

Control in an Information Rich World

Report of the Panel on Future Directions in
Control, Dynamics, and Systems

26 April 2002

Abstract

The field of *control* provides the principles and methods used to design physical and information systems that maintain desirable performance by sensing and automatically adapting to changes in the environment. Over the last forty years the field has seen huge advances, leveraging technology improvements in sensing and computation with breakthroughs in the underlying principles and mathematics. Automatic control systems now play critical roles in many fields, including manufacturing, electronics, communications, transportation, computers and networks, and many military systems.

As we begin the 21st Century, the opportunities to apply control principles and methods are exploding. Computation, communication and sensing are becoming increasingly cheap and ubiquitous, with more and more devices including embedded processors, sensors, and networking hardware. This will make possible the development of machines with a degree of intelligence and reactivity that will influence nearly every aspect of life on this planet, including not just the products available, but the very environment in which we live.

New developments in this increasingly information rich world will require a significant expansion of the basic tool sets of control. The complexity of the control ideas involved in the operation of the Internet, semi-autonomous systems, and enterprise-wide supply chain management, for example, are on the boundary of what can be done with available methods. Future applications in aerospace and transportation, information and networks, robotics and intelligent machines, biology and medicine, and materials and processing will create systems that are well beyond our current levels of complexity, and new research is required to enable such developments.

The purpose of this report is to spell out some of the prospects for control in the current and future technological environment, to describe the role the field will to play in military, commercial, and scientific applications over the next decade, and to recommend actions required to enable new breakthroughs in engineering and technology through application of control research.

Panel Membership

Richard M. Murray (chair)
California Institute of Technology

Karl J. Åström
Lund Institute of Technology

Pramod P. Khargonekar
University of Florida

Stephen P. Boyd
Stanford University

P. R. Kumar
University of Illinois

Siva S. Banda
Air Force Research Laboratory

P. S. Krishnaprasad
University of Maryland

Roger W. Brockett
Harvard University

Greg J. McRae
Massachusetts Institute of Technology

John A. Burns
Virginia Tech

Jerrold E. Marsden
California Institute of Technology

Munzer A. Dahleh
Massachusetts Institute of Technology

George Meyer
NASA Ames Research Center

John C. Doyle
California Institute of Technology

William F. Powers
Ford Motor Company

John Guckenheimer
Cornell University

Gunter Stein
Honeywell International

Charles J. Holland
Department of Defense

Pravin Varaiya
University of California, Berkeley

Additional Contributors

Jim Batterson, Richard Braatz, Dennis Bernstein, Joel Burdick, Raffaello D'Andrea, Michael Dickinson, Frank Doyle, Martha Gallivan, Jonathan How, Marc Jacobs, Jared Leadbetter, Jesse Leitner, Steven Low, Hideo Mabuchi, Dianne Newman, John Seinfeld, Eduardo Sontag, Anna Stefanopoulou, Allen Tannenbaum, Claire Tomlin, Kevin Wise

Contents

Preface	vii
1 Executive Summary	1
2 Overview of the Field	7
2.1 What is Control?	7
2.2 Control System Examples	13
2.3 The Increasing Role of Information-Based Systems	17
2.4 Opportunities and Challenges Now Facing Us	20
3 Applications, Opportunities, and Challenges	27
3.1 Aerospace and Transportation	29
3.2 Information and Networks	39
3.3 Robotics and Intelligent Machines	49
3.4 Biology and Medicine	58
3.5 Materials and Processing	65
3.6 Other Applications	72
4 Education and Outreach	79
4.1 The New Environment for Control Education	79
4.2 Making Control More Accessible	81
4.3 Broadening Control Education	83
4.4 The Opportunities in K-12 Math and Science Education	84
4.5 Other Opportunities and Trends	85
5 Recommendations	89
5.1 Integrated Control, Computation, Communications	89
5.2 Control of Complex Decision Systems	90
5.3 High-Risk, Long-Range Applications of Control	91
5.4 Support for Theory and Interaction with Mathematics	92
5.5 New Approaches to Education and Outreach	93
5.6 Concluding Remarks	94
A NSF/CSS Workshop on Education	97

Bibliography

101

Preface

This report documents the findings and recommendations of the Panel on Future Directions in Control, Dynamics, and Systems. This committee was formed in April 2000 under initial sponsorship of the Air Force Office of Scientific Research (AFOSR) to provide a renewed vision of future challenges and opportunities in the field, along with recommendations to government agencies, universities, and research organizations to ensure continued progress in areas of importance to the industrial and defense base. The intent of this report is to raise the overall visibility of research in control, highlight its importance in applications of national interest, and indicate some of the key trends that are important for continued vitality of the field.

The panel was chaired by Professor Richard Murray (Caltech) and was formed with the help of an organizing committee consisting of Professor Roger Brockett (Harvard), Professor John Burns (VPI), Professor John Doyle (Caltech) and Dr. Gunter Stein (Honeywell). The remaining panel members are Karl Åström (Lund Institute of Technology), Siva Banda (Air Force Research Lab), Stephen Boyd (Stanford), Munzer Dahleh (MIT), John Guckenheimer (Cornell), Charles Holland (DDR&E), Pramod Khargonekar (University of Florida), P. R. Kumar (University of Illinois), P. S. Krishnaprasad (University of Maryland), Greg McRae (MIT), Jerrold Marsden (Caltech), George Meyer (NASA), William Powers (Ford), and Pravin Varaiya (UC Berkeley). A writing subcommittee consisting of Karl Åström, Stephen Boyd, Roger Brockett, John Doyle, Richard Murray and Gunter Stein helped coordinate the generation of the report.

The Panel held a meeting on 16-17 July 2000 at the University of Maryland, College Park to discuss the state of the field and its future opportunities. The meeting was attended by members of the panel and invited participants from the academia, industry, and government. Additional meetings and discussions were held over the next 15 months, including presentations at DARPA and AFOSR sponsored workshops, meetings with government program managers, and writing committee meetings. The results of these meetings, combined with discussions amongst panel members and within the community at workshops and conferences, form the main basis for the findings and recommendations of this panel.

A web site has been established to provide a central repository for materials generated by the Panel:

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

Copies of this report, links to other sources of information, and presentation materials from the Panel workshop and other meetings can be found there.

Several similar reports and papers highlighting future directions in control came to the Panel's attention during the development of this report. Many members of the panel and participants in the June 2000 workshop were involved in the generation of the 1998 Fleming report [13] and a 1987 *IEEE Transactions on Automatic Control* article [22], both of which provided a roadmap for many of the activities of the last decade and continue to be relevant. More recently, the European Commission sponsored a workshop on future control systems [11] and several other more focused workshops have been held over the last several years [2, 1, 27]. Several recent papers and reports highlighted successes of control [28] and new vistas in control [9, 20]. The Panel also made extensive use of a recent NSF/CSS report on future directions in control engineering education [1], which provided a partial basis for Chapter 4 of the present report.

The Panel would like to thank the control community for its support of this report and the many contributions, comments, and discussions that help form the context and content for the report. We are particularly indebted to Dr. Marc Q. Jacobs for his initiative in the formation of the panel and for his support of the project through AFOSR.

Richard M. Murray

Pasadena, April 2002

Chapter 1

Executive Summary

Rapid advances in computing, communications, and sensing technology offer unprecedented opportunities for the field of control to expand its contributions to the economic and defense needs of the nation. This report presents the findings and recommendations of a panel of experts chartered to examine these opportunities. We present an overview of the field, review its successes and impact, and describe the new challenges ahead. We do not attempt to cover the entire field. Rather, we focus on those areas that are undergoing the most rapid change and that require new approaches to meet the challenges and opportunities that face the community.

Overview of Control

At its simplest, a control system is a device in which a sensed quantity is used to modify the behavior of a system through computation and actuation. Control systems engineering traces its roots to the industrial revolution, to devices such as the flyball governor, shown in Figure 1.1. This device used a flyball mechanism to sense the rotational speed of a steam turbine and adjust the flow of steam into the machine using a series of linkages. By thus regulating the turbine's speed, it provided the safe, reliable, consistent operation that was required to enable the rapid spread of steam-powered factories.

Control played an essential part in the development of technologies such as power, communications, transportation, and manufacturing. Examples include autopilots in military and commercial aircraft (Figure 1.2a), regulation and control of the electrical power grid, and high accuracy positioning of read/write heads in disk drives (Figure 1.2b). Feedback is an enabling technology in a variety of application areas and has been reinvented and patented many times in different contexts.

A modern view of control sees feedback as a tool for uncertainty management. By measuring the operation of a system, comparing it to a reference, and adjusting available control variables, we can cause the system to respond properly even if its dynamic behavior is not exactly known or if external disturbances tend to cause it to respond incorrectly. This is an essential feature in engineering systems since they




Figure 1.1. *The centrifugal governor (a), developed in the 1780s, was an enabler of the successful Watt steam engine (b), which fueled the industrial revolution. Figures courtesy of Cambridge University.*

must operate reliably and efficiently under a variety of conditions. It is precisely this aspect of control as a means of ensuring robustness to uncertainty that explains why feedback control systems are all around us in the modern technological world. They are in our homes, cars and consumer electronics, in our factories and communications systems, and in our transportation, military and space systems.

The use of control is extremely broad and encompasses a number of different applications. These include control of electromechanical systems, where computer-controlled actuators and sensors regulate the behavior of the system; control of electronic systems, where feedback is used to compensate for component variations and provide reliable, repeatable performance; and control of information and decision systems, where limited resources are dynamically allocated based on estimates of future needs. Control principles can also be found in areas such as biology, medicine, and economics, where feedback mechanisms are ever present. Increasingly, control is also a mission critical function in engineering systems: the systems will fail if the control system does not work.

Contributions to the field of control come from many disciplines, including pure and applied mathematics; aerospace, chemical, mechanical, and electrical engineering; operations research and economics; and the physical and biological sciences. The interaction with these different fields is an important part of the history and strength of the field.

Successes and Impact

Over the past 40 years, the advent of analog and digital electronics has allowed control technology to spread far beyond its initial applications, and has made it an enabling technology in many applications. Visible successes from past investment in control include:




Figure 1.2. Applications of control: (a) the Boeing 777 fly-by-wire aircraft and (b) the Seagate Barracuda 36ES2 disk drive. Photographs courtesy of the Boeing Company and Seagate Technology.

- Guidance and control systems for aerospace vehicles, including commercial aircraft, guided missiles, advanced fighter aircraft, launch vehicles, and satellites. These control systems provide stability and tracking in the presence of large environmental and system uncertainties.
- Control systems in the manufacturing industries, from automotive to integrated circuits. Computer controlled machines provide the precise positioning and assembly required for high quality, high yield fabrication of components and products.
- Industrial process control systems, particularly in the hydrocarbon and chemical processing industries. These maintain high product quality by monitoring thousands of sensor signals and making corresponding adjustments to hundreds of valves, heaters, pumps, and other actuators.
- Control of communications systems, including the telephone system, cell phones, and the Internet. Control systems regulate the signal power levels in transmitters and repeaters, manage packet buffers in network routing equipment, and provide adaptive noise cancellation to respond to varying transmission line characteristics.

These applications have had an enormous impact on the productivity of modern society.

In addition to its impact on engineering applications, control has also made significant intellectual contributions. Control theorists and engineers have made rigorous use of and contributions to mathematics, motivated by the need to develop provably correct techniques for design of feedback systems. They have been consistent advocates of the “systems perspective,” and have developed reliable techniques for modeling, analysis, design, and testing that enable development and implemen-

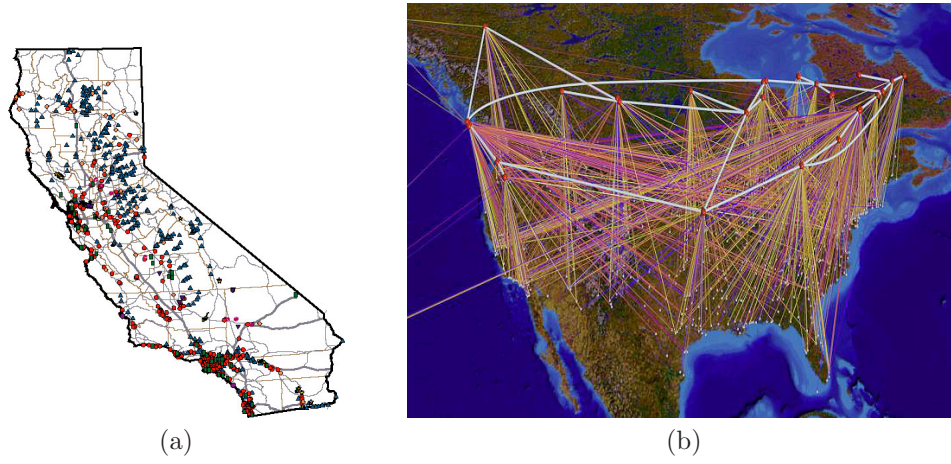


Figure 1.3. Modern networked systems: (a) the California power network and (b) the NSFNET Internet backbone. Figures courtesy of the state of California and the National Center for Supercomputer Applications (NCSA).

tation of the wide variety of very complex engineering systems in use today. Moreover, the control community has been a major source and training ground for people who embrace this systems perspective and who wish to master the substantial set of knowledge and skills it entails.

Future Opportunities and Challenges

As we look forward, the opportunities for new applications that will build on advances in control expand dramatically. The advent of ubiquitous, distributed computation, communication, and sensing systems has begun to create an environment in which we have access to enormous amounts of data and the ability to process and communicate that data in ways that were unimagined 20 years ago. This will have a profound effect on military, commercial and scientific applications, especially as software systems begin to interact with physical systems in more and more integrated ways. Figure 1.3 illustrates two systems where these trends are already evident. Control will be an increasingly essential element of building such interconnected systems, providing high performance, high confidence, and reconfigurable operation in the presence of uncertainties.

In all of these areas, a common feature is that system level requirements far exceed the achievable reliability of individual components. This is precisely where control (in its most general sense) plays a central role, since it allows the system to ensure that it is achieving its goal through correction of its actions based on sensing its current state. The challenge to the field is to go from the traditional view of control systems as a single process with a single controller, to recognizing control systems as a heterogeneous collection of physical and information systems, with intricate interconnections and interactions.

In addition to inexpensive and pervasive computation, communication, and sensing—and the corresponding increased role of information-based systems—an important trend in control is the move from low-level control to higher levels of decision making. This includes such advances as increased autonomy in flight systems (all the way to complete unmanned operation), and integration of local feedback loops into enterprise-wide scheduling and resource allocation systems. Extending the benefits of control to these non-traditional systems offers enormous opportunities in improved efficiency, productivity, safety, and reliability.

In order to realize the potential of control applied to these emerging applications, new methods and approaches must be developed. Among the challenges currently facing the field, a few examples provide insight into the difficulties ahead:

- *Control of systems with both symbolic and continuous dynamics.* Next generation systems will combine logical operations (such as symbolic reasoning and decision making) with continuous quantities (such as voltages, positions, and concentrations). The current theory is not well-tuned for dealing with such systems, especially as we scale to very large systems.
- *Control in distributed, asynchronous, networked environments.* Control distributed across multiple computational units, interconnected through packet-based communications, will require new formalisms for insuring stability, performance and robustness. This is especially true in applications where one cannot ignore computational and communications constraints in performing control operations.
- *High level coordination and autonomy.* Increasingly, feedback is being designed into enterprise-wide decision systems, including supply chain management and logistics, airspace management and air traffic control, and C4ISR systems. The advances of the last few decades in analysis and design of robust control systems must be extended to these higher level decision-making systems if they are to perform reliably in realistic settings.
- *Automatic synthesis of control algorithms, with integrated validation and verification.* Future engineering systems will require the ability to rapidly design, redesign and implement control software. Researchers need to develop much more powerful design tools that automate the entire control design process from model development to hardware-in-the-loop simulation, including system-level software verification and validation.
- *Building very reliable systems from unreliable parts.* Most large engineering systems must continue to operate even when individual components fail. Increasingly, this requires designs that allow the system to automatically reconfigure itself so that its performance degrades gradually rather than abruptly.

Each of these challenges will require many years of effort by some of the best minds in the field to make the results rigorous, practical, and widely available. They will require investments by funding agencies to ensure that current progress is continued and that forthcoming technologies are realized to their fullest.

Recommendations

To address these challenges and deliver on the promise of the control field, the Panel recommends that the following actions be undertaken:

1. Substantially increase research aimed at the *integration* of control, computer science, communications, and networking. This includes principles, methods and tools for modeling and control of high level, networked, distributed systems, and rigorous techniques for reliable, embedded, real-time software.
2. Substantially increase research in control at higher levels of decision making, moving toward enterprise level systems. This includes work in dynamic resource allocation in the presence of uncertainty, learning and adaptation, and artificial intelligence for dynamic systems.
3. Explore high-risk, long-range applications of control to areas such as nanotechnology, quantum mechanics, electromagnetics, biology, and environmental science. Dual investigator, interdisciplinary funding might be a particularly useful mechanism in this context.
4. Maintain support for theory and interaction with mathematics, broadly interpreted. The strength of the field relies on its close contact with rigorous mathematics, and this will be increasingly important in the future.
5. Invest in new approaches to education and outreach for the dissemination of control concepts and tools to non-traditional audiences. The community must do a better job of educating a broader range of scientists and engineers on the principles of feedback and the use of control to alter the dynamics of systems and manage uncertainty.

The impact of control is one which will come through many applications, in aerospace and transportation, information and networking, robotics and intelligent machines, materials and processing, and biology and medicine. It will enable us to build more complex systems and ensure that the systems we build are reliable, efficient, and robust. The Panel's recommendations are founded on the diverse heritage of rigorous work in control and are key actions to realize the opportunities of control in an information rich world.

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. At its core, control is an *information* science, and includes the use of information in both analog and digital representations.

A modern control system 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 insuring that the dynamics of the closed loop plant 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 sensing, computation, and actuation are clearly seen. In modern control systems, computation is typically implemented on a digital computer, requiring the use of analog-to-digital (A/D) and digital-to-analog (D/A) converters. Uncertainty enters the system through noise in sensing and actuation subsystems,

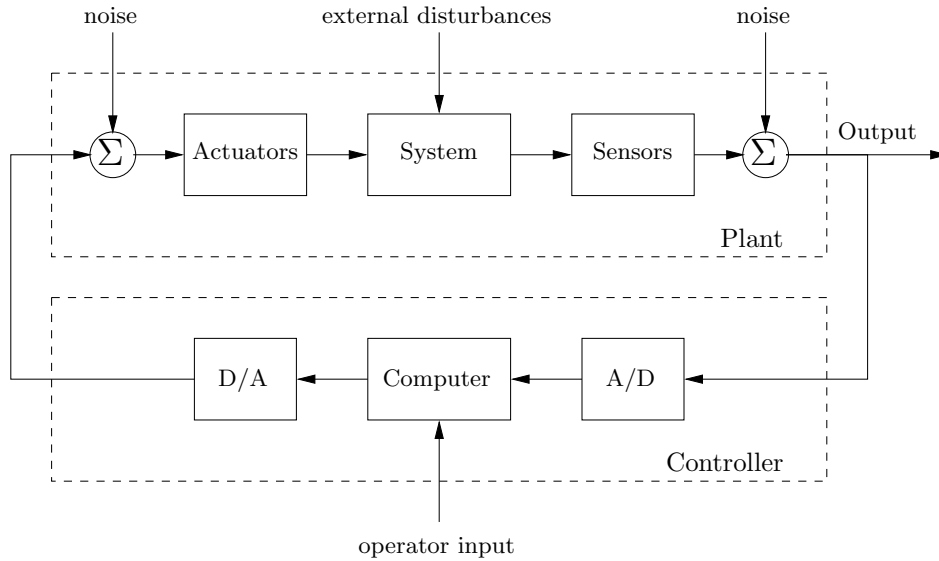


Figure 2.1. *Components of a modern control system.*

external disturbances that affect the underlying system physics, and uncertain dynamics in the physical system (parameter errors, unmodeled effects, etc).

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. 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 concise 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 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 us 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 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).

