Control in an Information Rich World

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

Richard M. Murray (chair)

DRAFT: v1.1, Sat Mar 31 11:53:24 2001

Abstract

The field of *automatic control* provides the principles and methods used to design systems that maintain desirable performance by automatically adapting to changes in the environment. Aircraft auto-landing systems and sophisticated guided weapons, for example, automatically adjust control surfaces to adapt to changing winds and changes in weight and loads, to accurately and reliably follow prescribed trajectories. Routing protocols in communication networks adapt to changing congestion, queue levels, and network traffic.

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. Over the last forty years, the field has seen huge advances, leveraging technology improvements with breakthroughs in the underlying principles and mathematics. Automatic feedback control systems now play critical roles in many fields, including manufacturing, electronics, communications, transportation, computers and networks, and many military systems. However, the enabling role automatic feedback control technology plays is sometimes unseen: the control methods are hidden in the computer code that carries out the au-

tomatic adjustments, or obscured by the difficult mathematics underlying it.

As we begin the 21st century, the opportunities for, and use of, automatic control principles and methods is exploding. Increasingly, computing will be ubiquitous, more and more devices will include embedded, cheap, high-performance processors and sensors, and wireless networks will greatly enhance information exchange. This will make possible the development of machines with a degree of intelligence and reactivity that will change everyone's life, both in terms of the goods available and the environment in which we live.

Of course new developments will require a significant expansion of the basic tool sets of automatic control. The complexity of the control ideas involved in the operation of the internet, autonomous systems, or an enterprise-wide supply chain system are on the boundary of what can be done with available methods, so new developments must be vigorously pursued.

The purpose of this report is to spell out some of the exciting prospects for the field in the current and future technological environment and to explain the critical role we expect it to play over the next decade.

Panel Membership

Richard M. Murray (chair)
California Institute of Technology

Karl Astrom
Lund Institute of Technology

Stephen Boyd
Stanford University

Siva Banda
Air Force Research Laboratory

Roger Brockett
Harvard

John Burns
Virginia Tech

Munzer Dahleh
Massachusetts Institute of Technology

John Doyle
California Institute of Technology

John Guckenheimer
Cornell University

Charles Holland
Department of Defense

Pramod Khargonekar
University of Michigan

P. R. Kumar
University of Illinois

P. S. Krishnaprasad
University of Maryland

Greg McRae
Massachusetts Institute of Technology

Jerrold Marsden
California Institute of Technology

George Meyer
NASA

William Powers
Ford

Gunter Stein
Honeywell

Pravin Varaiya
University of California Berkeley

Contents

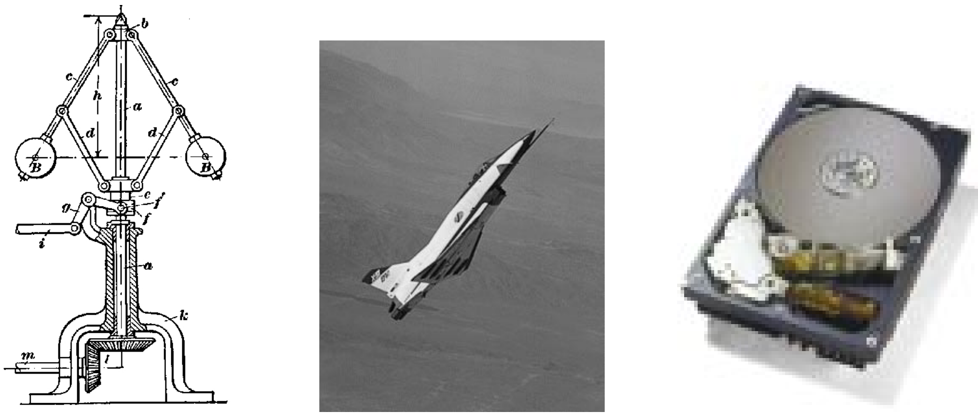

Figure 1: Applications of Control: (a) the Watt governor, (b) flight control and (c) disk drives.

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 conclusions and recommendations of a panel of experts chartered to examine these opportunities. We present an overview of the field, describe 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 require new approaches to meet the challenges and opportunities that face us.

