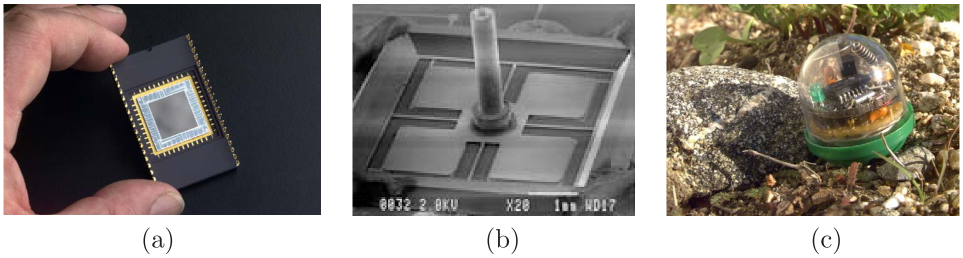

(a) (b) (c)

Figure 2.6. *Examples of current sensor technology: (a) 1024x1024 CCD array, (b) MEMS-based microgryoscope, and (c) sensor web pod. All photographs courtesy of Jet Propulsion Laboratory.*

will involve extreme interconnectedness of the type seen in the telephone system, the power grid, the Internet, and their descendants. Both types may, and probably will, depend on the use of control. The new ideas required to be successful in the two cases are, however, likely to be qualitatively different because we do not yet have a great deal of experience in building and operating safe, reliable, highly interconnected systems.

New Application Domains. In addition to the revolutionary changes in information technology, future control systems will involve interactions between physical, chemical, biological, and information sciences, integrated with algorithms and feedback. This will open up new application domains for control, such as biological engineering and quantum systems. While there are already researchers within the control community that are attacking problems in these areas, it will be necessary to educate new generations of researchers in both control and other disciplines in order to make advances in these applications. The possibilities for control are potentially very fundamental in nature, as illustrated in the following vignette.

Vignette: Quantum Measurement and Control (Hideo Mabuchi, Caltech)

To illustrate the applications of control in new domains, consider the research of Hideo Mabuchi, who is exploring the use of feedback and control in quantum systems and its implications for unifying quantum and classical physics:

> A grand enigma, which is perhaps our primary legacy from 20th Century physics, is that the states and dynamics we ascribe to microscopic (quantum) systems seem incompatible with macroscopic (classical) phenomenology. For example, physical theory claims that it should be illogical simultaneously to assign definite values to certain sets of measurable properties of a quantum system. And yet we want to believe that quantum mechanics is a correct description of microscopic physics, which evolves robustly into classical dynamics for systems of sufficiently large size and with a sufficiently high degree of interconnection among their manifold degrees of freedom.

> How can we understand the consistency of quantum mechanics, as a microscopic theory, with classical physics as a manifestly valid description of macroscopic phenomena?

Control theory provides a new set of tools for understanding quantum systems. One set of tools is through systematic techniques for model reduction:

> Viewed from a "multiscale" perspective, our challenge in explaining the quantum-classical transition will be to show that classical physics can rigorously be obtained as a robust and parsimonious approximation to the dynamics of certain aggregate degrees of freedom for generic complex quantum systems. In the language of control theory, one would like to derive classical physics as an optimal model reduction of quantum physics. A number of fundamental questions arise as soon as the problem is posed this way. How can this model reduction be so general and robust, depending only upon the structure of quantum theory and not the details of any particular dynamical system? What are the general parameters that control the error bounds on this model reduction? What impact will this program have, if successful, on our basic interpretation of quantum mechanics?

In addition, control can provide new techniques for doing experiments, allowing us to better explore physical understanding:

> ... we hope that feedback control will provide a crucial experimental methodology for scrutinizing the validity of quantum measurement theory in realistic laboratory scenarios, especially with regard to the equations for conditional evolution of a system under continuous observation. Such equations could be used as the starting point for controller synthesis, for example, and their validity would be assessed by comparison of experimentally observed closed-loop behavior with theoretical expectations.

Mabuchi's work illustrates the potential power of control theory as a disruptive technology for understanding the world around us.

Reliable Systems with Unreliable Parts. Most reasonably complex man-made systems are not rendered inoperable by the failure of any particular component and biological systems often demonstrate remarkable robustness in this regard. Simple redundancy, or the spare parts approach to such problems, is of limited effectiveness because it is uneconomical. Designs that allow the system to reconfigure itself when a component fails, even if this degrades the performance roughly in proportion to the magnitude of the failure, are usually preferred. Although computer memory chips and disk drive controllers often take advantage of strategies of this type, it is still true that the design of self healing systems is not well studied or analyzed.

This issue takes on considerable significance when dealing with interconnected systems of the complexity of the Internet. In this case there are billions of components and yet the system is so essential that little downtime can be tolerated.

Complexity. Air traffic control systems, power grid control systems and other large-scale, interconnected systems are typical of a class of problems whose complexity is fixed not by the designer but rather by economic considerations and the natural scale of the problem. An acceptable solution in this context must be capable of dealing with the given complexity. In deciding if a system can be built or not, it is important to correctly gauge the feasibility because there is no value in a product that “almost” works.

Every discipline has methods for dealing with some types of complexity. In the physical sciences, for example, the tools developed for studying statistical mechanics have lead to a very substantial body of literature, effective for solving some problems. However, in discussing complexity it is one thing to find a point of view from which aspects of the behavior is compressible (e.g., “the entropy of a closed system can only increase”) but it is another to have a “theory of complex systems”. The latter is something of an oxymoron, in that it suggests that the system is not really complex. On the other hand, it does make sense to seek to understand and organize the methodologies which have proven to be useful in the design of highly interconnected systems and to study naturally occurring systems with this in mind. Engineers looking at the immune system may very well be able to suggest new methods to defeat Internet viruses and ideas from neuroscience may inspire new developments in building reliable systems using unreliable components.