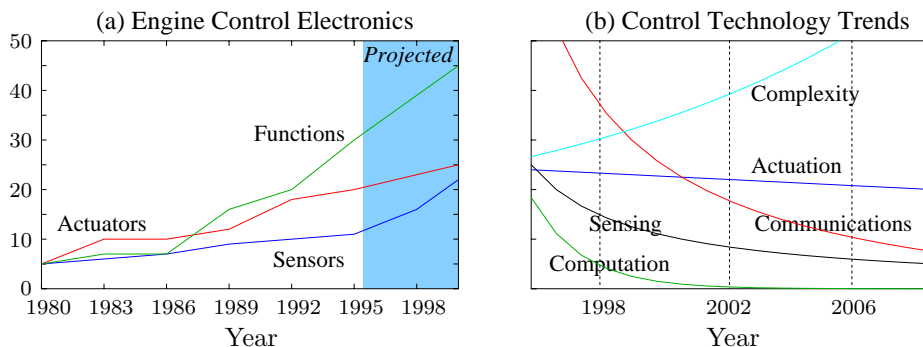


Figure 2.2. Trends in control technology: (a) the number of sensors, actuators and control functions in engine controls [6] and (b) illustration of cost/performance trends for component technologies.

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, computation and communications in automotive applications. As in many other application areas, the number of sensors, actuators, and microprocessors is increasing dramatically, as new features such as anti-lock 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 cheap 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. 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 us to design *robust* interconnections between subsystems, a feature that is crucial in the operation of all large, engineered systems.

Control is also closely associated with computer science, since virtually all modern control algorithms are implemented in software. However, control algorithms and software are very different from traditional computer software. The physics (dynamics) of the system are paramount in analyzing and designing them and their (hard) real-time nature dominates issues of their implementation. From a software-centric perspective, an F-16 is simply another peripheral, while from a




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 flyball governor) [19].

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 it off when temperature is too high. This explanation captures the essence of feedback, but it is a bit too simple

even for a basic device such as the thermostat. Actually, because lags and delays exist in the heating plant and sensor, a good thermostat does a bit of anticipation, turning the plant off before the error actually changes sign. This avoids excessive temperature swings and cycling of the heating plant.

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

There are many other control system examples, of course, that have developed over the years with progressively increasing levels of sophistication and impact. An early system with broad public exposure was the “cruise control” option introduced on automobiles in 1958 (see Figure 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 controls in oil refineries and chemical plants, to IC 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 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




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.*

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 [17] (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 the 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, 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 air 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, Boeing Corporation)

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 trade-off 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.

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 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 us with 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,




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

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 environment 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.




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.*

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 our 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 Now Facing Us

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 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 the 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 led 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 us 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 us to separate 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, we must develop a theory 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 abstraction. 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 we apply control to problems that cannot necessarily be expressed in terms of continuous variables, but rather have 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 abstraction, 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 experts 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 a high respect for careful thinking, often coupled with an emphasis on clear 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 vigorously pursue 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 our 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 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.

Chapter 3

Applications, Opportunities, and Challenges

In this chapter, we consider some of the opportunities and challenges for control in different application areas. These areas are not comprehensive, but represent some of the areas in which control has been historically important as well as some of the emerging areas that will drive control theory, technology and practice forward.

The Panel chose to explore five major application areas to identify the overarching themes that would guide its recommendations. These five areas are:

- Aerospace and transportation
- Information and networks
- Robotics and intelligent machines
- Biology and medicine
- Materials and processing

In addition, several additional areas arose over the course of the Panel's deliberations, including environmental science and engineering, economics and finance, and molecular and quantum systems. Taken together, the areas represent an enormous collection of applications and demonstrate the breadth of applicability of ideas from control.

The opportunities and challenges in each of these application areas form the basis for the major recommendations in this report. In each area, we have sought the advice and insights not only of control researchers in the area, but also experts in the application domains who might not consider themselves to be control researchers. In this way, we hoped to identify the true challenges in each area, rather than simply identifying interesting control problems that may not have a substantial opportunity for impact. We hope that the findings in these areas will be of interest not only to control researchers, but also to scientists and engineers seeking to understand how control tools might be applied to their discipline.

There were several overarching themes that arose across all of the areas considered by the Panel. The use of systematic and rigorous tools is considered critical

to future success and is an important trademark of the field. At the same time, the next generation of problems will require a paradigm shift in control research and education. The increased information available across all application areas requires more integration with ideas from computer science and communications, as well as improved tools for modeling, analysis, and synthesis for complex decision systems that contain a mixture of symbolic and continuous dynamics. The need to continue research in the theoretical foundations that will underly future advances was also common across all of the applications.

In each section that follows we attempt to give a brief description of the background and history of control in that domain, followed by a selected set of topics which are used to explore the future potential for control and the technical challenges that must be addressed. As in the rest of the report, we do not attempt to be comprehensive in our choice of topics, but rather highlight some of the areas where we see the greatest potential for impact. Throughout these sections, we have limited the references to those that provide historical context, future directions, or broad overviews in the topic area, rather than specific technical contributions (which are too numerous to properly document).

3.1 Aerospace and Transportation

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

Wilbur Wright, lecturing to the Western Society of Engineers in 1901 [25].

Aerospace and transportation encompasses a collection of critically important application areas where control is a key enabling technology. These application areas represent a very large part of the modern world's overall technological capability. They are also a major part of its economic strength, and they contribute greatly to the well being of its people. The historical role of control in these application areas, the current challenges in these areas, and the projected future needs all strongly support the recommendations of this report.

The Historical Role

In aerospace, specifically, control has been a key technological capability tracing back to the very beginning of the 20th Century. Indeed, the Wright Brothers are correctly famous not simply for demonstrating powered flight—they actually demonstrated *controlled* powered flight. Their early Wright Flyer incorporated moving control surfaces (vertical fins and canards) and warpable wings that allowed the pilot to regulate the aircraft's flight. In fact, the aircraft itself was not stable, so continuous pilot corrections were mandatory. This early example of controlled flight is followed by a fascinating success story of continuous improvements in flight control technology, culminating in the very high performance, highly reliable automatic flight control systems we see on modern commercial and military aircraft today (see Flight Control Vignette, page 16).

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

Similarly, control technology has played a key role in the continuing improvement and evolution of transportation—in our cars, highways, trains, ships and air transportation systems. Control's contribution to the dramatic increases of safety, reliability and fuel economy of the automobile is particularly noteworthy. Cars have advanced from manually tuned mechanical/pneumatic technology to computer controlled operation of all major functions including fuel injection, emission control, cruise control, braking, cabin comfort, etc. Indeed, modern automobiles carry

dozens of individual processors to see to it that these functions are performed accurately and reliably over long periods of time and in very tough environments. A historical perspective of these advances in automotive applications is provided in the following vignette.

Vignette: Emissions Requirements and Electronic Controls for Automotive Systems (Mark Barron and William Powers, Ford Motor Co)

One of the major success stories for electronic controls is the development of sophisticated engine controls for reducing emissions and improving efficiency. Mark Barron and Bill Powers described some of these advances in an article written in 1996 for the inaugural issue of the *IEEE/ASME Transactions on Mechatronics* [6].

In their article, Barron and Powers describe the environment that led up to the introduction of electronic controls in automobile engines:

Except for manufacturing technology, the automobile was relatively benign with respect to technology until the late 1960s. Then two crises hit the automotive industry. The first was the environmental crisis. The environmental problems led to regulations which required a reduction in automotive emissions by roughly an order of magnitude. The second crisis was the oil embargo in the early 1970s which created fuel shortages, and which led to legislation in the U.S. requiring a doubling of fuel economy. ...

Requirements for improved fuel efficiency and lower emissions demanded that new approaches for controlling the engine be investigated. While today we take for granted the capabilities which have been made possible by the microprocessor, one must remember that the microprocessor wasn't invented until the early 1970s. When the first prototype of a computerized engine control system was developed in 1970, it utilized a minicomputer that filled the trunk of a car. But then the microprocessor was invented in 1971, and by 1975 engine control had been reduced to the size of a battery and by 1977 to the size of a cigar box.

These advances in hardware allowed sophisticated control laws that could deal with the complexities of maintaining low emissions and high fuel economy:

The introduction in the late 1970s of the platinum catalytic converter was instrumental in reducing emissions to meet regulations. The catalytic converter is an impressive passive device which operates very effectively under certain conditions. One of the duties of the engine control system is to maintain those conditions by patterning the exhaust gases such that there are neither too many hydrocarbons nor too much oxygen entering the catalyst. If the ratio of air to fuel entering the engine is kept within a very tight range (i.e., a few percent) the catalyst can be over 90% efficient in removing hydrocarbons, carbon monoxide, and oxides of nitrogen. However, the

catalyst isn't effective until it has reached a stable operating temperature greater than 600°F (315°C), and a rule of thumb is that 80% of emissions which are generated under federal test procedures occur during the first two minutes of operation while the catalyst is warming to its peak efficiency operating temperature. On the other hand if the catalyst is operated for an extended period of time much above 1000°F (540°C) it will be destroyed. Excess fuel can be used to cool the catalyst, but the penalty is that fuel economy gets penalized. So the mechatronic system must not only control air-fuel ratios so as to maintain the catalyst at its optimum operating point, it must control the engine exhaust so that there is rapid lightoff of the catalyst without overheating, while simultaneously maintaining maximum fuel efficiency.

The success of control in meeting these challenges is evident in the reduction of emissions that has been achieved over the last 30 years [30]:

US, European and Japanese Emission Standard continue to require significant reductions in vehicle emissions. Looking closely at US passenger car emission standards, the 2005 level of hydrocarbon (HC) emissions is less than 2% of the 1970 allowance. By 2005, carbon monoxide (CO) will be only 10% of the 1970 level, while the permitted level for oxides of nitrogen will be down to 7% of the 1970 level.

Furthermore, the experience gained in engine control provided a path for using electronic controls in many other applications [6]:

Once the industry developed confidence in on-board computer control, other applications rapidly followed. Antilock brake systems, computer controlled suspension, steering systems and air bag passive restraint systems are examples. The customer can see or feel these systems, or at least discern that they are on the vehicle, whereas the engine control system is not an application which is easily discernible by the customer. Computers are now being embedded in every major function of the vehicle, and we are seeing combinations of two or more of these control systems to provide new functions. An example is the blending of the engine and antilock brake system to provide a traction control system, which controls performance of the vehicle during acceleration whereas antilock brakes control performance of the vehicle during deceleration.

Another important consequence of the use of control in automobiles was its success in demonstrating that control provided safe and reliable operation. The cruise control option introduced in the 1950s and 1960s was one of the first servo systems receiving very broad public exposure. Our society's inherent trust in control technology traces back to the success of such early control systems.

Certainly, each of these successes owes its debt to improvements in many



Figure 3.1. *The F/A-18 aircraft, one of the first production military fighters to use “fly-by-wire” technology, and the X-45 (UCAV) unmanned aerial vehicle.*

technologies, e.g. propulsion, materials, electronics, computers, sensors, navigation instruments, etc. However, they also depend in no small part on the continuous improvements that have occurred over the century in the theory, analysis methods and design tools of control. As an example, “old timers” in the flight control engineering community still tell the story that early control systems (circa World War II) were designed by manually tuning feedback gains in flight—in essence, trial-and-error design performed on the actual aircraft. Dynamic modeling methods for aircraft were in their infancy at that time, and formal frequency-domain design theories to stabilize and shape single-input single-output feedback loops were still only subjects of academic study. Their incorporation into engineering practice revolutionized the field, enabling successful feedback systems designed for ever more complex applications, consistently, with minimal trial-and-error, and with reasonable total engineering effort.

Of course, the formal modeling, analysis and control system design methods described above have advanced dramatically since mid-century. As a result of significant R&D activities over the last fifty years, the state of the art today lets us design controllers for much more than single-input single-output systems. The theory and tools handle many inputs, many outputs, complex uncertain dynamic behavior, difficult disturbance environments, and ambitious performance goals. In modern aircraft and transportation vehicles, dozens of feedback loops are not uncommon, and in process control the number of loops reaches well into the hundreds. Our ability to design and operate such systems consistently, reliably, and cost effectively rests in large part on the investments and accomplishments of control over the latter half of the century.

Current Challenges and Future Needs

Still, the control needs of some engineered systems today and those of many in the future outstrip the power of current tools and theories. This is so because our current tools and theories apply most directly to problems whose dynamic behaviors are smooth and continuous, governed by underlying laws of physics and represented

mathematically by (usually large) systems of differential equations. Most of the generality and the rigorously provable features of our methods can be traced to this nature of the underlying dynamics.

Many new control design problems no longer satisfy these underlying characteristics, at least in part. Design problems have grown from so-called “inner loops” in a control hierarchy (e.g. regulating a specified flight parameter) to various “outer loop” functions which provide logical regulation of operating modes, vehicle configurations, payload configurations, health status, etc [3]. For aircraft, these functions are collectively called “vehicle management.” They have historically been performed by pilots or other human operators and have thus fallen on the other side of the man-machine boundary between humans and automation. Today, that boundary is moving!

There are compelling reasons for the boundary to move. They include economics (two, one or no crew members in the cockpit versus three), safety (no operators exposed to dangerous or hostile environments), and performance (no operator-imposed maneuver limits). A current example of these factors in action is the growing trend in all branches of the military services to field unmanned vehicles. Certain benign uses of such vehicles are already commonplace (e.g. reconnaissance and surveillance), while other more lethal ones are in serious development (e.g. combat UAVs for suppression of enemy air defenses). Control design efforts for such applications must necessarily tackle the entire problem, including the traditional inner loops, the vehicle management functions, and even the higher-level “mission management” functions coordinating groups of vehicles intent on satisfying specified mission objectives.

Today’s engineering methods for designing the upper layers of this hierarchy are far from formal and systematic. In essence, they consist of collecting long lists of logical if-then-else rules from experts, programming these rules, and simulating their execution in operating environments. Because the logical rules provide no inherent smoothness (any state transition is possible) only simulation can be used for evaluation and only exhaustive simulation can guarantee good design properties. Clearly, this is an unacceptable circumstance—one where the strong system-theoretic background and the tradition of rigor held by the control community can make substantial contributions.

One can speculate about the forms that improved theories and tools for non-smooth (hybrid) dynamical systems might take. For example, it may be possible to impose formal restrictions on permitted logical operations, to play a regularizing role comparable to laws of physics. If rigorously obeyed, these restrictions could make resulting systems amenable to formal analyses and proofs of desired properties. This approach is similar to computer language design, and provides support for one of the recommendations of this report, namely that the control and computer science disciplines need to grow their intimate interactions. It is also likely that our traditional standards of formal rigor must expand to firmly embrace computation, algorithmic solutions, and heuristics.

However, we must not ever lose sight of the key distinguishing features of the control discipline, including the need for hard real time execution of control laws and the need for ultra-reliable operation of all hardware and software control compo-




Figure 3.2. *Battle space management.*

nents. Many controlled systems today (auto-land systems of commercial transports, launch boosters, F-16 and B-2 aircraft, certain power plants, certain chemical process plants, etc.) fail catastrophically in the event of control hardware failures, and many future systems, including the unmanned vehicles mentioned above, share this property. But the future of aerospace and transportation holds still more complex challenges. We noted above that changes in the underlying dynamics of our control design problems from continuous to hybrid are well under way. An even more dramatic trend on the horizon is a change in dynamics to large collections of distributed entities with local computation, global communication connections, very little regularity imposed by laws of physics, and no possibility to impose centralized control actions. Examples of this trend include the national airspace management problem, automated highway and traffic management, and the command and control for future battlefields (Figure 3.2).

The national airspace problem is particularly significant today, with eventual gridlock and congestion threatening the integrity of the existing air transportation system. Even with today's traffic, ground holds and airborne delays in flights due to congestion in the skies have become so common that airlines automatically pad their flight times with built-in delays. The structure of the Air Traffic Control (ATC) system is partially blamed for these delays: the control is distributed from airspace region to airspace region, yet within a region the control is almost wholly centralized, with sensory information from aircraft sent to a human air traffic controller who uses ground-based navigation and surveillance equipment to manually route aircraft along sets of well-traveled routes. In today's system, bad weather, aircraft failure, and runway or airport closure have repercussions throughout the whole country. Efforts are now being made to improve the current system by developing cockpit "sensors" such as augmented GPS navigation systems and datalinks for aircraft

to aircraft communication. Along with these new technologies, new hierarchical control methodologies are being proposed, which automate some of the functionality of ATC. This opens up a set of new challenges: the design of information-sharing mechanisms and new, distributed, *verified* embedded control schemes for separation assurance between aircraft, and the design of dynamic air traffic network topologies which aid in the safe routing of aircraft from origin to destination and which adapt to different traffic flows, are two areas which provide a tremendous opportunity to researchers in the control community.

Finally, it is important to observe that the future also holds many applications that fall under our traditional control design paradigm, yet are worthy of research support because of their great impact. Conventional “inner loops” in automobiles, but for non-conventional power plants, are examples. Hybrid cars combining electrical drives and low-power internal combustion engines and fuel cell powered cars combining electrical drives with fuel cell generation both depend heavily of well-designed control systems to operate efficiently and reliably. Similarly, increased automation of traditional transportation systems such as ships and railroad cars, with added instrumentation and cargo-tracking systems will rely on advanced controls and schedule optimization to achieve maximum economic impact. Another conventional area is general aviation, where control systems to make small aircraft easy and safe to fly and increased automation to manage them are essential needs.

Other Trends in Aerospace and Transportation

In addition to the specific areas highlighted above, there are many other trends in aerospace and transportation that will benefit from and inform new results in control. We briefly describe a few of these here.

Automotive Systems With 60 million vehicles produced each year, automotive systems are a major application area for control. Emission control regulations passed in the 1970s created a need for more sophisticated engine control systems that could provide clean and efficient operation in a variety of operating environments and over the lifetime of the car. The development of the microprocessor at that same time allowed the implementation of sophisticated algorithms that have reduced the emissions in automobiles by as much as a factor of 50 from their 1970 levels.

Future automobile designs will rely even more heavily on electronic controls [30]. Figure 3.3 shows some of the components that are being considered for next generation vehicles. Many of these components will build on the use of control techniques, including radar-based speed and spacing control systems, chassis control technologies for stability enhancement and improved suspension characteristics, active control of suspension and braking, and active restraint systems for safety. In addition, more sophisticated use of networking and communications devices will allow enhanced energy management between components and vehicle diagnostics with owner/dealer notification.

These new features will require highly integrated control systems that combine multiple components to provide overall stability and performance. Systems such as chassis control will require combining steering, braking, powertrain and suspension




Figure 3.3. Major future components for basic vehicle functions [30].

subsystems, along with adding new sensors. One can also imagine increased interaction between vehicles and the roadway infrastructure, as automated highways and self-controlled vehicles move from the research lab into applications. These latter applications are particularly challenging since they begin to link heterogeneous vehicles through communications systems that will experience varying bandwidths and latency depending on the local environment. Providing safe, reliable, and comfortable operation for such systems is a major challenge for control and one that will have application in a variety of consumer, industrial, and military applications.

Aircraft Propulsion Systems Much more effective use of information in propulsion systems is possible as the price/performance ratio of computation and sensing continues to drop. Intelligent turbine engines will ultimately lower lifetime operating and maintenance costs, similar to current and upcoming automotive systems. They will provide advanced health, performance, and life management by embedding models of their operation and optimizing based on condition and mission. They will be more flexible and more tolerant of component faults, and will integrate into the owners asset management system, lowering maintenance and fleet management costs by making engine condition information available to the owner on demand

and ensuring predictable asset availability.

Detection of damage (diagnostics) and prediction of the implications (prognostics) are the heart of an intelligent engine. Detailed modeling of the thermofluid, structural, and mechanical systems, as well as the operational environment, is needed for such assessments. To allow on-product use accounting for system interactions, physics-based models will be constructed using advanced techniques in reduced-order modeling. This approach significantly extends recent engine component modeling.

Embedded models can also be used for on-line optimization and control in real time. The benefit is the ability to customize engine performance to changes in operating conditions and the engine's environment through updates in the cost function, onboard model, and constraint set. Many of the challenges of designing controllers that are robust to a large set of uncertainties can thus be embedded in the online optimization, and robustness through a compromise design is replaced by always-optimal performance.

Flow Control Flow control, whether active or passive, involves the use of reactive devices for modifying fluid flow for the purposes of enhanced operability. Sample applications for flow control include increased lift and reduced drag on aircraft wings, engine nacelles, compressor fan blades, and helicopter rotor blades; higher performance diffusers in gas turbines, industrial heaters and chillers, and engine inlets; wake management for reduction of resonant stress and blade vortex interaction; and enhanced mixing for combustion and noise applications. A number of devices have been explored in the past several years for actuation of flow fields. These range from novel air injection mechanisms for control of rotating stall and separation, to synthetic jets developed for mixing enhancement and vectoring, to MEMS devices for modulating boundary layers and flow around stagnation points. In addition, new sensing technology, such as micro anemometers, is also becoming available.