What is Feedback Control?

At its simplest, a control system represents a feedback loop 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 Watt flyball governor, shown in Figure 1. This device used a flyball mechanism to sense the rotational speed of a steam turbine, adjusting the flow of steam into the machine using 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.

A more 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. It is precisely this aspect of Control as a means for ensuring robustness to uncertainty that explains the fact that feedback control systems are all around us in the modern technological world. They are in our homes, cars and toys, in our factories and communications systems, and in our transportation, military and space systems.

The use of Control is extremely broad, but can be roughly broken down into three basic categories:

- 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
- Control of information and decision systems, where limited resources are dynamically allocated based on estimates of future needs.

The 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.

Why Does It Matter?

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

- Guidance and control systems for aerospace vehicles, including commercial aircraft, guided missiles, advanced fighter aircraft, launch vehicles, and satellites. All provide stability and tracking in face of environmental and system uncertainties.

- Control systems in the manufacturing industries, from automotive to integrated circuits. In the latter, the positioning stages of IC steppers rapidly position and align wafers to the extraordinary degree of accuracy required for high yield fabrication.

- Industrial process control systems, particularly in the hydrocarbon and chemical processing industries, that maintain high product quality by monitoring hundreds of sensor signals and making corresponding adjustments to hundreds of valves, heaters, pumps, and other actuators.

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

In addition to its impact on modern 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 implementation 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 master the substantial set of knowledge and skills it entails.

Control Will Be Even More Important in the Future

As we look forward, the opportunities for new applications and new 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 Control, especially as our software systems begin to interact with our physical systems in much more integrated ways. Figure 2 illustrates several systems where these trends are evident.

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 feedback (in its most general sense) plays a central role, since it

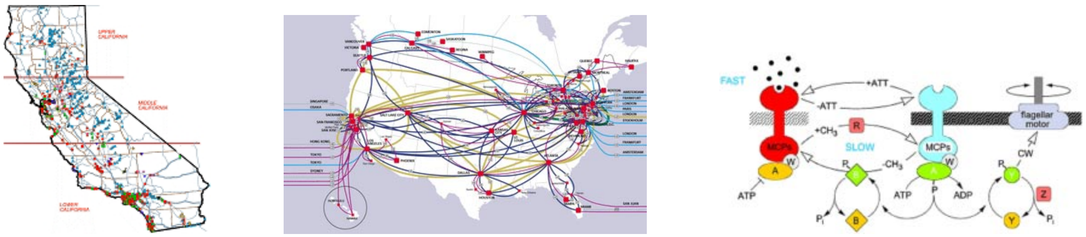

Figure 2: Modern networked systems: (a) the California power network, (b) UUNET's North American backbone and (c) Chemotaxis in *E. Coli*.

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 an interconnected set of components, to realizing control systems as a heterogeneous collection of physical and information systems, with intricate interconnections and interactions.

In addition to cheap and ubiquitous computation, communication and sensing, and the shift from physics-based systems to information-based systems, an important trend in Control is the move from low-level control to higher levels of abstraction. This includes such things as increased autonomy in flight systems (all the way to complete unmanned operation) integration of local feedback loops into enterprise-wide scheduling and resource allocation systems, and control of systems with linguistic and symbolic descriptions. Extending the benefits of Control to these larger scale systems offers enormous opportunities in improved efficiency, productivity, safety, and reliability.

... But It Won't Be Easy

In order to realize the potential of Control applied to these new 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 in distributed, asynchronous, networked environments
- High level coordination and autonomy

- Automatic synthesis of control laws, with integrated validation and verification
- Building very reliable systems from unreliable parts.

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 insure that our current progress is continued and that the opportunities are realized.