Vision for the Future

This new environment for control presents many challenges, but also many opportunities for impact across a broad variety of application areas. The future directions in control, dynamics, and systems must continue to address fundamental issues, guided by new applications.

One of the biggest challenges facing the field is the integration of computation, communications, and control. As computing, communications, and sensing become more ubiquitous, the use of control will become increasingly ubiquitous as well. However, many of the standard paradigms that allow the separation of these different disciplines will no longer be valid. For example, the ability to separate the computational architecture from the functions that are being computed is already beginning to unravel as we look at distributed systems with redundant, intermittent, and sometimes unreliable computational elements. Beyond simply looking at hybrid systems, a theory must be developed that integrates computer science and control.

Similarly, the simplification that two nodes that are connected can communicate with sufficient reliability and bandwidth such that the properties of the communications channel can be ignored no longer holds in the highly networked environment of the future. Control must become more integrated with the protocols of communications so that high response feedback loops are able to use the same channels as high throughput, lower bandwidth information, without interfering with each other.

Another element of the future of control is to begin to understand analysis and synthesis of control using higher levels of decision making. Traditionally control has dealt with the problem of keeping a few variables constant (regulation) or making variables follow specified time functions (tracking). In robotics, control was faced with more complicated problems such as obstacle avoidance and path planning (task-based control). Future systems will require that control be applied to problems that cannot necessarily be expressed in terms of continuous variables, but rather have symbolic, linguistic, or protocol-based descriptions. This is required as we move to more sophisticated autonomous and semi-autonomous systems that require high-level decision making capabilities.

At the same time as control moves to higher levels of decision making, it will also move to new domains that are only beginning to emerge at the present time. This includes biological, quantum and environmental systems; software systems; enterprise level systems; and economic and financial systems. In all of these new problem domains, it will be necessary to develop a *rigorous* theory of control. This has been a historical strength of the field and has allowed it to be successful in an enormous number of systems.

Finally, we envision an increased awareness of control principles in science and engineering, including much more exposure to feedback systems in math and science education.

Approach

The opportunities and challenges describe here should be addressed on two fronts. There is a need for a broadly supported, active research program whose goals are to explore and further develop methodologies for design and operation of reliable and robust highly interactive systems, and there is a need to make room in the academic programs for material specific to this area.

The research program must be better integrated with research activities in other disciplines and include scientists trained in software engineering, molecular biology, statistical mechanics, systems engineering and psychology. Control researchers must continue to branch out beyond traditional discipline boundaries and become participants and contributors in areas such as computer science, biology, economics, environmental science, materials science and operations research. There is particular need for increased control research in information-based systems, including communications, software, verification and validation, and logistics.

To support this broader research program, a renewed academic program must also be developed. This program should strengthen the systems view and stretch across traditional discipline boundaries. To do so, it will be necessary to provide better dissemination of tools to new communities and provide a broader education for control engineers and researchers. This will require considerable effort to present current knowledge in a more compact way and to allow new results in software, communications, and emerging application domains to be added, while maintaining the key principles of control on which new results will rest. Simultaneously, the control community must seek to increase exposure to feedback in math and science education at all levels, even K-12. Feedback is a fundamental principle that should be part of every technically literate person's knowledge base.

One of the characteristics of the control field has been an emphasis on theory and mathematical formulations of the problems being considered. This discipline has resulted in a body of work that is reliable and unambiguous. Moreover, because this style appeals to some very able graduate students, it has been an important factor in maintaining the flow of talent into the field. However, for engineers and scientists this has been a barrier to entry and can make it difficult for outsiders to assimilate and use the work in their own field. In addition, it has sometimes had a chilling effect on the development of ideas that are not easily translated into mathematical form. The challenge presented by the need to steer a course between the possible extremes here is not new, it has always been present. What is new is that the availability of easily used simulation tools has made the use of heuristic reasoning both more appealing and more reliable. In particular, optimization involving problems that are so large and/or so badly non-convex that rigorous analysis is infeasible can now be approached using principled heuristics. Because of the software and computing power now available this may be the most effective way to proceed. It is important to find a place for effective heuristics in the training of students and the highest level professional meetings of the field.

Finally, experimentation on representative systems must be an integral part of the control community's approach. The continued growth of experiments, both in education and research, should be supported and new experiments that reflect the new environment will need to be developed. These experiments are important for the insight into application domains that they bring, as well as the development of software and algorithms for applying new theory. But they also form the training ground for systems engineers, who learn about modeling, robustness, interconnection, and data analysis through their experiences on real systems.

The recommendations of the Panel, detailed in Chapter 5, provide a high level plan for implementing this basic approach. The recommendations focus on the need to pursue vigorously new application domains and, in particular, those domains in which the principles of control will be essential for future progress. They also highlight the need to maintain the field's strong theoretical base and historical rigor, while at the same time finding new ways to broaden the exposure and use of control to a broader collection of scientists and engineers.

The new environment that control faces is one with many new challenges and an enormous array of opportunities. Advancing the state of the art will require that that the community accelerate its integration across disciplines and look beyond the current paradigms to tackle the next generation of applications. In the next chapter, we explore some of the application areas in more detail and identify some of the specific advancements that will be required.