These changes in sensing and actuation technology are enabling new applications of control to unstable shear layers and separated flow, thermoacoustic instabilities, and compression system instabilities such as rotating stall and surge (see [8] for a recent survey). An emerging area of interest in hypersonic flight systems, where flow control techniques could provide a larger toolbox for design of vehicles, including drag reduction, novel methods for producing control forces, and better understanding of the complex physical phenomena at these speeds.

Space Systems¹ The exploitation of space systems for civil, commercial, defense, scientific, or intelligence purposes gives rise to a unique set of challenges in the area of control. For example, most space missions cannot be adequately tested on the ground prior to flight, which has a direct impact on many dynamics and control problems. A three-pronged approach is required to address these challenging space system problems: (1) detailed modeling, including improved means of

¹The Panel would like to thank Jonathan How and Jesse Leitner for their contributions to this section.

characterizing, at a very small scale, the fundamental physics of the systems; (2) flight demonstrations to characterize the behavior of representative systems; and (3) design of navigation and control approaches that maintain performance (noise rejection and tracking) even with uncertainties, failures, and changing dynamics.

There are two significant areas that can revolutionize the achievable performance from future space missions: flexible structure analysis and control, and space vehicle formation flying. These both impact the allowable size of the effective aperture, which influences the “imaging” performance, whether it is optical imaging or the collection of signals from a wide range of wavelengths. There are fundamental limitations that prevent further developments with monolithic mirrors (with the possible exception of inflatable and foldable membranes, which introduce their own extreme challenges) and the various segmented approaches—deployed arrays, tethered or freeflyer formations—provide the only solution. However, these approaches introduce challenging problems in characterizing the realistic dynamics and developing sensing and control schemes to maintain the necessary optical tolerances.

A significant amount of work has been performed in the area of flexible structure dynamics and control under the auspices of the Strategic Defense Initiative Organization (SDIO) in the 1970s and 80s. However, at the performance levels required for future missions (nanometers), much research remains to develop models at the micro-dynamics level and control techniques that can adapt to system changes at these small scales.

Similar problems exist with formation control for proposed imaging interferometry missions. These will require the development of control algorithms, actuators, and computation and communication networks. Sensors will also have to be developed to measure deflections on the scale of nanometers over distances hundreds of meters through kilometers. Likewise, actuation systems of various types must be developed that can control on the scale of nanometers to microns with very low noise levels and fine resolution. The biases and residuals generally accepted due to particular approximations in navigation and control algorithms will no longer be acceptable. Furthermore, the simulation techniques used for verification must, in some cases, maintain precision through tens of orders of magnitude separation in key states and parameters, over both long and short time-scales, and with stochastic noise inputs. In summary, in order to enable the next generations of advanced space systems, we must address the micro- and nano-scale problems in analysis, sensing, control, and simulation, for individual elements and integrated systems.

3.2 Information and Networks

A typical congested gateway looks like a fire hose connected to a soda straw through a small funnel (the output queue). If, on average, packets arrive faster than they can leave, the funnel will fill up and eventually overflow. RED [Random Early Detection] is [a] simple regulator that monitors the level in the funnel and uses it to match the input rate to the output (by dropping excess traffic). As long as its control law is monotone non-decreasing and covers the full range of 0 to 100% drop rate, RED works for any link, any bandwidth, any type of traffic.

Van Jacobson, North American Network Operators' Group meeting, 1998 [18].

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

Control of Networks

Control of networks is a large area, spanning many topics, a few of which are briefly described here. The basic problems in control of networks include routing the flow of packets through the network,

Several features of these control problems make them very challenging. The dominant feature is the extremely large scale of the system; the Internet is probably the largest feedback control system man has ever built. Another is the decentralized nature of the control problem: local decisions must be made quickly, and based only on local information. Stability is complicated by the presence of varying time lags, as information about the network state can only be observed or relayed to controllers after a time delay, and the effect of a local control action can be felt throughout the network after substantial delay. Uncertainty and variation in the network, through network topology, transmission channel characteristics, traffic demand, available resources, etc. may change constantly and unpredictably. Another complicating issues is the diverse traffic characteristics, in terms of arrival statistics at both the packet and flow time scales, and different requirements for quality of service, in terms of delay, bandwidth, and loss probability, that the network must support.

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

To illustrate these characteristics, we briefly describe the control mechanisms that can be invoked in serving a file request from a client: network caching, congestion control, routing and power control. Figure 3.4 shows a typical map for the networking infrastructure that is used to process such a request.

The problem of optimal network caching is to copy documents (or services) that are likely to be accessed often, from many different locations, on multiple servers. When the document is requested, it is returned by the nearest server. Here, proximity may be measured by geographical distance, hop count, network congestion, server load or a combination. The goal is to reduce delay, relieve server

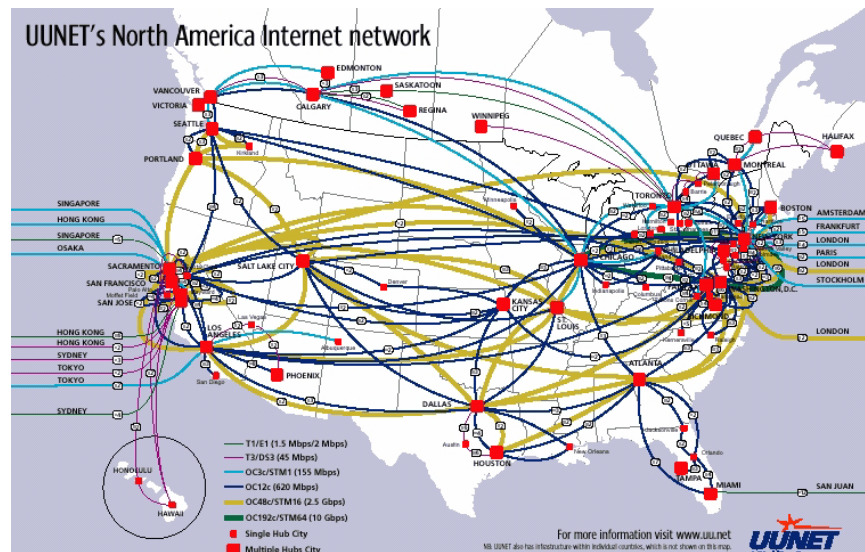


Figure 3.4. *UUNET network backbone for North America. Figure courtesy WorldCom.*

load, balance network traffic, and improve service reliability. If changes are made to the source document, those changes (at a minimum) must be transmitted to the servers, which consume network bandwidth.

The control problem is to devise a decentralized scheme for how often to update, where to cache copies of documents, and to which server a client request is directed, based on estimation and prediction of access patterns, network congestion, and server load. Clearly, current decisions affect the future state, such as future traffic on links, future buffer levels, delay and congestion, and server load. Thus a web of caches is a decentralized feedback system that is spatially distributed and interconnected, where control decisions are made asynchronously based on local and delayed information.

When a large file is requested, the server that is selected to return the file breaks it into a stream of packets and transports them to the client in a rate-adaptive manner. This process is governed by the Transport Control Protocol (TCP). The client acknowledges successful reception of each packet and the stream of acknowledgment carries congestion information to the server. Congestion control is a distributed algorithm to share network resources among competing servers. It consists of two components: a source algorithm that dynamically adjusts the server rate in response to congestion in its path, and a router algorithm that updates a congestion measure and sends it back to sources that go through that router. Examples of congestion measures are loss probability and queuing delay. They are implicitly updated at the routers and implicitly fed back to sources through delayed end-to-end observations of packet loss or delay. The equilibrium and dynamics of the network depends on the pair of source and router algorithms.

A good way to understand the system behavior is to regard the source rates as primal variables and router congestion measures as dual variables, and the process of congestion control as an asynchronous distributed primal-dual algorithm carried out by sources and routers over the Internet in real time to maximize aggregate source utility subject to resource capacity constraints. Different protocols all solve the same prototypical problem, but they use different utility functions and implement different iterative rules to optimize them. Given any source algorithm, it is possible to derive explicitly the utility function it is implicitly optimizing.

While TCP controls the rate of a packet flow, the path through the network is controlled by the Internet Protocol (IP). In its simplest form, each router must decide which output link a given packet will be sent to on its way to its final destination. Uncertainties include varying link congestion, delays, and rates, and even varying network topology (e.g., a link goes down, or new nodes or links become available), as well as future traffic levels. A routing algorithm is an asynchronous distributed algorithm executed at routers that adapts to node and link failures, balances network traffic and reduces congestion. It can be decomposed into several time scales, with very fast decisions made in hardware using lookup tables, which in turn are updated on a slower time scale. At the other extreme in time scale from the routing problem, we have optimal network planning, in which new links and nodes are proposed to meet predicted future traffic demand.

The routing problem is further exacerbated in wireless networks. Nodes with wireless modems may be mobile, and the address of a node may neither indicate where it is located nor how to reach it. Thus the network needs to either search for a node on demand, or it must keep track of the changing locations of nodes. Further, since link capacities in wireless networks may be scarce, routing may have to be determined in conjunction with some form of load balancing. This gives rise to the need for distributed asynchronous algorithms which are adaptive to node locations, link failures, mobility, and changes in traffic flow requirements.

Finally, if the client requesting the file accesses it through an *ad hoc* wireless network, then there also arises the problem of power control: at what transmission power level should each packet broadcast be made? Power control is required because *ad hoc* networks do not come with ready made links; the topology of the network is formed by individual nodes choosing the power levels of their broadcasts. This poses a conceptual problem in the current protocol hierarchy of the Internet since it simultaneously affects the physical layer due to its effect on signal quality, the network layer since power levels determine which links are available for traffic to be routed, and the transport layer since power levels of broadcasts affect congestion. Power control is also a good challenge for multi-objective control since there are many cost criteria involved, such as increasing the traffic carrying capacity of the network, reducing the battery power used in relaying traffic, and reducing the contention for the common shared medium by the nodes in geographical vicinity.

Control of networks extends beyond data and communication networks. Optimal routing and flow control of commercial aircraft (with emphasis on guaranteeing safe inter-vehicle distances) will help maximize utilization of airports. The (network and software) infrastructure for supply chain systems is being built right now, and simple automated supply chain management systems are beginning to be deployed.

In the near future, sophisticated optimization and control methods can be used to direct the flow of goods and money between suppliers, assemblers and processors, and customers.

Control over Networks

While the advances in information technology to date have led to a global Internet that allows users to exchange information, it is clear that the next phase will involve much more interaction with the physical environment. Networks of sensory or actuator nodes with computational capabilities, connected wirelessly or by wires, can form an orchestra which controls our physical environment. Examples include automobiles, smart homes, large manufacturing systems, intelligent highways and networked city services, and enterprise-wide supply and logistics chains. Thus, this next phase of the information technology revolution is the convergence of communication, computing and control. The following vignette describes a major architectural challenge in achieving this convergence.

Vignette: The importance of abstractions and architecture for the convergence of communications, computing, and control (P. R. Kumar, Univ. of Illinois, Urbana-Champaign)

Communication networks are very diverse, running over copper, radio, or optical links, various computers, routers, etc. However, they have an underlying architecture which allows one to just plug-and-play, and not concern oneself with what lies underneath. In fact, one reason for the anarchic proliferation of the Internet is precisely this architecture—a hierarchy of layers together with peer-to-peer protocols connecting the layers at different nodes. On one hand, nodes can be connected to the Internet without regard to the physical nature of the communication link, whether it be infrared or copper, and this is one reason for the tremendous growth in the number of nodes on the Internet. On the other hand, the architecture allows plug-and-play at all levels, and thus each layer can be designed separately, allowing a protocol at one level to be modified over time without simultaneously necessitating a redesign of the whole system. This has permitted the Internet protocols to evolve and change over time.

This raises the issue: What is the right architecture for the convergence of communication, control, and computing? Is there an architecture which is application and context independent, one which allows proliferation, just as the OSI architecture did for communication networks? What are the right abstraction layers? How does one integrate information, control, and computation? If the overall design allows us to separate algorithms from architecture, then this convergence of control with communication and computation will rapidly proliferate.

As existing networks continue to build out, and network technology becomes cheaper and more reliable than fixed point-to-point connections, even in small localized systems, more and more control systems will operate over networks. We can foresee sensor, actuator, diagnostic, and command and coordination signals all

traveling over data networks. The estimation and control functions can be distributed across multiple processors, also linked by data networks. (For example, smart sensors can perform substantial local signal processing before forwarding relevant information over a network.)

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

Can we develop a theory and practice for control systems that operate in a distributed, asynchronous, packet-based environment?

It is very interesting to compare current control system technology with current packet-based data networks. Data networks are extremely robust to gross, unpredicted changes in topology (such as loss of a node or a link); packets are simply re-sent or re-routed to their destination. Data networks are self-configuring: we can add new nodes and links, and soon enough packets are flowing through them. One of the amazing attributes of data networks is that, with good architecture and protocol design, they can be far more reliable than their components. This is in sharp contrast with modern control systems, which are only as reliable as their weakest link. Robustness to component failure must be designed in, by hand (and is, for safety critical systems).

Looking forward, we can imagine a marriage of current control systems and networks. The goal is an architecture, and design and analysis methods, for distributed control systems that operate in a packet-based network. If this is done correctly, we might be able to combine the good qualities of a robust control system, i.e., high performance and robustness to parameter variation and model mismatch, with the good qualities of a network: self-configuring, robust to gross topology changes and component failures, and reliability exceeding that of its components.

One can imagine systems where sensors asynchronously burst packets onto the network, control processors process the data and send it out to actuators. Packets can be delayed by varying amounts of time, or even lost. Communication links can go down, or become congested. Sensors and actuators themselves become unavailable or available. New sensors, actuators, and processors can be added to the system, which automatically reconfigures itself to make use of the new resources. As long as there are enough sensors and actuators available, and enough of the packets are getting through, the whole system works (although we imagine not as well as with a dedicated, synchronous control system). This is of course very different from any existing current high performance control system.

It is clear that for some applications, current control methods, based on syn-

chronous clocked systems and networks that guarantee arrival and bounded delays for all communications, are the best choice. There is no reason not to configure the controller for a jet engine as it is now, i.e., a synchronous system with guaranteed links between sensors, processors, and actuators. But for consumer applications not requiring the absolute highest performance, the added robustness and self-reconfiguring abilities of a packet-based control system could make up for the lost performance. In any case what will emerge will probably be something in between the two extremes, of a totally synchronous system and a totally asynchronous packet-based system.

Clearly, several fundamental control concepts will not make the transition to an asynchronous, packet-based environment. The most obvious casualty will be the transfer function, and all the other concepts associated with linear time invariant (LTI) systems (impulse and step response, frequency response, spectrum, bandwidth, etc.) This is not a small loss, as this has been a foundation of control engineering since about 1930. With the loss goes a lot of intuition and understanding. For example, Bode plots were introduced in the 1930s to understand and design feedback amplifiers, were updated to handle discrete-time control systems in the 1960s, and were applied to robust MIMO control systems in the 1980s (via singular value plots). Even the optimal control methods in the 1960s, which appeared at first to be quite removed from frequency domain concepts, were shown to be nicely interpreted via transfer functions.

So what methods will make the transition? Many of the methods related to optimal control, and optimal dynamic resource allocation will likely transpose gracefully to an asynchronous, packet-based environment. A related concept that is likely to survive is also one of the oldest: Lyapunov functions (which were introduced in 1890). The following vignette describes some of the possible changes to control that may be required.

Vignette: Lyapunov Functions in Networked Environments (Stephen Boyd, Stanford)

Here is an example of how an “old” concept from control will update gracefully. The idea is that of the Bellman value function, which gives the optimal value of some control problem, posed as an optimization problem, as a function of the starting state. It was studied by Pontryagin, Bellman, and other pioneers of optimal control in the 1940s, and has recently had a resurgence (in generalized form) under the name of control Lyapunov function. It is a key concept in dynamic programming.

The basic idea of a control Lyapunov function (or the Bellman value function) is this: If you knew the function, then the best thing to do is to choose current actions that minimize the value function in the current step, without any regard for future effects. (In other words, we ignore the dynamics of the system.) By doing this we are actually carrying out an optimal control for the problem. In other words, the value function is the cost function whose greedy minimization actually yields the optimal control for the original problem, taking the system dynamics into account. In the work of the 1950s

and 60s, the value function is just a mathematical stepping stone toward the solution of optimal control problems.

But the idea of value function transposes to an asynchronous system very nicely. If the value function, or some approximation, were broadcast to the actuators, then each actuator could take independent and separate action, i.e., each would do whatever it could to decrease the value function. If the actuator were unavailable, then it would do nothing. In general the actions of multiple actuators has to be carefully coordinated; simple examples show that turning on two feedback systems, each with its own sensor and actuator, simultaneously, can lead to disastrous loss of performance, or even instability. But if there is a value or control Lyapunov function that each is separately minimizing, everything is fine; the actions are automatically coordinated (via the value function).

Another idea that will gracefully extend to asynchronous packet-based control is model predictive control. The basic idea is to carry out far more computation at run time, by solving optimization problems in the real-time feedback control law. Model predictive control has played a major role in process control, and also in supply-chain management, but not (yet) in other areas, mainly owing to the very large computational burden it places on the controller implementation. The idea is very simple: at each time step we formulate the optimal control problem, up to some time horizon in the future, and solve for the whole optimal trajectory (say, using quadratic programming). We then use the current optimal input as the actuator signal. We use the sensor signals to update the model, and carry the same process out again. A major extension required to apply model predictive control in networked environments would be the distributed solution of the underlying optimization problem.

Other Trends in Information and Networks

While we have concentrated in this section on the role of control in communications and networking, there are many problems in the broader field of information science and technology for which control ideas will be important.

Vigilant, high confidence software systems Modern information systems are required to operate in environments where the users place high confidence on the availability and correctness of the software programs. This is increasingly difficult due to the networked and often adversarial environment in which these programs operate. One approach that is being explored by the computer science community is to provide confidence through *vigilance*. Vigilance refers to continuous, pervasive, multi-faceted monitoring and correction of system behavior, i.e., control.

The key idea in vigilant software is to use fast and accurate sensing to monitor the execution of a system or algorithm, compare the performance of the algorithm to an embedded model of the computation, and then modify the operation of the algorithm (through adjustable parameters) to maintain the desired performance. This “sense-compute-act” loop is the basic paradigm of feedback control and pro-




Figure 3.5. *An example of a vigilant high confidence software system: distributed sorting using feedback.*

vides a mechanism for online management of uncertainty. Its power lies in the fact that rather than considering every possible situation at design time, we react to specific situations as they occur. An essential element of the strategy is the use of either an embedded model, through which an appropriate control action can be determined, or a predefined control strategy that is analyzed offline to ensure stability, performance, and robustness.

As an indication of how vigilance might be used to achieve high confidence, consider an example of feedback control for distributed sorting, as shown in Figure 3.5. We envision a situation in which we have a collection of partial sort algorithms that are interconnected in a feedback structure. Suppose that each sorter has multiple inputs, from which it chooses the best sorted, and a single output, to which it sends an updated list that is more ordered. By connecting these modules together in a feedback loop, it is possible to get a completely sorted list at the end of a finite number of time steps.

While unconventional from a traditional computer science perspective, this approach gives robustness to failure of individual sorters, as well as self-tuning operation. Robustness comes because if an individual module *unsorts* its data, this data will not be selected from the input streams by the other modules. Further, if the modules have different loads (perhaps due to other processing being done on a given processor), the module with the most time available will automatically take on the load in performing the distributed sorting. Other properties such as disturbance rejection, performance, and stability could also be studied by using tools from control.

Verification and validation of protocols and software The development of complex software systems is increasing at a rapid rate and our ability to design such systems so that they give provably correct performance is increasingly strained. Current methods for verification and validation of software systems require large amounts of testing and many errors are not discovered until late stages of development or even product release. Formal methods for verification of software are used for systems of moderate complexity, but do not scale well to large software systems.

Control theory has developed a variety of techniques for giving provably correct behavior by using upper and lower bounds to effectively break computational complexity bounds. Recent results in convex optimization of semialgebraic problems

(those that can be expressed by polynomial equalities and inequalities) are providing new insights into verification of a diverse set of continuous and combinatorial optimization problems [29]. In particular, these new techniques allow a systematic search for “simple proofs” of continuous and discrete problems and offer ideas for combining formal methods in computer science with stability and robustness results in control.