What Needs to Be Done

We recommend that the following actions be undertaken to address these challenges and deliver on the promise of the Control field:

1. Substantially increase research aimed at the *integration* of control, computer science, communications, and networking. This includes principles, methods and tools for 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 abstraction, moving toward enterprise level systems. This includes work in dynamic resource allocation in the presence of uncertainty, and learning, adaptation, and artificial intelligence for dynamic systems.
3. Explore high-risk, long-range applications of Control to areas such as nanotechnology, quantum mechanics, biology, and environmental science. Dual investigator funding might be particularly useful mechanism in this context.
4. Maintain support for theory and interaction with mathematics, broadly interpreted (including areas such as dynamical systems, graph theory, combinatorics, complexity theory, queuing theory, etc). 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 basic ideas to nontraditional audiences. For Control to realize

its potential, we must do a better job of educating those outside Control on the principles of feedback and its use as a tool for altering the dynamics of systems and managing uncertainty.

These actions build upon the rich heritage of rigorous work in Control, extending that work to cover ever more complex and significant technological problems. They are key actions to realize the opportunities of Control in the future information-rich world.

2 Overview of the Field

Control is a field with broad application to a number of engineering applications. Its impact on modern society is both profound and 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 loops in electronic amplifiers, proportional integral derivative (PID) controllers in chemical 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 typical example of a modern control system is shown in Figure **??**. In this figure ... [describe key elements]

It is important to note that while feedback is a central element of controls engineering, feedback as a phenomenon is ubiquitous in science and nature. Biological systems are one example. [Add figure of biogical regulation system plus comments.]

While ideas and tools from control theory 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 are most in need of new ideas and new techniques.

Control theory

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 guarantees in performance in the presence of uncertainty. This rigor in approach is a hallmark of modern control theory 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. [Add 1–2 paragraphs on history]

A closely related area is the problem of dynamic allocation of resources under uncertainty, which arises in many problems in operations research. Although not normally considered part of Control, at least in the traditional sense, it relies on many of the same underlying tools (e.g., optimization and game theory) and is an increasingly important aspect as control applications move from physics-based systems to information-based systems.

Control technology

Control *technology* includes modeling, sensing, actuation and computation, used together to produce a working system. A modern control loop 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. 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.

Figure **??** shows some of the trends in sensing, actuation, computation and communications in automotive applications. [Add a chart showing how the number of sensors, actuators, microprocessors, and communications bandwidth is evolving in automobiles]

Feedback as a tool

Feedback is a powerful tool in making engineering systems behave in a desired fashion. Through feedback, we can *alter the dynamics* of a system to meet the needs of an application. Thus, 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. In addition, feedback can be used to *manage uncertainty*. By measuring the operation of the system and comparing it to a reference, we can force the system to respond properly even if the systems dynamics are not exactly known or external disturbances are present that would normally force the system to respond incorrectly.

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).

Control 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 do 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.

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.

Comparison with physics, computer science, and 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 very different from these subjects, in both insights and approach.

Modeling of physical systems is common across all engineering and scientific disciplines. One of the key differences in control-oriented modeling 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. For example, disturbance rejection is naturally described by input/output models. 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.

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

2.2 Control System Examples

Control systems are all around us in the modern technological world. They are in our homes, our cars, our toys, our factories, transportation and communication systems, and of course, 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.

The proliferation of control in engineered systems has occurred primarily in the latter half of the 20th Century. There are some familiar exceptions, of course, such as the Watt governor, designed in the mid-1800s to regulate the speed of steam engines, and the thermostat, 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 the plant “on” when the error is negative and “off” again as it grows positive. This explanation captures the essence of feedback, but it is a bit too simple even for a basic device such as the thermostat. Actually, because lags and delays exist in the heating plant and sensor, a good thermostat does a bit of anticipation, turning the plant off before the error actually changes sign. This avoids excessive temperature swings and cycling of the heating plant.

This modification illustrates that, even in simple cases, good control system design it not entirely trivial. It must take into account the dynamic behavior of the object being controlled in order to do a good job. The more complex the dynamic behavior 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 very broad public exposure was the "cruise control" option introduced on automobiles in the mid-1960s. 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 system operating reliably, robustly and well, 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, of course, 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, and to power generation and distribution over the power grid. They now also play critical roles in the routing of messages across the Internet (TCP/IP) and in power management on wireless communication systems.