Real-time supply chain management As increasing numbers of enterprise systems are connected to each other across networks, there is an enhanced ability to perform enterprise level, dynamic reconfiguration of high availability assets for achieving efficient, reliable, predictable operations. As an example of the type of application that one can imagine, consider the operation of a network of HVAC systems for a regional collection of buildings, under the control of a single operating company. In order to minimize overall energy costs for its operation, the company makes a long-term arrangement with an energy broker to supply a specified amount of electrical power that will be used to heat and cool the buildings. In order to get the best price for the energy it purchases, the company agrees to purchase a fixed amount of energy across its regional collection of buildings and to pay a premium for energy usage above this amount. This gives the energy broker a fixed income as well as a fixed (maximum) demand, for which it is willing to sell electricity at a lower price (due to less uncertainty in future revenue as well as system loading).

Due to the uncertainty in the usage of the building, the weather in different areas across the region, and the reliability of the HVAC subsystems in the buildings, a key element in implementing such an arrangement is a distributed, real-time command and control system capable of performing distributed optimization of interconnected assets. The power broker and the company must be able to communicate information about asset condition and mission between the control systems for their electrical generation and HVAC systems and the subsystems must react to sensed changes in the environment (occupancy, weather, equipment status) to optimize the fleet level performance of the network.

Realization of enterprise-wide optimization of this sort will require substantial progress in a number of technical areas: distributed, embedded modeling tools that allow low resolution modeling of the external system combined with high resolution modeling of the local system, resident at each node in the enterprise; distributed optimization algorithms that make use of the embedded modeling architecture to produce near optimal operations; fault-tolerant, networked control systems that allow control loops to operate across unreliable network connections; and low-cost, fault-tolerant, reconfigurable hardware and software architectures.

A very closely related problem is that of C4ISR (command, control, communications, computers, intelligence, surveillance, and reconnaissance) in military systems. Here also, networked systems are revolutionizing the capabilities for continuous planning and asset allocation, but new research is needed in providing robust solutions that give the required performance in the presence of uncertainty and adversaries. The underlying issues and techniques are almost identical to enterprise level resource allocation, but the environment in which they must perform is

much more challenging for military applications. Control concepts will be essential tools for providing robust performance in such dynamic, uncertain, and adversarial environments.

3.3 Robotics and Intelligent Machines

It is my thesis that the physical functioning of the living individual and the operation of some of the newer communication machines are precisely parallel in their analogous attempts to control entropy through feedback. Both of them have sensory receptors as one stage in their cycle of operation: that is, in both of them there exists a special apparatus for collecting information from the outer world at low energy levels, and for making it available in the operation of the individual or of the machine. In both cases these external messages are not taken neat, but through the internal transforming powers of the apparatus, whether it be alive or dead. The information is then turned into a new form available for the further stages of performance. In both the animal and the machine this performance is made to be effective on the outer world. In both of them, their performed action on the outer world, and not merely their intended action, is reported back to the central regulatory apparatus.

Norbert Wiener, from *The Human Use of Human Beings: Cybernetics and Society*, 1950 [34].

Robotics and intelligent machines refers to a collection of applications involving the development of machines with human-like behavior. While early robots were primarily used for manufacturing, modern robots include wheeled and legged machines capable of competing in robotic competitions and exploring planets, unmanned aerial vehicles for surveillance and combat, and medical devices that provide new capabilities to doctors. Future applications will involve both increased autonomy and increased interaction with humans and with society. Control is a central element in all of these applications and will be even more important as the next generation of intelligent machines are developed.

Background and History

The goal of cybernetic engineering, already articulated in the 1940s and even before, has been to implement systems capable of exhibiting highly flexible or “intelligent” responses to changing circumstances. In 1948, the MIT mathematician Norbert Wiener gave a widely read, albeit completely non-mathematical, account of cybernetics [33]. A more mathematical treatment of the elements of engineering cybernetics was presented by H. S. Tsien in 1954, driven by problems related to control of missiles [32]. Together, these works and others of that time form much of the intellectual basis for modern work in robotics and control.

The early applications leading up to today’s robotic systems began after World War II with the development of remotely controlled mechanical manipulators, which used master-slave mechanisms. Industrial robots followed shortly thereafter, starting with early innovations in Computer Numerically Controlled (CNC) machine tools. Unimation, one of the early robotics companies, installed its first robot in a General Motors plant in 1961. Sensory systems were added to allow robots to respond to changes in their environment and by the 1960s many new robots were capable of grasping, walking, seeing (through binary vision), and even responding to simple voice commands.

The 1970s and 80s saw the advent of computer controlled robots and the field of robotics became a fertile ground for research in computer science and me-




Figure 3.6. *The Mars Sojourner and Sony AIBO robots. Photographs courtesy of Jet Propulsion Laboratory and Sony.*

chanical engineering. Manufacturing robots became commonplace (led by Japanese companies) and a variety of tasks ranging from mundane to high precision, were undertaken with machines. Artificial intelligence (AI) techniques were also developed to allow higher level reasoning, including attempts at interaction with humans. At about this same time, new research was undertaken in mobile robots for use on the factory floor and remote environments.

Two accomplishments that demonstrate the successes of the field are the Mars Sojourner robot and the Sony Aibo robot, shown in Figure 3.6. Sojourner successfully maneuvered on the surface of Mars for 83 days starting in July 1997 and sent back live pictures of its environment. The Sony AIBO robot debuted in June of 1999 and was the first “entertainment” robot that was mass marketed by a major international corporation. It was particularly noteworthy because of its use of AI technologies that allowed it to act in response to external stimulation and its own judgment.

It is interesting to note some of the history of the control community in robotics. The IEEE Robotics and Automation Society was jointly founded in the early 1980s by the Control Systems Society and the Computer Society, indicating the mutual interest in robotics by these two communities. Unfortunately, while many control researchers were involved active in robotics, the control community did not play a leading role in robotics research throughout much of the 1980s and 90s. This was a missed opportunity since robotics represents an important collection of applications that combines ideas from computer science, artificial intelligence, and control. New applications in (unmanned) flight control, underwater vehicles, and satellite systems are generating renewed interest in robotics and many control researchers are becoming active in this area.

Despite the enormous progress in robotics over the last half century, the field is very much in its infancy. Today’s robots still exhibit extremely simple behaviors compared with humans and their ability to locomote, interpret complex sensory inputs, perform higher level reasoning, and cooperate together in teams is limited.

Indeed, much of Wiener's vision for robotics and intelligent machines remains unrealized. While advances are needed in many fields to achieve this vision—including advances in sensing, actuation, and energy storage—the opportunity to combine the advances of the AI community in planning, adaptation, and learning with the techniques in the control community for modeling, analysis, and design of feedback systems presents a renewed path for progress. This application area is strongly linked with the Panel's recommendations on the integration of computing, communication and control, development of tools for higher level reasoning and decision making, and maintaining a strong theory base and interaction with mathematics.

Challenges and Future Needs

The basic electromechanical engineering and computing capabilities required to build practical robotic systems have evolved over the last half-century to the point where today there exist rapidly expanding possibilities for making progress toward the long held goals of intelligence and autonomy. The implementation of principled and moderately sophisticated algorithms is already possible on available computing hardware and more capability will be here soon. The successful demonstration of vision guided automobiles operating at high speed, the use of robotic devices in manufacturing, and the commercialization of mobile robotic devices attest to the practicality of this field.

Robotics is a broad field; the perspectives afforded by computer science, control, electrical engineering, mechanical engineering, psychology, and neuroscience all yield important insights. Even so, there are pervasive common threads, such as the understanding and control of spatial relations and their time evolution. The emergence of the field of robotics has provided the occasion to analyze, and to attempt to replicate, the patterns of movement required to accomplish useful tasks. On the whole, this has been sobering experience. Just as the ever closer examination of the physical world occasionally reveals inadequacies in our vocabulary and mathematics, roboticists have found that it is quite awkward to give precise, succinct descriptions of effective movements using the syntax and semantics in common use. Because the motion generated by a robot is usually its *raison d'être*, it is logical to regard motion control as being a central problem. Its study has raised several new questions for the control engineer relating to the major themes of feedback, stability, optimization, and estimation. For example, at what level of detail in modeling (i.e. kinematic or dynamic, linear or nonlinear, deterministic or stochastic, etc.) does optimization enter in a meaningful way? Questions of coordination, sensitivity reduction, stability, etc. all arise.

In addition to these themes, there is the need for development of appropriate software for controlling the motion of these machines. At present there is almost no transportability of robotic motion control languages. The idea of vendor independent languages that apply with no change to a wide range of computing platforms and peripherals has not yet been made to work in the field of robotics. The clear success of such notions when applied to operating systems, languages, networks, disk drives and printers makes it clear that this is a major stumbling block. What is missing is a consensus about how one should structure and standardize a "motion

description language.” Such a language should, in addition to other things, allow one to implement compliance control in a general and natural way.

Another major area of study is adaptation and learning. As robots become more commonplace, they will need to become more sophisticated in the way they interact with their environment and reason about the actions of themselves and others. The robots of science fiction are able to learn from past experience, interact with humans in a manner that is dependent on the situation, and reason about high level concepts to which they have not been previously exposed. In order to achieve the vision of intelligent machines that are common in our society, major advances in machine learning and cognitive systems will be required. Robotics provides an ideal testbed for such advances: applications in remote surveillance, search and rescue, entertainment, and personal assistance are all fertile areas for driving forward the state of the art.

In addition to better understanding the actions of individual robots, there is also considerable interest and opportunity in cooperative control of teams of robots. The U.S. military is considering the use of multiple vehicles operating in a coordinated fashion for surveillance, logistical support, and combat, to offload the burden of dirty, dangerous, and dull missions from humans. Over the past decade, several new competitions have been developed in which teams of robots compete against each other to explore these concepts. Perhaps the best known of these is RoboCup, which is described briefly in the following vignette.

Vignette: RoboCup: a testbed for autonomous collaborative behavior in adversarial environments (Raffaello D’Andrea, Cornell University)

RoboCup is an international collection of robotics and artificial intelligence (AI) competitions. The competitions are fully autonomous (no human intervention) head to head games, whose rules are loosely modeled after the human game of soccer; each team must attempt to score more goals than the opponent, subject to well defined rules and regulations (such as size restrictions, collision avoidance, etc.) The three main competitions are known as the Simulation League, the F2000 League, and the F180 League,

The F180 League is played by 6 inch cube robots on a 2 by 3 meter table (see Figure 3.7, and can be augmented by a global vision system; the addition of global vision shifts the emphasis away from object localization and computer vision, to collaborative team strategies and aggressive robot maneuvers. In what follows, we will describe Cornell’s experience in the F180 League at the 1999 competition in Stockholm, Sweden and the 2000 competition in Melbourne, Australia.

Cornell was the winner of the F180 League in both 1999, the first year it entered the competition, and 2000. The team’s success can be directly attributed to the adoption of a systems engineering approach to the problem, and by emphasizing system dynamics and control. The systems engineering approach was instrumental in the complete development of a competitive team in only 9 months (for the 1999 competition). Twenty-five students, a mix of first year graduate students and seniors representing computer sci-




Figure 3.7. *F180 league RoboCup soccer. Photograph courtesy Raffaello D'Andrea.*

ence, electrical engineering, and mechanical engineering, were able to construct two fully operational teams by effective project management, by being able to capture the system requirements at an early stage, and by being able to cross disciplinary boundaries and communicate among themselves. A hierarchical decomposition was the means by which the problem complexity was rendered tractable; in particular, the system was decomposed into estimation and prediction, real time trajectory generation and control, and high level strategy.

Estimation and prediction entailed relatively simple concepts from filtering, tools known to most graduate students in the area of control. In particular, smoothing filters for the vision data and feed-forward estimators to cope with system latency were used to provide an accurate and robust assessment of the game state. Trajectory generation and control consisted of a set of primitives that generated feasible robot trajectories; various relaxation techniques were used to generate trajectories that (1) could quickly be computed in real time (typically less than 1000 floating point operations), and (2) took full advantage of the inherent dynamics of the vehicles. In particular, feasible but aggressive trajectories could quickly be generated by solving various relaxations of optimal control problems. These primitives were then used by the high level strategy, essentially a large state-machine.

The high-level strategy was by far the most ad-hoc and heuristic component of the

Cornell RoboCup team. The various functions that determined whether passes and interceptions were possible were rigorous, in the sense that they called upon the provably effective trajectory and control primitives, but the high level strategies that determined whether a transition from defense to offense should be made, for example, or what play should be executed, relied heavily on human judgment and observation. As of the writing of this summary, most of the efforts at Cornell have shifted to understanding how the design and verification of high level strategies that respect and fully utilize the system dynamics can take place.

Certain robotic applications, such as those that call for the use of vision systems to guide robots, now require the use of computing, communication and control in an integrated way. The computing that is to be done must be opportunistic, i.e. it must be tailored to fit the needs of the specific situation being encountered. The data compression that is needed to transmit television signals to a computer must be done with a view toward how the results will be used by the control system. It is both technologically difficult and potentially dangerous to build complex systems that are controlled in a completely centralized way. For this reason we need to decide how to distribute the control function over the communication system. Recent work on the theory of communication protocols has made available better methods for designing efficient distributed algorithms. This work can likely be adapted in such a way as to serve the needs of robotic applications.

Finally, we note the need to develop robots that can operate in highly unstructured environments. This will require considerable advances in visual processing and understanding, complex reasoning and learning, and dynamic motion planning and control. Indeed, a framework for reasoning and planning in these unstructured environments will likely require new mathematical concepts that combine dynamics, logic, and geometry in ways that are not currently available. One of the major applications of such activities is in the area of remote exploration (of the earth, other planets, and the solar system), where human proxies will be used for continuous exploration to expand our understanding of the universe.

Other Trends in Robotics and Intelligent Machines

In addition to the challenges and opportunities described above, there are many other trends that are important for advances in robotics and intelligent machines and that will drive new research in control.

Mixed Initiative Systems and Human Interfaces It seems clear that more extensive use of computer control, be it for factories, automobiles or homes, will be most effective if it comes with a natural human interface. Having this goal in mind, one should look for interfaces which are not only suitable for the given application but which are sufficiently general so that, with minor modification, they can serve in related applications as well. Progress in this area will not only require new insights into processing of visual data (described above), but a better understanding of the *interactions* of humans with machines and computer controlled systems.

One program underway in the United States is exploring the use of “variable autonomy” systems, in which machines controlled by humans are given varying levels of command authority as the task evolves. Such systems involve humans that are integrated with a computer-controlled system in such a way that the humans may be simultaneously receiving instructions from and giving instructions to a collection of machines. One application of this concept is a semi-automated air traffic control system, in which command and control computers, human air traffic controllers, flight navigation systems, and pilots have varying levels of responsibility for controlling the airspace. Such a system has the possibility of combining the strengths of machines in rapid data processing with the strengths of humans in complex reasoning, but will require substantial advances in understanding of man-machine systems.

Control Using High Data-Rate Sensors Without large expenditure, we are able to gather and store more pictures and sounds, temperatures and particle counts, than we know how to use. We continue to witness occasional catastrophic failures of our man-machine systems, such as those used for transportation, because we do not correctly interpret or appropriately act on the information available to us. It is apparent that in many situations collecting the information is the easy part. Feedback control embodies the idea that performance can be improved by coupling measurement directly to action. Physiology provides many examples attesting to the effectiveness of this technique. However, as engineers and scientists turn their attention to the highly automated systems currently being built by the more advanced manufacturing and service industries, they often find that the direct application of feedback control is frustrated by a web of interactions which make the smallest conceptual unit too complex for the usual type of analysis. In particular, vision guided systems are difficult to design and often fail to be robust with respect to lighting conditions and changes in the environment. In order to proceed, it seems, design and performance evaluation must make more explicit use of ideas such as adaptation, self-diagnosis and self-optimization.

Indications are that the solution to the problems raised above will involve active feedback control of the perceptual processes, an approach which is commonplace in biology. One area that has received considerable attention is the area of active vision in which the vision sensor is controlled on the basis of the data it generates. Other work involves tuning the vision processing algorithms on basis of the data collected. The significant progress now being made toward the resolution of some of the basic problems results, in large part, from the discovery and aggressive use of highly nonlinear signal processing techniques. Examples include the variational theories that have been brought to bear on the image segmentation problem, the theories of learning based on computational complexity, and information theoretic based approaches to perceptual problems. Attempts to incorporate perceptual modules into larger systems, however, often raise problems about communication and distributed computation which are not yet solved.

Related to this is the problem of understanding and interpreting visual data. The technology for recognizing voice commands is now sophisticated enough to

see use in many commercial systems. However, the processing and interpretation of image data is in its infancy, with very few systems capable of decision making and action based on visual data. One specific example is understanding of human motion, which has many applications in robotics. While it is possible for robots to react to simple gestures, we do not yet have a method for describing and reasoning about more complex motions, such as a person walking down the street, stooping to pick up a penny, and being bumped by someone that did not see them stop. This sort of interpretation requires representation of complex spatial and symbolic relationships that are beyond our currently available tools in areas such as system identification, state estimation, and signal to symbol translation.

Medical Robotics While the word “revolution” is overused, one can honestly say that computer and robotic technology holds the potential to have a revolutionary impact on the practice of medical surgery. By extending surgeons’ ability to plan and carry out surgical interventions more accurately and in a minimally invasive manner, computer-aided and robotic surgical systems should reduce surgical and hospital costs, improve clinical outcomes, and improve the efficiency of health care delivery. The ability to consistently carry out surgical procedures and to comprehensively log key patient and procedure outcome data should also lead to long term improvements in surgical practice.

Robotic technology is useful in a variety of surgical contexts. For example, the “RoboDoc” surgical assistant uses the precision positioning and drilling capabilities of robots to improve the fit of implants during total hip replacement [4]. The improved fit leads to significantly fewer complications and longer lasting implants. Similarly, 3-dimensional imaging data can drive the precision movement of robot arms during stereotactical brain surgery, thereby reducing the risk of collateral brain damage. The DaVinci system from Intuitive Surgical uses teleoperation and force-reflecting feedback methods to enable minimally invasive coronary procedures that would otherwise require massively invasive chest incisions [26]. Figure 3.8 shows the ZEUS system developed by Computer Motion, Inc. a modified version of which was used in 2001 to allow a surgeon in New York to operate on a 68 year old woman in Strasbourg, France [23] These are only a few of the currently approved robotic surgical systems, with many, many more systems in clinical trials and laboratory development.

While medical robotics is becoming a reality, there are still many open research and development questions. Clearly, medical robotics will benefit from the same future advances in computing, communication, sensing, and actuation technology that will broadly impact all future control systems. However, the issue of system and software reliability is fundamental to the future of medical robotics. Formal methods for system verification of these highly nonlinear, hybrid, and uncertain systems, as well as strategies for extreme fault tolerance are clearly needed to ensure rapid and widespread adoption of these technologies. Additionally, for the foreseeable future, robotic medical devices will be assistants to human surgeons. Consequently, their human/machine interfaces must be able to deal with the complex contexts of crowded operating rooms in an absolutely reliable way, even during unpredictable



Figure 3.8. *The ZEUS (tm) Robotic Surgical System, developed by Computer Motion Inc., is capable of performing minimally invasive microsurgery procedures from a remote location. Photograph courtesy of Computer Motion Inc.*

surgical events.

3.4 Biology and Medicine

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

Mahlon B. Hoagland and B. Dodson, *The Way Life Works* [15].

At a variety of levels of organization—from molecular to cellular to organismal—biology is becoming more accessible to approaches that are commonly used in engineering: mathematical modeling, systems theory, computation, and abstract approaches to synthesis. Conversely, the accelerating pace of discovery in biological science is suggesting new design principles that may have important practical applications in man-made systems. This synergy at the interface of biology and engineering offers unprecedented opportunities to meet challenges in both areas. The principles of control are central to many of the key questions in biological engineering and will play an enabling role in the future of this field.

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

The Panel also identified potential roles for control in medicine and biomedical research. These included intelligent operating rooms and hospitals, from raw data to decisions; image guided surgery and therapy; hardware and soft tissue integration; fluid flow control for medicine and biological assays; and the development of physical and neural prosthesis. Many of these areas have substantial overlap with robotics and some have been discussed already in Section 3.3.

We focus in this section on three interrelated aspects of biological systems: molecular biology, integrative biology, and medical imaging. These areas are representative of a larger class of biological systems and demonstrate how principles from control can be used to understand nature and to build engineered systems.

Molecular Biology²

The life sciences are in the midst of a major revolution that will have fundamental implications in biological knowledge and medicine. Genomics has as its objective the complete decoding of DNA sequences, providing a “parts list” for the proteins present in every cell of the organism being studied. Proteomics is the study of the three-dimensional structure of these complex proteins. The shape of a protein determines its function: proteins interact with each other through “lego-like” fitting

²The Panel would like to thank Eduardo Sontag for his contributions to this section, based on his Reid Prize plenary lecture at the 2001 SIAM Annual Meeting.

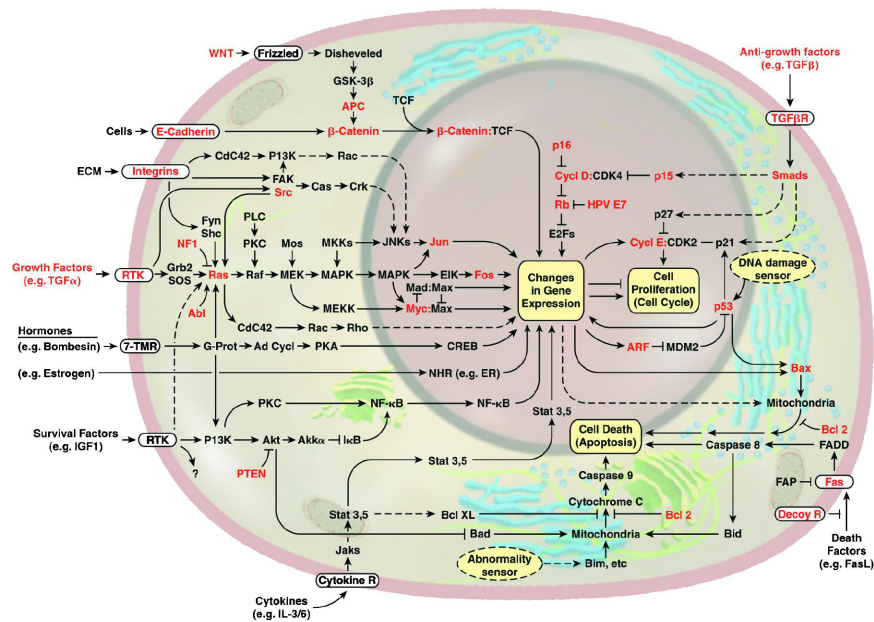


Figure 3.9. *The wiring diagram of the growth signaling circuitry of the mammalian cell [14].*

of parts in “lock and key” fashion, and their conformation also enhances or represses DNA expression through selective binding.

One may view cell life as a huge “wireless” network of interactions among proteins, DNA, and smaller molecules involved in signaling and energy transfer. As a large system, the external inputs to a cell include physical signals (UV radiation, temperature) as well as chemical signals (drugs, hormones, nutrients). Its outputs include chemicals that affect other cells. Each cell can be thought of, in turn, as composed of a large number of subsystems involved in cell growth, maintenance, division, and death. A typical diagram describing this complex set of interactions is shown in Figure 3.9.

The study of cell networks leads to the formulation of a large number of questions. For example, what is special about the information-processing capabilities, or input/output behaviors, of such biological networks? Can one characterize these behaviors in terms familiar to control theory (e.g., transfer functions or Volterra series)? What “modules” appear repeatedly in cellular signaling cascades, and what are their system-theoretic properties? Inverse or “reverse engineering” issues include the estimation of system parameters (such as reaction constants) as well as the estimation of state variables (concentration of protein, RNA, and other chemical substances) from input/output experiments. Generically, these questions may be viewed respectively as the identification and observer (or filtering) problems which are at the center of much of control theory.

One can also attempt to better understand the stability properties of the

various cascades and feedback loops that appear in cellular signaling networks. Dynamical properties such as stability and existence of oscillations in such networks are of interest, and techniques from control theory such as the calculation of robustness margins will play a central role in the future. At a more speculative (but increasingly realistic) level, one wishes to study the possibility of using control strategies (both open and closed loop) for therapeutic purposes, such as drug dosage scheduling.

The need for mathematical models in cellular biology has long been recognized, and indeed many of the questions mentioned above have been studied for the last 20 or 30 years. What makes the present time special is the availability of huge amounts of data—generated by the genomics and proteomics projects, and research efforts in characterization of signaling networks—as well as the possibility for experimental design afforded by genetic engineering tools (gene knock-outs and insertions, PCR) and measurement technology (green fluorescent protein and other reporters, and gene arrays). Control-oriented modeling and analysis of feedback interconnections is an integral component of building effective models of biological systems.

Feedback and uncertainty. From a theoretical perspective, feedback serves to minimize uncertainty and increase accuracy in the presence of noise. The cellular environment is extremely noisy in many ways, while at the same time variations in levels of certain chemicals (such as transcriptional regulators) may be lethal to the cell. Feedback loops are omnipresent in the cell and help regulate the appropriate variations. It is estimated, for example, that in *E. coli* about 40% of transcription factors self-regulate. One may ask whether the role of these feedback loops is indeed that of reducing variability, as expected from principles of feedback theory. Recent work tested this hypothesis in the context of tetracycline repressor protein (TetR) [7]. An experiment was designed in which feedback loops in TetR production were modified by genetic engineering techniques, and the increase in variability of gene expression was correlated with lower feedback “gains,” verifying the role of feedback in reducing the effects of uncertainty. Modern experimental techniques will afford the opportunity for testing experimentally (and quantitatively) other theoretical predictions, and this may be expected to be an active area of study at the intersection of control theory and molecular biology.

Necessity of embedded structures in regulation loops. Another illustration of the interface between feedback theory and modern molecular biology is provided by recent work on chemotaxis in bacterial motion. *E. coli* moves, propelled by flagella, in response to gradients of chemical attractants or repellents, performing two basic types of motions: *tumbles* (erratic turns, with little net displacement) and *runs*. In this process, *E. coli* carries out a stochastic gradient search strategy: when sensing increased concentrations it stops tumbling (and keeps running), but when it detects low gradients it resumes tumbling motions (one might say that the bacterium goes into “search mode”).

The chemotactic signaling system, which detects chemicals and directs motor actions, behaves roughly as follows: after a transient nonzero signal (“stop tumbling, run toward food”), issued in response to a change in concentration, the system adapts and its signal to the motor system converges to zero (“OK, tumble”). This adaptation happens for any constant nutrient level, even over large

ranges of scale and system parameters, and may be interpreted as robust (structurally stable) rejection of constant disturbances. The internal model principle of control theory implies (under appropriate technical conditions) that there must be an embedded integral controller whenever robust constant disturbance rejection is achieved. Recent models and experiments succeeded in finding, indeed, this embedded structure [5, 35].

This work is only one of the many possible uses of control theoretic knowledge in reverse engineering of cellular behavior. Some of the deepest parts of the theory concern the necessary existence of embedded control structures, and in this manner one may expect the theory to suggest appropriate mechanisms and validation experiments for them.

Genetic circuits. Biomolecular systems provide a natural example of hybrid systems, which combine discrete and logical operations (a gene is either turned on or off for transcription) and continuous quantities (such as concentrations of chemicals) in a given cell or in a cell population). Complete hybrid models of basic circuits have been formulated, such as the lysogeny/lysis decision circuit in bacteriophage λ [24].

Current research along these lines concerns itself with the identification of other naturally occurring circuits, as well as with the engineering goal of *designing* circuits to be incorporated into delivery vehicles (bacteria, for example), for therapeutic purposes. This last goal is, in principle, mathematically in the scope of realization theory, that branch of systems theory which deals with the synthesis of dynamical systems which implement a specified behavior.

Integrative Biology³

Control also has a role to play in understanding larger scale organisms, such as insects and animals. The components of these *integrative* biological systems are becoming much better understood and, like molecular systems, it is becoming evident that systems principles are required to build the next level of understanding. This understanding of natural systems will enable new approaches to engineered systems, as we begin to build systems with the efficiency, robustness, and versatility of the natural world. We focus here on the problem of locomotion, for which there has been substantial recent work (see [12] for a recent review).

Integrative studies of locomotion have revealed several general principles that underly a wide variety of organisms. These include energy storage and exchange mechanisms in legged locomotion and swimming, nonpropulsive lateral forces that benefit stability and maneuverability, and locomotor control systems that combine mechanical reflexes with multimodal sensory feedback and feedforward control. Locomotion, especially at the extremes of what is found in nature, provides a rich set of examples that have helped elucidate a variety of structure-function relationships in biological systems.

Control systems and feedback play a central role in locomotion. A suite of neurosensory devices are used within the musculoskeletal system and are active throughout each cycle of locomotion. In addition, the viscoelastic dynamics of the

³The Panel would like to thank Michael Dickinson for his contributions to this section.




Figure 3.10. Overview of flight behavior in a fruit fly, *Drosophila*. (A) Cartoon of the adult fruit fly showing the three major sensor structures used in flight: eyes, antennae, and halteres (detect angular rotations). (B) Example flight trajectories over a 1 meter circular arena, with and without internal targets. (C) A schematic control model of the flight system. Figure and description courtesy of Michael Dickinson.

musculoskeletal system play a critical role in providing rapid feedback paths that enable stable operation. Rapid feedback from both mechanical and neural pathways is integrated with information from eyes, ears, noses and other sensing organs used to control the overall motion of an animal and provide robust operation in a wide variety of environments.

The process that gives rise to locomotion is a complex one, as illustrated in Figure 3.10 for the flight behavior of a fruit fly. Each element of the flight control system has enormous complexity in itself, with the interconnection (grossly simplified in the figure) allowing for a very rich set of behaviors. The sensors, actuators, and control systems for insects such as the fly are highly evolved, so that the dynamics of the system play strongly into the overall capabilities of the organism.

From the perspective of control theory, the performance, robustness, and flaw tolerance of the fly's flight control system represents a gold standard by which all

autonomous systems might be judged. Flies can manage to stay in the air with torn wings, missing legs, blind eyes, or when burdened with twice their weight in additional mass. The fact that the control algorithms yielding this behavior reside in a brain the size of a sesame seed raises the bar for any biomimetic effort attempting to match its performance. If the principles that engender a fly with such robust agility could be discovered and formalized for general use, the results might catalyze a revolution in the design, fabrication, and implementation of control systems.

Similarly, the use of control tools to understand the fly's flight control system represents a systems approach to biology that will be important for understanding the general principles of locomotion systems and allow new understanding of integrative biological principles.

This synergy between biology and control in insect flight is but one example of many that are possible and that form a rich source of scientific and engineering activity. Additional areas of overlap include multiresolution modeling and analysis of (nongeneric, designed, evolved, heterogeneous) multiscale systems, and integrated communications and computing for control of and with pervasive, distributed, embedded networks. Biological systems are also exceptionally capable of transforming raw data into information and knowledge, and eventually into decision and action. These are precisely the problems that confront us in building engineering systems and the interaction of biologists and control researchers is destined to be fruitful.

Medical Systems⁴

Control is also an essential element in the burgeoning field of biomedicine. Some of these applications, such as robot surgery, have already been discussed in the context of robotics and intelligent machines (see Section 3.3). We consider two additional examples here: image guided therapy (IGT) and image guided surgery (IGS).

Image guided therapy and surgery provide illustrations of how to use biomedical engineering principles to develop general-purpose software methods that can be integrated into complete therapy delivery systems. Such systems will support more effective delivery of many image-guided procedures—biopsy, minimally invasive surgery, and radiation therapy, among others. A key element is controlled active vision. To understand its role in the therapeutic process, and to appreciate the current usage of images before, during, and after treatment, one must consider the four main components of IGT and IGS: localization, targeting, monitoring and control.

To use controlled active imaging one must first develop robust algorithms for *segmentation*, automated methods that create patient-specific models of relevant anatomy from multi-modal imagery, and *registration*, automated methods that align multiple data sets with each other and with the patient. These technologies must then be integrated into complete and coherent image guided therapy delivery systems and validated using performance measures established in particular application areas. Control enters at almost every stage of the process. For example, control-theoretic methods can be essential for the success of the *deformable or active*

⁴The Panel would like to thank Allen Tannenbaum for his contributions to this section.

contours technique in active vision for therapeutic and surgical procedures. These are autonomous processes that employ image coherence in order to track features of interest over time. They have been used for segmentation and edge detection as well. For dynamically changing imagery in a surgical environment, Kalman filtering has been important in estimating the position of an active contour at a given time given its previous position. This estimated data may be used then in a closed loop visual tracker.

Further, speed and robustness are very important in interventional magnetics, which uses magnetic resonance imagery (MRI) during surgery. Here surgeons can operate in an open MRI device, and use the images to guide their procedure. Fast segmentation is of paramount importance, and one can use active contours very effectively when coupled with an estimation scheme to extract key features (such as a brain tumor or breast cyst).

Image registration is the process of establishing a common geometric reference frame between two or more data sets from the same or different imaging modalities possibly taken at different times. Multimodal registration proceeds in several steps. First, each image or data set to be matched should be individually calibrated, corrected from imaging distortions, and cleaned from noise and imaging artifacts. Next, a measure of dissimilarity between the data sets must be established, so we can quantify how close an image is from another after transformations are applied to them. Once features have been extracted from each image, they must be paired to each other. Then, a similarity measure between the paired features is formulated which can be formulated as an optimization problem of the type many times used in control.

Optimal transport methods have proved very useful for this. Optimal transport ideas have been used in nonlinear stability analysis, and very similar concepts lead to a measure of similarity between images which can be employed in registration and data fusion.

In general, IGT and IGS will benefit enormously from systems oriented ideas. At this point most of the control is being done by the computer vision and medical imaging community. By building stronger ties between these groups and the control community, it will be possible to make more rapid progress and to leverage advances from other applications. In addition, the specific features of this class of problems will drive new advances in control theory and technology, which can then be exported to other areas.




Figure 3.11. *Intel Pentium IV wafer and die.*

3.5 Materials and Processing⁵

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

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

Background and History

At least one materials or chemicals process is involved in the manufacture of nearly every commercial product, including microprocessors, consumer products such as detergents and shampoo, books, diskettes, disk drives, video cassette recorders, food, pharmaceuticals, adhesives, automobile dashboards, and aircraft interiors. Feedback controllers for these processes provide improved product quality, reduced

⁵The Panel would like to thank Richard Braatz and Frank Doyle for their contributions to this section.

materials and energy usage, reduced environmental impact, improved safety, and the reduced costs needed for U.S. industry to be competitive in the global economy.

By the late 1960s, process control had been implemented liberally to chemical and materials processes, primarily in the form of single-loop controllers with little communications between controllers. Multi-variable control began to be implemented in the 1970s, including some rather large scale processes such as the control of uniformity in plastic film and paper machines. The use of multi-variable control grew rapidly throughout the 1980s and 1990s. Over the last 25 years, multi-variable optimal control in the form of model predictive control has become a standard control technique in the process industries for addressing actuator and state constraints, which are quite prevalent in chemicals and materials processes. Model predictive control explicitly takes constraints into account during the on-line calculation of the manipulated variable actions. In 1996, more than 2000 applications of model predictive control were reported by the control vendors of that time (e.g., DMC Corporation, Setpoint Inc., Honeywell, Adersa, and Treiber Controls). Applications have been reported in a wide range of industries including refining, petrochemical, pulp and paper, air separation, food processing, furnaces, aerospace, and automotive. In recent years model predictive control algorithms have been developed that enable their application to very large scale process control problems.

This should not be taken, however, to indicate that all process control problems have been solved. New control techniques are needed that address all of the characteristics of the most challenging chemicals and materials processes.

Current Challenges and Future Needs

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

Advances in materials and processing are important for a variety of industries in which control of complex process systems enables growth in the world economy. One example is the microelectronics industry, which has an average annual growth of 20%, with sales of \$200 billion in 2001. As described by the International Technology Roadmap for Semiconductors,⁶ high performance feedback control will be needed to achieve the small length scales required for the next generation of microelectronic devices that are predicted (and hence demanded) by Moore's Law.

A second example is the pharmaceuticals industry, which is growing at 10-20% annually, with sales of \$150 billion in 2000. The primary bottleneck to the operation of production-scale drug manufacturing facilities is associated with difficulties in controlling the size and shape distribution of crystals produced by complex crys-

⁶<http://public.itrs.net/>




Figure 3.12. *Microscope image of paracetamol crystals (paracetamol is the active ingredient in Tylenol [right]), which shows the variability in crystal shape that can occur at a single time instance in a pharmaceutical crystallizer. Image courtesy of Richard Braatz.*

tallization processes (see Figure 3.12). Crystallization processes typically involve growth, agglomeration, nucleation, and attrition mechanisms which can be affected by particle-particle collisions. Poor control of this crystal size distribution can completely halt the production of pharmaceuticals, causing both economic and medical hardship.

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

As in many other application areas, new sensor technology is creating new opportunities for control. On-line sensors—including laser backscattering, video microscopy, ultraviolet, infrared, and Raman spectroscopy—are becoming more robust and less expensive, and are appearing in more manufacturing processes. Many of these sensors are already being used by current process control systems, but more sophisticated signal processing and control techniques are needed to more effectively use the real-time information provided by these sensors. Control engineers can also contribute to the design of even better sensors which are still needed, for example, in the microelectronics industry. As elsewhere, the challenge is making use of the large amounts of data provided by these new sensors in an effective manner. In addition, a control-oriented approach to modeling the essential physics of the underlying processes is required to understand fundamental limits on observability of the internal state through sensor data.

Another common feature in materials and process control is the inherent com-

plexity of the underlying physical processing. Modern process systems exhibit very complex nonlinear dynamics, including substantial model uncertainty, actuator and state constraints, and high dimensionality (usually infinite). These systems are often best described by tightly coupled systems of algebraic equations and stochastic partial integrodifferential equations with widely varying time and length scales and significant nonlinearities. This is especially true in the microelectronics industry, where hundreds of stiff partial differential equations can be required for predicting product quality, for example, during the modeling of cluster formation and dissolution during fast-ramp annealing after ion bombardment. Other processes are best described by kinetic Monte Carlo simulations, with or without coupling to continuum equations, which can be run on serial or parallel computers. Both identification and control algorithms are needed that can *simultaneously* address the high complexity, high nonlinearity, and high dimensionality of these complex process systems. Furthermore, there is significant uncertainty associated with many of the kinetic parameters, even with improved sensors, so these algorithms need to be robust to model uncertainties.

Two specific areas that illustrate some of the challenges and future needs are control of particulate systems and biotechnology.

Control of Particulate Systems

Particulate processes are prevalent in a number of process industries including agricultural, chemical, food, minerals, and pharmaceutical. By some estimates, 60% of the products in the chemical industry are manufactured as particulates with an additional 20% using powders as ingredients. One of the key attributes of such systems is the distributed characterization of physical and chemical properties such as size, shape, morphology, porosity, molecular weight, etc. The underlying mechanisms which describe the evolution of such systems are captured by population balance models, which are coupled sets of hyperbolic partial differential and algebraic equations.

There are a number of challenges in the numerical solution of such equations, particularly when considering real-time applications such as model-based control. Critical in such models are the kernels or driving forces (e.g., nucleation, growth, agglomeration, breakup, etc.) that are typically not well characterized, and are often determined from process data via various identification techniques. These problems become increasingly complex as one considers higher-dimension population balance models (e.g., size and shape), where the number of parameters in the kernels grows rapidly with the increase in additional degrees of freedom. At the same time, there have been substantial advances in the domain of sensor technology, such that attributes like the particle size distribution can be measured in real-time by a variety of techniques including light scattering, ultrasound spectroscopy and hydrodynamic capillary separation. This leads to control formulations involving distributed measurement variables, highly nonlinear process models, nonlinear operating constraints, and complex hierarchical operating objectives.

To explore some of the major challenges, we consider three selected application areas—polymerization, granulation, and profile control.

Emulsion Polymerization. Increasing global competition for the production of higher quality polymer products at lower costs, coupled with a general trend away from new capital investments in the U.S., has placed considerable pressure on the process engineers in the U.S. to operate the existing polymer plants more efficiently and to use the same plant for the production of many different polymer products. Lack of sufficient controllability is a barrier to better product quality control in some polymer processes. In many polymer processes, better product quality requires minimizing/maximizing several product quality indices simultaneously. This multi-objective requirement may result in narrow ranges of process trajectories, putting a premium on the controllability of the process. For instance, in coatings, the product's composition, molecular weight, and particle size distributions should be maintained simultaneously in limited ranges to ensure the coating has a desired level of film formation, film strength, and gloss.