Similarly, control systems have been critical enablers in the aerospace and military world. We are all familiar, for example, with the saturation bombing campaigns of World War II, needing to drop 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 very sophisticated control systems, combining inertial guidance sensors, radar and infra-red homing seekers, satellite navigation updates from the global positioning system, and very sophisticated processing of the “feedback error”, all combined in an affordably disposable package.

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

Another notable success story for control in the aerospace world comes from the control of flight. More dramatically than many others, 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. These adjustments were critical because the Wright Flyer itself was unstable, unable to 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, early analog electronic systems to construct and process feedback errors, and early 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 (differential equations, transfer functions, etc.) that accurately describe aircraft dynamics, and they developed three increasingly powerful generations of control analysis and design methods to design control laws. The first generation, known collectively as "classical control", provides methods for single-input single-output feedback loops. The second generation, known as "modern control", provides methods for multi-variable systems where many strongly coupled loops must be designed simultaneously and finally, the third generation, known as "robust multi-variable control", adds powerful formal methods to guarantee desired closed loop properties in the face of uncertainties.

First generation methods enabled the systematic design of early stability augmentation systems, while second and third generation methods are critical in all of today's modern flight control 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).

2.3 The Shift to Information-Based Systems

Early applications of control focused on the physics of the system being controlled. As computers and communications become more ubiquitous, control of *information systems* is taking on a more significant role. Recent examples include new protocols for routers on the Internet, power control in wireless systems, and dynamic optimization of supply chains.

In this context, Control can be described as the science and engineering of optimal dynamic resource allocation under uncertainty. The term "dynamic" means that current actions or decisions will affect the future behavior of the system, so future effects must be taken into account when deciding on current actions. The term "resource allocation" means that our decisions often can be interpreted as managing a trade-off between competing goals, or limits on the actions we can choose from.

But the last term, "uncertainty", is critical: it means that there is some uncertainty in the system's behavior, so that decisions have to be carefully made taking this uncertainty 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 or noise and disturbance signals that affect the system.

This abstract view of control includes a wide variety of applications, from traditional electro-mechanical systems to control of networks, supply chain management, and even optimal investment strategies.

2.4 Opportunities and challenges now facing us

Although in the past there were often a relevant precedents to serve as guide, the rapid evolution of computation and communication capability has led to a situation in which past solutions are less likely to provide useful models.

Characteristics of the New Environment

Computation, Communication, 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 whose outputs can be made available via communication networks. It will take decades to take full advantage of these developments. Some innovation will involve stand alone, or "receiver" type items, and some will involve extreme interconnectedness of the type seen in the telephone system and its descendants. Both types may, and probably will, depend on the use of feedback 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.

Complexity Air traffic control systems, power grid control systems, etc. are typical of a class of systems 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 here is no value in a product that almost works.

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

Internet viruses and ideas from neuroscience may inspire new developments in building reliable systems using unreliable components.

Reliable Systems with Unreliable Parts Most reasonably complex man-made systems are not rendered inoperable by a 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 which 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 down time can be tolerated.

Visualizing the Future

The conventional engineering view of the future is often encapsulated by the words, "better, faster, cheaper". This evolutionary attitude has resulted in amazing progress and continues to represent an important part of engineering. However, this is not the point of view that will lead to a new air traffic control system or the development of a quantum computer. The recognition of the potential of systems that require and exploit carefully managed complexity has been slow in coming even as the number of truly complex systems integral to our lives is growing rapidly. The systems used for managing the electric power grid or the control of traffic on the Internet are, perhaps, only preliminary examples of what we will see in another decade.

Approach

The problems addressed here should be addressed on two fronts. There is a need for a broadly supported, active research program whose goals are to research and further develop methodologies for design and operating 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 should include scientists trained in software engineering, molecular biology, statistical mechanics, systems engineering and psychology. Concrete goals for this work would include the description and rationalization of particular forms of discipline to be used in the design and operation, guidance about the percentage of the resources that should be allocated to security and privacy issues, considerations of the evolvability and expandability of the system, etc. Of course it will be impossible to completely separate these questions from the more obvious "quality of service" issues, so these must be considered at the same time.

Discipline, Rigor and Heuristics

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 mathematics 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 nonconvex 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 find a place for effective heuristics in the training of students and the highest level professional meetings of the field.