The critical link between these product quality indices and the operating process variables is often the distributed attribute such as the size distribution. In the past, such attributes were controlled indirectly using inferential control schemes, but on-line sensor technology brings the promise of real-time control of these properties. This motivates the development of refined quantitative relationships between the distributed quantities and the quality variables. While experimental techniques have been used to develop relationships that hold for limited operating conditions, these descriptions do not readily lend themselves to optimization, either in terms of productivity or reduction in variance.

Granulation. Granulation is a key step in many particulate processes where fine particles are agglomerated with the aid of a liquid binder into larger granules. It is often used to improve the visual appearance and/or taste of materials, improve the flowability of the materials, enable compaction and tableting, and reduce dustiness. The granulation process exhibits many characteristics common to other particulate processes such as crystallization and emulsion polymerization. Typically, a desired product quality can be inferred from the Particle Size Distribution (PSD) of a process. The ability to manipulate a PSD allows for control of the end product quality, but PSD control can pose a very difficult control problem due to the significant multi-variable interacting character of PSD systems. In some situations, values of the measured PSD may be constrained to a specified acceptable region in order to achieve a desired product quality.

As with many particulate processes, there is a rich interplay between mechanisms at the microscopic, mesoscopic and macroscopic levels in granulation, however, the fundamental knowledge to link these mechanisms for use in model-based control is rather limited. In particular, the trade-offs between model quality and complexity for various model uses have not been investigated systematically, leading to inadequate selections of model forms. Furthermore, granulation is a complex multi-scale process, including multi-number, dimension and time scales. The current status of granulation research clearly shows significant gaps between microscopic-level studies and plant-scale modeling, and also between the model forms and the use of models. Given such models and the already existing sensor technology, one can realize the tight regulation of this complex unit operation.

Profile Control. Though the systems described in this area are not strictly particulate processes, they share the attribute that a distributed variable is directly tied to product performance, hence many of the underlying mathematical constructs required for control are common to both classes of problems. The problems of controlling a “profile” arise in a number of rather different process industry unit operations, including polymer extrusion, cross direction control (paper, aluminum, glass, etc.), tubular chemical reactors, and advanced materials processing (photo-voltaic, microelectronic, etc.), to name a few. In some instances the properties of interest are measured in the cross direction (CD) giving rise to a 1-D profile control problem, or in other cases the quality attribute is measured in both the machine direction (MD) and CD, giving rise to a 2-D sheet control problem. In reaction unit operations, the extent of reaction across a spatial direction is a critical parameter that controls important quality indices. For example, in a pulp digester, the control of reaction extent profile (measured by the Kappa number) along the axial direction in the reactor enables the tight regulation of critical fiber properties, such as strength, which depend on the reaction path as well as the final conversion.

One of the interesting challenges that arises, for example, in the paper machine CD control problem is that hundreds of input/output variables are involved, complex temporal and spatial constraints must be maintained, and real-time requirements dictate solution times on the order of seconds. This is further complicated by non-ideal process behavior owing to paper shrinkage, lane shifting, production grade and rate changes—all of which give rise to significant plant-model mismatch, and hence a robust controller is required. As with the particulate problems, the sensor technology is changing rapidly, enabling richer formulations of controlled operation.

Biotechnology

While process control has played a relatively minor role in the biotechnology industries in past years, its value as an enabling technology is increasing, owing to novel developments in sensor technology coupled with advances in the mathematical characterization of intracellular behavior. Furthermore, the potential to realize efficient process design by accounting for operational issues (i.e., combined design and control), brings promise of reducing the development time for new plants, and maximizing the production interval in a patent window.

Classical bioreactor control focused on the manipulation of residence time, nutrient feed and the reactor environment (agitation, temperature, pH, etc.) in response to averaged system measurements (dissolved oxygen, temperature, pH, limited metabolite concentrations, etc.) Advances in sensor technology have enabled direct measurement and manipulation of intracellular mechanisms, and recent advances in quantitative modeling of bioprocesses allow a much more thorough understanding of the underlying biochemical processes. A number of the resulting model structures are inspired from systems engineering concepts, such as the cybernetic model which incorporates optimal control regulation of cellular processes, or the flux balance analysis which accounts for the convex constraint space available from a metabolic network. Population balance models also find application in this area, for example, in the description of age distributions in oscillating microbial cell

cultures. As with particulate systems, one can construct high-order population balance descriptions by accounting for the various elements of the physiological state space (DNA, RNA, protein, etc.) Commensurate with this increase in structural complexity is the possibility to achieve refined control objectives for bioprocessing, such as the control of distinct metabolic pathways.

An example of the opportunities that emerge from such increased understanding is the use of recombinant organizations to produce enzymes and proteins. Typically, the genes corresponding to the desired product are inserted into the microorganism through a plasmid. The first phase in the recombinant protein production involves increasing the cell productivity (biomass concentration), as the expression of the foreign protein imposes additional metabolic burden on the organism and decreases the growth rate. Once a sufficiently high biomass concentration is achieved, the inducer that expresses the inserted gene is added in the second phase resulting in the synthesis of the recombinant product. Therefore, the concentration of the inducer and the time at which the inducer is added are key variables in maximizing the amount of the recombinant protein formed. One specific example is the recombinant bioprocess involving chloramphenicol acetyltransferase (CAT) production in the genetically engineered strain of *E. coli* JM105. This strain produces both the green fluorescent protein (GFP) and CAT, when the metabolized inducer, arabinose, is added to the bioreactor. The objective is the maximization of the amount of CAT formed at the end of a batch. The manipulated variables are the glucose feed rate and the feed rate of arabinose, the inducer which turns on the expression of the desired product.

The use of GFP and its variants have revolutionized experimental biology by enabling accurate real-time reporting of intracellular location and interactions, which has proven valuable in determining the role and interactions of a number of proteins. GFP was cloned from the jellyfish, *Aequorea victoria*, in 1992 and has since been successfully expressed in a number of host organisms as a reporter of expression, as well as to identify the location of proteins within cells. GFP and its variants have been successfully used to quantitate intracellular expression levels in insect larvae, bacterial systems and mammalian cells. Owing to the optical nature of the signal, the development of sensing devices for industrial application is direct.

3.6 Other Applications

The previous sections have described some of the major application areas discussed by the Panel. However, there are certainly many more areas where ideas from control are being applied or could be applied. In this section we collect a few such areas which are more specialized than those discussed to this point. As before, these areas are not meant to be exhaustive, but rather representative of some of the new and exciting directions within control.

Environmental Science and Engineering⁷

It is now indisputable that human activities have altered the environment on a global scale. Problems of enormous complexity challenge researchers in this area and first among these is to understand the feedback systems that operate on the global scale. One of the challenges in developing such an understanding is the multi-scale nature of the problem, with detailed understanding of the dynamics of microscale phenomena such as microbiological organisms being a necessary component of understanding global phenomena, such as the carbon cycle. Two specific areas where control is relevant are atmospheric systems and microbiological ecosystems.

Atmospheric systems and pollution. Within the last few years “inverse modeling” has become an important technique in atmospheric science when there are unknown sources or sinks of a species. The essential problem is to infer an optimal global source (or sink) distribution of an atmospheric trace species from a set of global observations. This is equivalent to the following control problem: given a system governed by a set of PDEs and a set of noisy observations of the system, determine the optimal set of inputs that match the model to the data. Such a problem has relevance to atmospheric chemical transport models, of which CO₂ is perhaps the most important at the present time.

At present, inverse modeling for atmospheric species has been applied only to those compounds that are inert in the atmosphere or only react via simple mechanisms. One area that offers promise is the development of techniques for inverse modeling to trace species that undergo nonlinear atmospheric processes, such as ozone. The inverse modeling problem is closely related theoretically to the sensitivity analysis problem, wherein one seeks the sensitivity of spatially and temporally varying concentrations to uncertainties in input functions and variables. Atmospheric inverse modeling is an important application of ideas from control to estimate global source (and sink) distributions of trace species based on noisy, usually sparse measurements.

Microbiological ecosystems. To illustrate how ideas from control can play a role in microbiological ecosystems, consider the example of microbial biofilms. It is widely recognized that microbial biofilms are ubiquitous, resilient, responsive to their environment, and able to communicate through chemical signaling. Furthermore, specific genes, gene products, and regulatory networks that control how bacteria

⁷The Panel would like to thank Jared Leadbetter, Dianne Newman, and John Seinfeld for their contributions to this section.

communicate have been described in a variety of bacteria. To date, studies of biofilm development have been largely limited to studies of pure cultures. While much has been learned regarding the genetic pathways taken by a variety of organisms when transitioning from the planktonic to the sessile phase, little is known about how these pathways change in response to changes in the environment. Researchers now believe that at the scale most relevant to bacteria (the microscale), one of the most important environmental factors that affect biofilm development by a given species is the presence of other organisms. The study of such ecological networks is at the forefront of research in this area and the tools of control can play a major role developing systematic understanding of their complex interactions.

Another example is in the area of bacterial cells that live inside organisms. Although they have limited conventional sensing and decision making abilities, bacterial cells are able to rapidly assess and respond to changes in their metabolism by monitoring and maintaining relative pool sizes of an extraordinary number (thousands) of cellular building blocks/intermediates. A common theme that has emerged in understanding how this works is related to timing. Many changes in physiology are effected by responses to pauses brought about by a binding site of an enzyme not being occupied by a given building block. If a certain building block is depleted, the enzyme that would incorporate it into cellular material pauses “in wait.” Paused enzymes will often do or allow things that an occupied one does not. On one hand, this might result in the increased production of the missing metabolite to bring up the depleted pool to better reflect the size of the other building block pools, keeping things in balance. On the other, some enzymes pause as a result of many pools being depleted in concert, this signals to the cell that it has begun to exhaust its total resources and moves it into a starvation survival phase.

To control for the possible overproduction of certain pools, many enzymes involved in the early stages of building block synthesis become inactive if a binding site becomes occupied by a later or final product. With the knowledge of this control mechanism, industrial microbiologists have been able to obtain feedback inhibition mutant bacteria that over produce almost any desired amino acid.

In the natural world, work done on termites provides a model system for studying the role of feedback and control in such microbiological ecosystems. There is every reason to believe that termites can control the delivery of oxygen to, and the consumption of it within differing zones of the gut epithelium. By doing so, the termite should be able to protect and even control the activities of its oxygen sensitive microbiota—but the forms of feedback that the tissue receives and processes from the gut and atmosphere are not known. One could envision several ways in which the gut tissue might respond to oxygen and acetate concentrations to control oxygen delivery to, and diffusion into the gut compartment. An important question, and one which control can help provide an answer, is how the insect and gut tissues create, control, and maintain a very complex and fragile ecosystem.

Economics and Finance⁸

Many control tools have found applications in economics and there are many common mathematical methods between the fields in areas such as game theory, stochastic modeling and control, and optimization and optimal control.

Control theory also became an important tool in finance with the introduction of the Black-Scholes-Merton formula for pricing options in the early 1970s. In essence, they showed that the dynamic control of a portfolio of risky assets could be used to change its risk profile. In the extreme case of options, the risk could be completely eliminated by a dynamic trading strategy, and this led to a fair price for the option.

The general problem of pricing and hedging an option is one of optimal stochastic control, and involves dynamically trading financial assets to achieve desired payoffs or risk profiles. When placed in this control theory framework, the quantities of various assets held in a portfolio become the decision variables, the price movements (often random) of the assets are the dynamics of the system, and achieving a desired risk profile is the objective. In structure, they tend to deviate from control problems involving physical systems due to the fact that the dynamics of the system are dominated by uncertainty. That is, the movement of prices is modeled in a highly stochastic manner.

Control problems in finance, especially those related to hedging and pricing derivative securities, present a number of interesting challenges for the control community.

The securities being offered in the financial marketplace are becoming increasingly complex. That means that the pricing and hedging of these securities is also becoming increasingly complex. Examples already in existence include options that allow the holder to decide when to exercise the option, options on averages of prices, options on baskets of securities, options on options, etc. and these options can be written on stocks, futures, interest rates, the weather, earthquakes and catastrophes, etc. Hedging of these options is a challenging and rather daunting task for stochastic control theory.

The lack of robustness of dynamic schemes in use during the 1987 crash was another critical factor. Since modeling is itself a difficult problem, it is important that control schemes work well in the presence of modeling errors. This is especially true in finance, where uncertainties can be large and time varying. Often this uncertainty is handled in a stochastic manner. For instance, some models in finance assume that the volatility of an asset is stochastic. This has been used to explain certain deviations between markets and theory. More recently, researchers have been developing control and hedging schemes which explicitly account for model errors and uncertainty, and are designed to work well in their presence. This will be an area in which robust control theory has a great amount to contribute.

⁸The Panel would like to thank Jim Primbs for his contributions to this section.

Electromagnetics

The development of adaptive optics and phased array antennas have created new opportunities for active wavefront control in a variety of applications. Cancellation of atmospheric effects is already being used in scientific and military applications and more sophisticated uses of control are now being considered. One potential area for new work is in the area of active electromagnetic nulling for stealth applications.

To avoid detection and targeting, great strides have been achieved in reducing the radar cross section of military systems. Perhaps the best known advance has been the use of angularity and radar absorbing materials to minimize the detection of fighter aircraft. The narrow forward profile of the stealth fighter is very effective in minimizing radar reflection. However, the radar cross section increases whenever the pilot performs a banking turn. The radar absorbing materials used are expensive and are susceptible to moisture. multistatic radar systems, that can increasingly be built inexpensively, effectively track a stealth fighter. Additionally, engine exhaust infrared signatures represent serious system vulnerabilities.

Rather than use angularity to deflect incoming tracking or targeting radiation, a different approach is to develop inexpensive antenna arrays that will actively null incoming radiation. The use of ferrite structures in antennas could allow extremely rapidly change their radiating and receive properties. This would in turn allow arrays of such antennas to be used to intelligently respond to the surrounding electromagnetic environment by increasing the self-absorption of impinging radiation and by in turn radiating a field that will null further incoming radiation.

Several challenges must be overcome to achieve this goal, including distributed control system theory to define the currents applied to the radiating antenna to null the incoming radiation. The problem of field sensing and prediction in order to control its subsequent evolution is a significant mathematical and electrical engineering challenge. Advances in this area could have other applications in cellular phone communications systems, adaptive multistatic radar systems, and other directed energy systems.

Infrared exhaust signatures are another possible application area for active nulling. The complex flame dynamics in a gas turbine result in a statistically rich infrared field. A fundamental question is whether a finite array of infrared sensors and laser diodes could be used to sense, characterize and control this electromagnetic structure.

Molecular, Quantum and Nanoscale Systems

Control of molecular, quantum and nanoscale systems is receiving increased attention as our ability to sense and actuate systems at this scale improves. Recent progress in computational chemistry and physics has enabled the predictive simulation of nanoscale materials behavior and processing, in systems ranging from nanoparticles to semiconductor heterostructures to nanostructured bulk materials. With this physical understanding and a mathematical model, it is now possible for formulate optimization and control questions for nanoscale materials and systems. Applications include design of nanostructured materials, precision measurement,

and quantum information processing.

Macroscopic materials are well-described by their bulk properties, but as a structure's size shrinks to nanometers, bulk descriptions no longer capture the relevant physics. Surface effects become increasingly important and alter the electronic properties. These new properties may be exploited in a variety of engineering applications, from quantum dot lasers to ultra-hard coatings. A major challenge in exploiting these unique features is the ability to manufacture materials at the nanometer scale using high-throughput manufacturing processes. Improved first-principles models, new techniques for data-rich sensing and in-situ diagnostics, design of new actuation approaches, and algorithms for controlling microscale phenomena are required and the control community can be a major contributor to progress in this area.

There are many open questions in the control of phenomena at this scale. Brown and Rabitz [10] divide these into three categories: control law design, closed loop implementation, and identification of the system Hamiltonian. New results in controllability, optimizal control theory, adaptation and learning, and system identification are required to make progress in this area. What makes the problem difficult is the use of quantum wave interference as a mechanism for achieving prescribed control objectives such as the selective dissociation of a polyatomic molecule or the manipulation of wavepackets in semiconductors. Recent experimental successes (see [10] for more details) include cleaving and rearranging selected chemical bonds, control of fluorescence in polyatomic molecules and enhanced radiative emission in high harmonic generation.

Control of quantum systems also provides a new set of tools for understanding nature, as described in the vignette on quantum measurement and control (page 21).

Energy Systems

Control has always been a central element in the design of large scale energy systems. From their origins as single generators connected to associated loads, power systems began around 70 years ago to evolve into more broadly interconnected systems, motivated among other things by reliability (loads are now not totally dependent on any particular generator). Recent outage events have highlighted, however, that reliability is a more subtle question, and in fact global connectivity can lead to the multiplication of failures.

At same time, the industry is currently undergoing deregulation, which could easily lead to a loosening of control and a shortage of system information (even about neighbors!), elements which are key to the successful containment of failures. There is a significant risk that without a technological effort to improve the reliability of the network against failure, we can expect increased vulnerability of this fundamental infrastructure. One aspect of this effort concerns the design and management policies for the grid. This includes network expansions, or the deployment of new technological capabilities such as Flexible AC Transmission Systems (FACTS), and the decisions on load distribution to service the required demand.

Another area where fundamental research can have significant impact is real-time control for correct dynamic regulation and event-management for the contain-

ment of failures. There are increased linkages between computer and communications systems with the electric power grid, including the Internet and communications by the utilities to the in-orbit satellite network used for the Wide Area Measurement System (WAMS). This increased connectivity allows the possibility for significant local processing power at each generator, connected to a global data network providing information about the current state of the grid.

The technological challenges posed by such a system are multiple. A first challenge, even assuming a free and instantaneous flow of information, is to develop simulation capabilities for analysis of faults with the potential for large-scale cascading failures. Note that after a major outage it is possible to go back over the data and obtain a simple explanation. However, going beyond post mortems and into prediction will require truly revolutionary advances in basic simulation capability, indeed a completely new way of doing simulations of complex systems, going beyond both the traditional top-down approach that has dominated scientific computing, as well as the more recent bottom-up style of agent-based methods.

Secondly, distributed software protocols must be implemented for event-flow management and distributed simulation. A great challenge in this regard is that both the autonomy and the economic interests of the individual power companies must be protected. In other words, the distributed software system must be structured so that clear boundaries are drawn between serving the common good and protecting proprietary and economic information. Finally, note that with this network we are actually increasing interconnectedness, a driving factor of cascading failures. We should be wary of this process, which could lead to failures propagating through both networks. In particular, software systems should incorporate robustness against these kinds of events.

Progress in this area will require research progress in control theory and dynamical systems, along with the insight of statistical physics and new results in numerical methods, distributed software, computer networks and power systems. Each of these areas is relevant, and experts in each domain must expand their vision to tie to the other areas. As we evolve toward a “mega-infrastructure” that includes energy (including the electric grid, water, oil and gas pipelines), telecommunications, transportation, and electronic commerce, concepts from control will become even more important and will present significant challenges to the research community.

Chapter 4

Education and Outreach

Control education is an integral part of the community's activities and one of its most important mechanisms for transition and impact. In 1998, the National Science Foundation (NSF) and the IEEE Control Systems Society (CSS) jointly sponsored a workshop in control engineering education which made a number of recommendations for improving control education (see [1] and Appendix A). This section is based on the findings and recommendations of that report, and on discussions between Panel members and the control community. The Panel would particularly like to thank Jim Batterson for his contributions to this chapter.

4.1 The New Environment for Control Education

Control is traditionally taught within the various engineering disciplines that make use of its tools, allowing a tight coupling between the methods of control and their applications in a given domain. It is also taught almost exclusively within engineering departments, especially at the undergraduate level. Graduate courses are often shared between various departments and in some places are part of the curriculum in applied mathematics or operations research (particularly in regards to optimal control and stochastic systems). This approach has served the field well for many decades and has trained an exceptional community of control practitioners and researchers.

Increasingly, the modern control engineer is put in the role of being a systems engineer, responsible for linking together the many elements of a complex product or system. This requires not only a solid grounding in the framework and tools of control, but also the ability to understand the technical details of a wide variety of disciplines, including physics, chemistry, electronics, computer science, and operations research.

In addition, control is increasingly being applied outside of its traditional domains in aeronautics, chemical engineering, electrical engineering and mechanical engineering. Biologists are using ideas from control to model and analyze cells and animals; computer scientists are applying control to the design of routers and

embedded software; physicists are using control to measure and modify the behavior of quantum systems; and economists are exploring the applications of feedback to markets and commerce.

This change in the use of control presents a challenge to the community. In the United States, discipline boundaries within educational institutions are very strong and it is difficult to maintain a strong linkage between control educators and researchers across these boundaries. While the control community is large and prosperous, control is typically a small part of any given discussion on curriculum since these occur within the departments. Hence it is difficult to get the resources needed to make major changes in the control curriculum. In addition, many of the new applications of control are outside of the traditional disciplines that teach control and it is hard to justify developing courses that would appeal to this broader community and integrate those new courses into the curricula of those other disciplines (e.g., biology, physics, or medicine).

In order for the opportunities described elsewhere in this report to be realized, control education must be restructured to operate in this new environment. Several universities have begun to make changes in the way that control is taught and organized and these efforts provide some insights into how this restructuring might be done successfully.

Often the first step is establishing a cross-disciplinary research center, where there is a larger critical mass of control researchers. Examples include the Coordinated Science Laboratory (CSL) at the University of Illinois, Urbana-Champaign, the Center for Control Engineering and Computation (CCEC) at the University of California, Santa Barbara, and the Institute of Systems Research (ISR) at the University of Maryland. These centers coordinate research activities, organize workshops and seminars, and provide mechanisms for continuing interactions between control students and faculty in different departments.

A second step is the establishment of shared courses between the disciplines, often at the graduate level. These shared courses encourage a broader view of control since the students come from varying backgrounds. They also provide an opportunity for the larger control community at the university to establish active dialogs and provide a mechanism for sharing students and building joint research activities. Many U.S. universities have adopted this model, especially for theory oriented courses.

Finally, some schools have established a separate M.S. or Ph.D. program in control. These are common in Europe, but have been much less prevalent in the US, partly due to the traditional discipline structure around which most universities are organized. Examples in the U.S. include the Control and Dynamical Systems (CDS) program at Caltech and the Department of Systems Science and Mathematics (SSM) at Washington University. The advantage of a separate graduate program in control is that it gives the faculty better control over the curriculum and allows a less discipline-centric approach to control.

One other mechanism, popular in Europe but not yet established in the United States, is the creation of regional control alliances that build critical mass by linking together multiple universities in a geographic region. This mechanism is used very effectively, for example, in the Netherlands through the Dutch Institute of Systems

and Control (DISC).¹

4.2 Making Control More Accessible

Coupled with this new environment for control education is the clear need to make the basic principles of feedback and control known to a wider community. As the main recommendations of the Panel illustrate, many of the future opportunities for control are in new domains and the community must develop the educational programs required to train the next generation of researchers who will address these challenges.

A key need is for new books and courses that emphasize feedback concepts and the requisite mathematics, without requiring that students come from a traditional engineering background. As more students in biology, computer science, environmental science, physics, and other disciplines seek to learn and apply the methods of control, we must explore new ways of providing the background necessary to understand the basic concepts and apply some of the advanced tools that are available. Textbooks that are aimed at this more general audience could be developed and used in courses that target first year biology or computer science graduate students, who may have very little background in continuous mathematics beyond a sophomore course in scalar ODEs and linear algebra.

The following vignette describes one attempt to make control more accessible to a broader community of research scientists and engineers.

Vignette: CDS 110: Introduction to Control Concepts, Tools, and Theory (Kristi Morgansen and Richard Murray, Caltech)

The Control and Dynamical Systems Department at Caltech has recently undertaken a revision of its entry level graduate courses in control to make them accessible to students who do not have a traditional background in chemical, mechanical, or electrical engineering. The current course, CDS 110, is taken by senior undergraduates and first year graduate students from all areas of engineering, but has traditionally not been easily accessible to students in scientific disciplines, due to its heavy engineering slant. With the increased interest in control from these communities, it was decided to revise the course so that it could not only continue to serve its traditional role, but also provide an introduction to control concepts for first year graduate students in biology, computer science, environmental engineering, and physics.

The goal of the course is to provide an understanding of the principles of feedback and their use as a tool for altering the dynamics of systems and managing uncertainty. The main topics of the course are modeling, dynamics, interconnection, and robustness of feedback systems. On completion of the course, students are able to construct control-oriented models for typical systems found in engineering and the sciences, specify and describe performance for feedback systems, and analyze open loop and feedback

¹<http://www.disc.tudelft.nl>

behavior of such systems. Key themes throughout the course include input/output response, modeling and model reduction, linear versus nonlinear models, and local versus global behavior.

The updated version of the course has two “tracks”: a conceptual track and an analytical track. The conceptual track is geared toward students who want a basic understanding of feedback systems and the computational tools available for modeling, analyzing, and designing feedback systems. The analytical track is geared toward a more traditional engineering approach to the subject, including the use of tools from linear algebra, complex variables, and ordinary differential equations. Both tracks share the same lectures, but the supplemental reading and homework sets differ.

In addition to the main lectures, optional lectures are given by faculty from other disciplines whose research interests include control. Hideo Mabuchi (Physics) and Michael Dickinson (Biology) are two such lecturers and they provide examples of some applications of feedback to a variety of scientific and engineering applications. These lectures are used to emphasize how the concepts and tools are applied to real examples, drawn from areas such as aerospace, robotics, communications, physics, biology, and computer science.

The first iteration of the course, taught in 2001–02, succeeded in developing a set of conceptual lectures (given as the first lecture in the week) that introduced the main ideas of control with minimal mathematical background. Based on these lectures, students are able to use the tools of control (e.g., MATLAB and SIMULINK) and understand the results. Two additional lecture hours per week are used to provide the more traditional mathematical underpinnings of the subject and to derive the various results that are presented in the conceptual lectures.

In the second iteration of the course, to be taught in 2002–03, we intend to refine the lectures and put more effort into dividing the class into sections based on research interests. Individual lectures in the sections will then be used to build the necessary background (for example, providing a refresher on linear algebra and ODEs for biologists and computer scientists) or to provide additional perspectives (for example, linking transfer functions to Laplace transforms in a more formal manner).

In addition to changes in specific courses on control, universities could also integrate modules on dynamics and control into their undergraduate mathematics and science curricula. Any linear algebra course could be strengthened by the addition of a short lesson on linear systems, eigenvalues, and their physical interpretation and manipulation through feedback. Freshman physics could be enriched by extending lessons on mechanical oscillators to the problem of balancing an inverted pendulum or the stability of person riding a bicycle.

The control community also must continue to implement its tools in software, so that they are accessible to users of control technology. While this has already occurred in some areas of control (such as classical and modern linear control theory), there are very few general purpose software packages available for analysis and design of nonlinear, adaptive, and hybrid systems—and many of these are not available on general purpose platforms (such as MATLAB). These tools can be used to

allow non-experts to apply the most advanced techniques in the field without requiring that they first obtain a Ph.D. in control. Coupled with modeling and simulation tools, such as SIMULINK and Modelica, these packages will be particularly useful in teaching the principles of feedback by allowing exploration of relevant concepts in a variety of domains.

4.3 Broadening Control Education

In addition to changes in the curriculum designed to broaden the accessibility of control, it is important that control students also have a broader grasp of engineering, science, and mathematics. Modern control involves the development and implementation of a wide variety of very complex engineering systems and the control community has been a major source of training for people who embrace a systems perspective. The curriculum in control needs to reflect this role and provide students with the opportunity to develop the skills necessary for modern engineering and research activities.

At the same time, the volume of work in control is enormous and so effort must be placed on unifying the existing knowledge base into a more compact form. There is a need for new books that systematically introduce a wide range of control techniques in an effective manner. This will be a major undertaking, but is required if future students of control are to receive a concise but thorough grounding in the fundamental principles underlying control, so that they can continue to extend the research frontier beyond its current boundary.

Increasingly, control engineers are playing the role of “systems integrator” in large engineering projects. This occurs in part because they bring systems insight that is required for successful operation of a complex engineering product, but also because control is often the glue that ties together the components of the system (often in the form of embedded control software). Unfortunately, most control curricula do not emphasize the types of leadership and communications skills that are critical for success in these environments.

A related aspect of this is strengthening the skills required for working in teams. All modern systems design is done in interdisciplinary teams and it requires certain skills to understand how to effectively interact with domain experts from a wide variety of disciplines. Project courses are an effective mechanism for developing this type of insight and these should be more aggressively incorporated into control curricula at both undergraduate and graduate levels. Another effective mechanism is participation in national competitions where control tools are required, such as RoboCup² and FIRST³.

It is also important that control students be provided with a balance between theory, applications, and computation. In particular, it is essential that control students build a deep domain knowledge in one or more disciplines, so that they understand how this knowledge interacts with the control methodology. Independent of the specific domain chosen, this approach provides a context for understanding

²<http://www.robocup.org>

³<http://www.usfirst.org>

other engineering domains and developing control practices and tools that bridge application areas.

Experiments continue to form an important part of a control education and projects should form an integral part of the curriculum for both undergraduate and graduate students. Shared laboratories within individual colleges or universities as well as laboratories shared among different universities could be used to implement this (with additional benefits in building cross-disciplinary and cross-university interactions). New experiments should be developed that explore the future frontiers of control, including increased use of computing, communications and networking, as well as exploration of control in novel application domains.

4.4 The Opportunities in K-12 Math and Science Education

Much as computer literacy has become commonplace in our K-12 curriculum, an understanding of the requirements, limits, and capabilities of control should become part of every scientifically literate citizen's knowledge. Whether it is understanding why you should not pump anti-skid brakes or why you need to complete a regimen of antibiotics through the final pills even after symptoms disappear, an understanding of dynamics and control is essential. The development of inexpensive microprocessors, high-level computer languages, and GUIs has made the development of test apparatus and small laboratories for rudimentary control experiments and demonstrations available within the budgets of all school districts. The U.S. National Science Foundation recognizes the importance of its funded programs impacting the general public through its "Criterion 2" (Broader Impacts) in the evaluation of all submitted proposals. Because of the broad applications of control to the public good and standards of living, it is important to develop a curriculum for inclusion in pre-college (K-12) education.

Currently, mathematics, science, and computer technology are taught in separate departments in the vast majority of K-12 curricula. Even sciences are compartmentalized at many schools. As at universities, the multidisciplinary nature of control is very much antithetical to that traditional thinking and structure in K-12 education. However, there is some evidence of advances toward application and integration of mathematics with science. The Consortium for Mathematics and Its Applications (COMAP) is one example.⁴ Indeed, the leveraging of efforts with COMAP could prove fruitful and the control community could work with COMAP to enhance the current books and curricula that have been developed by that consortium over the past two decades.

In the control arena, simple experiments involving governors, thermostats, and "see-saws" can be performed as early as schools to illustrate the basic concepts of control. As mathematical sophistication increases through middle school and high school, quantitative analysis can be added and experimentally verified. Some schools are beginning to teach calculus in the junior year and so a post-calculus course in applied mathematics of differential equations and dynamical systems could

⁴<http://www.comap.com/>

be created bridging chemistry, physics, biology, and mathematics.

Complementary to the development of educational materials and experiments, it is also important to provide teachers with the opportunities to learn more about control. As an example of how this could be done, NASA Langley Research Center sponsored a program for teachers under the auspices of the HPCCP (High Performance Computing and Communications Program) several years ago. In this program teachers from six school districts spent 8 weeks learning the state of the art in computer hardware and software for engineering and science. Most days were spent with new material delivered in a lecture or laboratory environment in the morning with a “homework” laboratory in the afternoons. Teachers were paid a fellowship that approximated the per diem rate of entry-level teachers. This type of residential environment allowed for a total immersion in the material. In addition to becoming familiar with research-grade hardware and software and the Internet, the participants formed partnerships with one another that promoted continued collaboration throughout the coming academic years.

There are numerous curriculum development and general education meetings and conferences throughout the country each year. In particular, most states have an active association of school boards and there is a National School Boards Association. A presentation at these meetings would communicate directly with the policy and decision makers. Such a presentation would have to be tailored for the layman but might produce a pull to match a push from one of the ideas above.

4.5 Other Opportunities and Trends

In addition to the specific opportunities for education and outreach described above, there are many other possible mechanisms to help expand the understanding and use of control tools.

Popular Books and Articles

In September of 1952, *Scientific American* published an entire issue dedicated to Automatic Control [31]. The issue highlighted the role that control was playing in the new advancements of the time, particularly in manufacturing. The introduction of cruise control (originally called Autopilot) a few years later provided direct experience with the main principles of feedback.

Since that time, control has become less and less visible to the general public, perhaps in part because of its success. Individuals interact with control systems and feedback many times every day, from the electronic amplifiers, tuners, and filters in television and radio, to congestion control algorithms that enable smooth Internet communications, to flight control systems for commercial aircraft. Yet most people are unaware of control as a discipline. Other fields, such as artificial intelligence, robotics, and computer science have often been given credit for ideas whose origins lie within the control community.

There is a great need to better educate the public on the successes and opportunities for control. This public awareness is increasingly important in the face of decisions that will need to be made by government funding agencies about support

for specific areas of research.

The use of any number of popular outlets for communication can reach this group. Many local newspapers now have a science page or section on a weekly basis. The development of a popular level series of articles on dynamics and control could be prepared for these pages. The New York Times publishes a science section every Tuesday. A series of articles could be developed for this section spanning several weeks. A number of science museums have been developed across the nation in recent years. Many of these museums are allied through professional associations. The development of interactive dynamics and control displays for these museums would be beneficial to the museum by giving them a new exhibit and the displays reach the entire age range of the public from children through adults.

Multimedia Tools

There is an increasing need for educational materials that can be used in a variety of contexts for communicating the basic ideas behind control. One possible mechanism is to develop a multimedia CDROM that would include materials on the history and concepts of control, as well as tutorial material on specific topics and public domain software tools for control analysis and design.

The fluid mechanics community has recently developed such a multimedia CDROM that can be used as a supplement to traditional courses in fluid mechanics [16]. It contains historical accounts of fluid mechanics, videos and animations of important concepts in fluids, and detailed descriptions of fundamental phenomena. It can be purchased through university bookstores or online from Amazon.com.

One initial activity in developing such tools for control has been made by Wilson J. Rugh at Johns Hopkins University, who has created a series of interactive demonstrations of basic concepts of control that can be executed over the web.⁵ Modules include Fourier analysis, convolution, the sampling theorem, and elementary control systems. One of the most sophisticated tools demonstrates robust stabilization, including the ability to specify an uncertainty weight by move poles and zeros of the weighting transfer function with the mouse. A controller can then be designed by dragging the compensator poles and zeros to achieve robust, closed loop stability.

Software

One of the success stories of control is the wide availability of commercial software for modeling, analyzing, designing, and implementing control systems. The Controls Toolbox in MATLAB provides the basic tools of classical and modern control and many other toolboxes are available for more implementing more specialized theory. These toolboxes are used throughout academia, government, and industry and give students, researchers, and practitioners access to powerful tools that have been carefully designed and tested.

Despite the impressive current state of the art, much of this software is restricted to a very small class of the systems typically encouraged in control and

⁵<http://www.jhu.edu/~signals/>

there are many gaps that will need to be filled to enable the types of applications described in the previous chapter. One area where substantial progress has been made recently is in modeling tools, where there are several software packages available for modeling, simulation, and analysis of large-scale, complex systems. One such example is *Modelica*,⁶ which provides an object oriented language for describing complex physical systems. Modelica is particularly noteworthy because it was designed to model systems with algebraic constraints, allowing a much richer class of systems to be represented.

Additional tools are needed for controls-oriented modeling, analysis, and synthesis of nonlinear and hybrid systems, particularly those that have a strong interaction with information rich systems, where good scaling properties are required. As yet, there is not a standard representational framework for such systems (beyond symbolic representations) and hence software tools for nonlinear or hybrid analysis are much less used than those for linear systems. One of the main issues here is to capture the relevant dynamics in a framework that is amenable to computation. Analysis and synthesis must be able to handle systems containing table lookups, logical elements, time delays, and models for computation and communication elements.

The payoff for investing in the development of such tools is clear: it brings the advanced theoretical techniques that are developed within the community to the people who can most use those results.

Interaction with Industry

Interaction with industry is an important component of any engineering research or educational activity. The control community has a strong history of impact on many important problems and industry involvement will be critical for the eventual success of the future directions described in this report. This could occur through cooperative Ph.D. programs where industrial researchers are supported half by companies and half by universities to pursue Ph D's (full-time), with the benefits of bringing more understanding of real-world problems to the university and transferring the latest developments back to industry. In addition, industry leaders and executives from the control community should continue to interact with the community and help communicate the needs of their constituencies.

The NSF/CSS workshop also recognized the important role that industry plays and recommended that educators and funding organizations

encourage the development of WWW-based initiatives for technical information dissemination to industrial users of control systems and encourage the transfer of practical industrial experience to the classroom [1].

The further recommended that cooperative efforts between academia and industry, especially in terms of educational matters, be significantly expanded.

The International Federation of Automatic Control (IFAC) is creating a collection of IFAC Professional Briefs. These Professional Briefs are aimed at a readership of general professional control engineers (industrial and academic), rather

⁶<http://www.modelica.org>

than specialist researchers. The briefs provide an introduction and overview of a “hot topic,” illustrative results, and a sketch of the underlying theory, with special attention given to providing information sources such as useful Internet sites, books, papers, etc. Eight titles have been selected to launch the Professional Briefs series:

- Computer Controlled Systems
- PID Auto-Tuning
- Control of Biotechnological Processes
- Control Busses and Standards
- Physical-Based Modeling of Mechatronic Systems
- Genetic Algorithms in Control Systems Engineering
- Low Cost Automation in Manufacturing
- Engineering Dependable Industrial Real-Time Software.

Chapter 5

Recommendations

Control continues to be a field rich in opportunities. In order to realize these opportunities, it is important that the next generation of control researchers receive the support required to develop new tools and techniques, explore new application areas, and reach out to new audiences. Toward this end, the Panel developed a list of five major recommendations for accelerating the impact of control.

5.1 Integrated Control, Computation, Communications

Cheap and ubiquitous sensing, communications, and computation will be a major enabler for new applications of control to large-scale, complex systems. Research in control over networks, control of networks, and design of safety critical, large-scale interconnected systems will generate many new research issues and theoretical challenges. A key feature of these systems is their robust yet fragile behavior, with cascade failures leading to large disruptions in performance.

A significant challenge will be to bring together the diverse research communities in control, computer science, and communications in order to build the unified theory required to make progress in this area. Joint research by these communities will be much more team-based and will likely involve groups of domain experts working on common problems, in addition to individual investigator-based projects.

To realize the opportunities in this area, the Panel recommends that government agencies and the control community

Substantially increase research aimed at the *integration* of control, computer science, communications, and networking.

In the United States, the Department of Defense has already made substantial investment in this area through the Multidisciplinary University Research Initiative (MURI) program and this trend should be continued. It will be important to create larger, multi-disciplinary centers that join control, computer science, and

communications and to train engineers and researchers who are knowledgeable in these areas.

Industry involvement will be critical for the eventual success of this integrated effort and universities should begin to seek partnerships with relevant companies. Examples include manufacturers of air traffic control hardware and software, and manufacturers of networking equipment.

The benefits of increased research in integrated control, communications, and computing will be seen in our transportation systems (air, automotive, and rail), our communications networks (wired, wireless, and cellular), and enterprise-wide operations and supply networks (electrical power, manufacturing, service and repair).

5.2 Control of Complex Decision Systems

The role of logic and decision making in control systems is becoming an increasingly large portion of modern control systems. This decision making includes not only traditional logical branching based on system conditions, but higher levels of abstract reasoning using high level languages. These problems have traditionally been in the domain of the artificial intelligence (AI) community, but the increasing role of dynamics, robustness, and interconnection in many applications points to a clear need for participation by the control community as well.

A parallel trend is the use of control in very large scale systems, such as logistics and supply chains for entire enterprises. These systems involve decision making for very large, very heterogeneous systems where new protocols are required for determining resource allocations in the face of an uncertain future. Although models will be central to analyzing and designing such systems, these models (and the subsequent control mechanisms) must be scalable to *very* large systems, with millions of elements that are themselves as complicated as the systems we currently control on a routine basis.

To tackle these problems, the Panel recommends that government agencies and the control community

Substantially increase research in control at higher levels of decision making, moving toward enterprise level systems.

The extension of control beyond its traditional roots in differential equations is an area that the control community has been involved in for many years and it is clear that some new ideas are needed. Effective frameworks for analyzing and designing systems of this form have not yet been fully developed and the control community must get involved in this class of applications in order to understand how to formulate the problem.

A useful technique may be the development of testbeds to explore new ideas. In the military arena, these testbeds could consist of collections of unmanned vehicles (air, land, sea and space), operating in conjunction with human partners and adversaries. In the commercial sector, service robots and personal assistants may be a fruitful area for exploration. And in a university setting, the emergence of robotic competitions is an interesting trend that control researchers should explore as a

mechanism for developing new paradigms and tools. In all of these cases, stronger links with the AI community should be explored, since that community is currently at the forefront of many of these applications.

The benefits of research in this area include replacing *ad hoc* design methods by systematic techniques to develop much more reliable and maintainable decision systems. It will also lead to more efficient and autonomous enterprise-wide systems and, in the military domain, provide new alternatives for defense that minimize the risk of human life.

5.3 High-Risk, Long-Range Applications of Control

The potential application areas for control are increasing rapidly as advances in science and technology develop new understanding of the importance of feedback, and new sensors and actuators allow manipulation of heretofore unimagined detail. To discover and exploit opportunities in these new domains, experts in control must actively participate in new areas of research outside of their traditional roots. At the same time, we must find ways to educate domain experts about control, to allow a fuller dialog and to accelerate the uses of control across the enormous number of possible applications.

In addition, many applications will require new paradigms for thinking about control. For example, our traditional notions of signals that encode information through amplitude and phase relationships may need to be extended to allow the study of systems where pulse trains or bio-chemical “signals” are used to trace information.

One of the opportunities in many of these domains is to export (and expand) the framework for systems-oriented modeling that has been developed in control. The tools that have been developed for aggregation and hierarchical modeling can be important in many systems where complex phenomena must be understood. The tools in control are among the most sophisticated available, particularly with respect to uncertainty management.

To realize some of these opportunities, the Panel recommends that government agencies and the control community

Explore high-risk, long-range applications of control to areas such as nanotechnology, quantum mechanics, electromagnetics, biology, and environmental science.

A challenge in exploring new areas is that experts in two (or more) fields must come together, which is often difficult under mainly discipline-based funding constructs. There are a variety of mechanisms that might be used to do this, including dual-investigator funding through control programs that pay for biologists, physicists, and others to work on problems side-by-side with control researchers. Similarly, funding agencies should broaden the funding of science and technology to include funding of the control community through domain-specific programs.

Another need is to establish “meeting places” where control researchers can join with new communities and each can develop an understanding of the principles and tools of the other. This could include focused workshops of a week or more to

explore control applications in new domains or 4–6 week short courses on control that are tuned to a specific applications area, with tutorials in that application area as well.

At universities, new materials are needed to teach non-experts who want to learn about control. Universities should also consider dual appointments between science and engineering departments that recognize the broad nature of control and the need for control to not be confined to a single disciplinary area. Cross-disciplinary centers (such as the CCEC at Santa Barbara) and programs in control (such as the CDS program at Caltech) are natural locations for joint appointments and can act as a catalyst for getting into new areas of control by attracting funding and students outside of traditional disciplines.

There are many areas ripe for the application of control and increased activity in new domains will accelerate the use of control and enable advances in many different domains. In many of these new application areas, the systems approach championed by the control community has yet to be applied, but it will be required for eventual engineering applications. Perhaps more important, control has the opportunity to revolutionize other fields, especially those where the systems are complicated and difficult to understand. Of course, these problems are extremely hard and many previous attempts have not always been successful, but the opportunities are great and we must continue to strive to move forward.

5.4 Support for Theory and Interaction with Mathematics

A core strength of control has been its respect for and effective use of theory, as well as contributions to mathematics driven by control problems. Rigor is a trademark of the community and one that has been key to many of its successes. Continued interaction with mathematics and support for theory is even more important as the applications for control become more complex and more diverse.

An ongoing need is making the existing knowledge base more compact so that the field can continue to grow. Integrating previous results and providing a more unified structure for understanding and applying those results is necessary in any field and has happened many times in the history of control. This process must be continuously pursued and requires steady support for theoreticians working on solidifying the foundations of control. It is also needed for control experts to expand the applications base by having the appropriate level of abstraction to identify new applications of existing theory.

To ensure the continued health of the field, the Panel recommends that the community and funding agencies

Maintain support for theory and interaction with mathematics, broadly interpreted.

Some possible areas of interaction include dynamical systems, graph theory, combinatorics, complexity theory, queuing theory, statistics, etc. Additional perspectives on the interaction of control and mathematics can be found in a recent survey article

by Brockett [9].

A key need is to identify and provide funding mechanisms for people to work on core theory. The proliferation of multi-disciplinary, multi-university programs threatens this base of individual investigators who are working on the theory that is required for future success. It is important to leave room for theorists on these applications-oriented projects and to better articulate the successes of the past so that support for the theory is appreciated. Program managers should support a balanced portfolio of applications, computation, and theory, with clear articulation of the importance of long term, theoretical results.

The linkage of control with mathematics should also be increased, perhaps through new centers and programs. Funding agencies should consider funding national institutes for control science that would engage the mathematics community, and existing institutes in mathematics should be encouraged to sponsor year-long programs on control, dynamics, and systems.

The benefits of this investment in theory will be a systematic design methodology for building complex systems and rigorous training for the next generation of researchers and engineers.

5.5 New Approaches to Education and Outreach

As many of the recommendations above indicate, applications of control are expanding and this is placing new demands on education. The community must continue to unify and compact the knowledge base by integrating materials and frameworks from the past 40 years. As important, material must be made more accessible to a broad range of potential users, well beyond the traditional base of engineering science students and practitioners. This includes new uses of control by computer scientists, biologists, physicists, and medical researchers. The technical background of these constituencies is often very different than traditional engineering disciplines and will require new approaches to education.

The Panel believes that control principles are now a required part of any educated scientist's or engineer's background and we recommend that the community and funding agencies

Invest in new approaches to education and outreach for the dissemination of control concepts and tools to non-traditional audiences.

As a first step toward implementing this recommendation, new courses and textbooks should be developed for both experts and non-experts. Control should also be made a *required* part of engineering and science curricula at major universities, including not only mechanical, electrical, chemical, and aerospace engineering, but also computer science, applied physics, and bio-engineering. It is also important that these courses emphasize the *principles* of control rather than simply providing tools that can be used in a given domain.

An important element of education and outreach is the continued use of experiments and the development of new laboratories and software tools. These are much easier to do than ever before and also more important. Laboratories and

software tools should be integrated into the curriculum, including moving beyond their current use in introductory control courses to increased use in advanced (graduate) course work. The importance of software cannot be overemphasized, both in terms of design tools (e. g., MATLAB toolboxes) and implementation (real-time algorithms).

Increased interaction with industry in education is another important step. This could occur through cooperative Ph.D. programs where industrial researchers are supported half by companies and half by universities to pursue Ph.D.'s (full-time), with the benefits of bringing more understanding of real-world problems to the university and transferring the latest developments back to industry. In addition, industry leaders and executives from the control community should continue to interact with the community and help communicate the needs of their constituencies.

Additional steps to be taken include the development of new teaching materials that can be used to broadly educate the public about control. This might include chapters on control in high school textbooks in biology, mathematics, and physics or a multimedia CD that describes the history, principles, successes, and tools for control. Popular books that explain the principles of feedback, or perhaps a "cartoon book" on control should be considered. The upcoming IFAC Professional Briefs for use in industry are also an important avenue for education.

The benefits of reaching out to broader communities will be an increased awareness of the usefulness of control, and acceleration of the benefits of control through broader use of its principles and tools. The use of rigorous design principles will result in safer systems, shorter development times, and more transparent understanding of key systems issues.

5.6 Concluding Remarks

The field of control has a rich history and a strong record of success and impact in commercial, military, and scientific applications. The tradition of rigorous use of mathematics combined with strong interaction with applications has produced a set of tools that are used in a wide variety of technologies. The opportunities for future impact are even richer than those of the past, and the field is well positioned to expand its tools to apply to new areas and applications.

The pervasiveness of communications, computing and sensing will enable many new applications of control but will also require a substantial expansion of the current theory and tools. The control community must embrace new, information rich applications and generalize existing concepts to apply to systems at higher levels of decisions making. Combined with new, long-range areas that are opening up to control techniques, the next decade promises to be a fruitful one for the field.

The payoffs for investment in control research are substantial. They include the successful development of systems that operate reliably, efficiently, and robustly; new materials and devices that are made possible through advanced control of manufacturing processes; and increased understanding of physical and biological systems through the use of control principles. Perhaps most important is the continued de-

velopment of individuals who embrace a systems perspective and provide technical leadership in modeling, analysis, design and testing of complex engineering systems.

Appendix A

NSF/CSS Workshop on New Directions in Control Engineering Education [1]

The National Science Foundation (NSF) and the IEEE Control Systems Society (CSS) held a workshop in October 1998 to identify the future needs in control systems education. The executive summary of the report is presented here. The full report is available from the CDS Panel homepage.

Executive Summary

The field of control systems science and engineering is entering a golden age of unprecedented growth and opportunity that will likely dwarf the advancements stimulated by the space program of the 1960s. These opportunities for growth are being spurred by enormous advances in computer technology, material science, sensor and actuator technology, as well as in the theoretical foundations of dynamical systems and control. Many of the opportunities for future growth are at the boundaries of traditional disciplines, particularly at the boundary of computer science with other engineering disciplines. Control systems technology is the cornerstone of the new automation revolution occurring in such diverse areas as household appliances, consumer electronics, automotive and aerospace systems, manufacturing systems, chemical processes, civil and environmental systems, transportation systems, and even biological, economic, and medical systems.

The needs of industry for well trained control systems scientists and engineers are changing, due to marketplace pressures and advances in technology. Future generations of engineering students will have to be broadly educated to cope with cross-disciplinary applications and rapidly changing technology. At the same time, the backgrounds of students are changing. Many come from nontraditional backgrounds; they often are less well prepared in mathematics and the sciences while being better prepared to work with modern computing technologies. The time is thus ripe for major renovations in control and systems engineering education.

To address these emerging challenges and opportunities, the IEEE Control Systems Society initiated the idea of holding a workshop that would bring together leading control systems researchers to identify the future needs in control systems

education. The workshop was held on the campus of the University of Illinois at Urbana-Champaign, October 2–3, 1998. It attracted sixty-eight participants.

This report summarizes the major conclusions and recommendations that emerged from the workshop. A slightly modified version of the main body of this report will also appear in the October, 1999, issue of the IEEE Control Systems Magazine. These recommendations, which cover a broad spectrum of educational issues, are addressed to several constituencies, including the National Science Foundation, control systems professional organizations, and control systems researchers and educators.

1. General Recommendation

- 1 *Enhance cooperation among various control organizations and control disciplines throughout the world to give attention to control systems education issues and to increase the general awareness of the importance of control systems technology in society.*

Mechanisms to accomplish this include joint sponsorship of conferences, workshops, conference sessions, and publications devoted to control systems education as well as the development of books, websites, videotapes, and so on, devoted to the promotion of control systems technology.

2. Additional Recommendations

- 2 *Provide practical experience in control systems engineering to freshmen to stimulate future interest and to introduce fundamental notions like feedback and the systems approach to engineering.*

This can be accomplished by incorporating modules and/or projects that involve principles of control systems into freshmen courses that already exist in many engineering schools and colleges.

- 3 *Encourage the development of new courses and course materials that will significantly broaden the standard first introductory control systems course at the undergraduate level.*

Such new courses would be accessible to all third year engineering students and would deal with fundamental principles of system modeling, planning, design, optimization, hardware and software implementation, computer aided control systems design and simulation, and systems performance evaluation. Equally important, such courses would stress the fundamental applications and importance of feedback control as well as the limits of feedback, and would provide a bridge between control systems engineering and other branches of engineering that benefit from systems engineering concepts such as networks and communications, biomedical engineering, computer science, economics. etc.

- 4 *Develop follow on courses at the undergraduate level that provide the necessary breadth and depth to prepare students both for industrial careers and for graduate studies in systems and control.*

Advanced courses in both traditional control methodologies, like digital control, and courses treating innovative control applications should be available to undergraduate students in order to convey the excitement of control systems engineering while still providing the fundamentals needed in practice.

5 Promote control systems laboratory development, especially the concept of shared laboratories, and make experimental projects an integral part of control education for all students, including graduate students.

Mechanisms to accomplish this include increased support for the development of hands-on control systems laboratories, as well as the development of benchmark control systems examples that are accessible via the Internet. Shared laboratories within individual colleges or universities as well as shared laboratories among different universities makes more efficient use of resources, increases exposure of students to the multidisciplinary nature of control, and promotes the interaction of faculty and students across disciplines.

The promotion of laboratory development also includes mechanisms for continued support. Too often, laboratories are developed and then abandoned after a few years because faculty do not have time or funds for continued support. It is equally important, therefore, to provide continuity of support for periodic hardware and software upgrades, maintenance, and the development of new experiments.

The National Science Foundation and IEEE Control Systems Society can also help realize this goal by developing workshops and short courses for laboratory development and instruction to promote interaction and sharing of laboratory development experiences among faculty from different universities.

6 Emphasize the integration of control systems education and research at all levels of instruction.

The National Science Foundation program, Research Experiences for Undergraduates, exemplifies an excellent mechanism to accomplish this at the undergraduate level and should be continued. Sponsorship of student competitions in control is another such mechanism that should be encouraged. At the graduate level control educators should take advantage of National Science Foundation programs such as the Integrative Graduate Education and Research Training Program (IGERT) and the Course, Curriculum, and Laboratory Improvement Program (CCLI).

7 Improve information exchange by developing a centralized Internet repository for educational materials.

These materials should include tutorials, exercises, case studies, examples, and histories, as well as laboratory exercises, software, manuals, etc. The IEEE Control Systems Society can play a leadership role in the development of such a repository by coordinating the efforts among various public and private agencies.

8 Promote the development of a set of standards for Internet based control systems materials and identify pricing mechanisms to provide financial compensation to Internet laboratory providers and educational materials providers.

A mechanism to accomplish this could be a National Science Foundation sponsored workshop devoted to Internet standards for control education materials and pricing.

9 Develop WWW-based peer reviewed electronic journal on control education and laboratory development.

Control systems professional organizations can play leadership roles, perhaps working with the American Society of Engineering Education (ASEE) to accomplish this goal.

10 Encourage the development of initiatives for technical information dissemination to industrial users of control systems and encourage the transfer of practical industrial experience to the classroom.

Mechanisms to accomplish this include special issues of journals and magazines devoted to industrial applications of control, programs to bring speakers from industry to the classroom, and programs that allow university faculty to spend extended periods of time in industry.

Bibliography

- [1] P. Antsaklis, T. Basar, R. DeCarlo, N. H. McClamroch, M. Spong, and S. Yurkovich, editors. *NSF/CSS Workshop on New Directions in Control Engineering Education*. National Science Foundation and IEEE Control Systems Society, 1998. Available at <http://robot0.ge.uiuc.edu/~spong/workshop>.
- [2] M. Athans, J. J. S. Sentieiro, and M. A. Dahleh. Workshop on future directions in systems and control theory. Cascais, Portugal, June 2000.
- [3] S. Banda, J. C. Doyle, R. M. Murray, J. Paduano, J. Speyer, and G. Stein. Research needs in dynamics and control for uninhabited aerial vehicles. Panel Report, November 1997. <http://www.cds.caltech.edu/~murray/notes/uav-nov97.html>.
- [4] W. L. Bargar, A. Bauer, and M. Borner. Primary and revision total hip replacement using the Robodoc (R) system. *Clinical Orthopaedics and Related Research*, (354):82–91, 1998.
- [5] N. Barkai and S. Leibler. Robustness in simple biochemical networks. *Nature*, 387(6636):913–917, 1997.
- [6] 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.
- [7] A. Becskei and L. Serrano. Engineering stability in gene networks by autoregulation. *Nature*, 405(6786):590–593, 2000.
- [8] T. R. Bewley. Flow control: new challenges for a new Renaissance. *Progress in Aerospace Sciences*, 37:21–58, 2001.
- [9] R. W. Brockett. New issues in the mathematics of control. In B. Engquist and W. Schmid an W. Schmid, editors, *Mathematics Unlimited—2001 and Beyond*, pages 189–220. Springer Verlag, 2000.
- [10] E. Brown and H. Rabitz. Some mathematical and algorithmic challenges in the control of quantum dynamics phenomena. *Journal of Chemical Physics*, 2002. To appear.

-
- [11] European Commission. Workshop on future and emerging control systems, November 2000. Available at ftp://ftp.cordis.lu/pub/ist/docs/ka4/report_controlws.pdf.
- [12] M. H. Dickinson, C. T. Farley, R. J. Full, M. A. R. Koehl, R. Kram, and S. Lehman. How animals move: An integrative view. *Science*, 288:100–106, 2000.
- [13] W. H. Fleming, editor. *Future Directions in Control Theory: A Mathematical Perspective*. SIAM, 1988.
- [14] D. Hanahan and R. A. Weinberg. The hallmarks of cancer. *Cell*, 100:57–70, 2000.
- [15] M. B. Hoagland and B. Dodson. *The Way Life Works*. Times Books, 1995.
- [16] G. M. Homsy, Aref H., K. Breuer, S. Hocqreb, and J. R. Koseff. *Multimedia Fluid Mechanics*. Cambridge University Press, 2000.
- [17] T. P. Hughes. *Elmer Sperry Inventor and Engineer*. The John Hopkins University Press, London, 1993.
- [18] V. Jacobson. Notes on using RED for queue management and congestion avoidance. Presentation, North American Network Operators’ Group meeting, June 1998. Available at <http://www.nanog.org/mtg-9806/ppt/vj-nanog-red.pdf>.
- [19] F. Rowsone Jr. What it’s like to drive an auto-pilot car. *Popular Science Monthly*, April 1958. Available at <http://www.imperialclub.com/ImFormativeArticles/1958AutoPilot>.
- [20] P. R. Kumar. New technological vistas for systems and control: The example of wireless networks. *Control Systems Magazine*, 21(1):24–37, 2001.
- [21] W. S. Levine. *Control System Fundamentals*. CRC Press, 1999.
- [22] A. H. Levis, S. I. Marcus, W. R. Perkins, P. Kokotovic, M. Athans, R. W. Brockett, and A. S. Willsky. Challenges to control: A collective view. *IEEE Transactions on Automatic Control*, AC-32(4):275–285, 1987.
- [23] J. Marescaux, J. Leroy, M. Gagner, F. Rubino, D. Mutter, M. Vix, S. E. Butner, and M. K. Smith. Transatlantic robot-assisted telesurgery. *Nature*, 413:379–380, 2001.
- [24] H. H. McAdamsz and L. Shapiro. Circuit simulation of genetic networks. *Science*, 269(5224):650–656, 1995.
- [25] M. W. McFarland, editor. *The Papers of Wilbur and Orville Wright*. McGraw Hill, 1953.

-
- [26] F. W. Mohr, V. Falk, A. Diegeler, and R. Autschbach. Computer-enhanced coronary artery bypass surgery. *Journal of Thoracic and Cardiovascular Surgery*, 117(6):1212–1214, 1999.
- [27] National Science Foundation. NSF workshop for high school teachers of mathematics and science. American Control Conference, June 2000. Available at http://www.math.ukans.edu/ksacg/nsf_hswkshop.html.
- [28] H. Nijmeijer and J. M. Schumacher. Four decades of mathematical system theory. In J. W. Polderman and H. L. Trentelman, editors, *The Mathematics of Systems and Control: From Intelligent Control to Behavioral Systems*, pages 73–83. Univ. of Groningen, 1999.
- [29] P. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, Control and Dynamical Systems, 2000.
- [30] W. F. Powers and P. R. Nicastrì. Automotive vehicle control challenges in the 21st century. *Control Engineering Practice*, 8:605–618, 2000.
- [31] Scientific American. Special issue on Automatic Control, September 1952.
- [32] H. S. Tsien. *Engineering Cybernetics*. McGraw-Hill, 1954.
- [33] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. John Wiley, 1948.
- [34] N. Wiener. *The Human Use of Human Beings: Cybernetics and Society*. Houghton Mifflin, 1950.
- [35] T. M. Yi, Y. Huang, M. I. Simon MI, 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.