Control in an Information Rich World

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

Cover photo here

Ideas:

- Collage: flyball, fbk loop, airplane, disk drive, plots (dynamics), network controlled car, chemotaxis, ATC. [Collect from NASA, Panel, etc; get Jessica and/or Art Center to put together.]
- Painting?
- ATC?

[Note] *Need to identify an exciting cover photo. Some ideas are to do a collage (old to new, across disciplines, etc) or the air traffic control problem (nice fit to main recommendations).*

Control in an Information Rich World

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

DRAFT: v2.4, Sun Mar 3 21:49:38 2002

Abstract

The field of *Control* provides the principles and methods used to design physical and information systems that maintain desirable performance by 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 feedback 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 for Control principles and methods are exploding. Increasingly, computation, communication and sensing will be cheap and ubiquitous, with more and more devices including embedded, high-performance processors, sensors, and networking hardware. 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.

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, autonomous systems, or enterprise-wide supply chain systems 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 Control to play over the next decade.

Panel Membership

Richard M. Murray (chair)
California Institute of Technology

Karl J. Åström
Lund Inst. of Tech.

Stephen P. Boyd
Stanford University

Siva S. Banda
Air Force Research Laboratory

Roger W. Brockett
Harvard University

John A. Burns
Virginia Tech

Munzer A. Dahleh
Massachusetts Inst. of Tech.

John C. Doyle
California Inst. of Tech.

John Guckenheimer
Cornell University

Charles J. Holland
Department of Defense

Pramod P. Khargonekar
University of Florida

P. R. Kumar
University of Illinois

P. S. Krishnaprasad
University of Maryland

Greg J. McRae
Massachusetts Inst. of Tech.

Jerrold E. Marsden
California Inst. of Tech.

George Meyer
NASA

William F. Powers
Ford

Gunter Stein
Honeywell

Pravin Varaiya
University of California Berkeley

Additional Contributors

Jim Batterson, Dennis Bernstein, Raffaello D'Andrea, Jonathan How, Jesse Leitner, Hideo Mabuchi, Eduardo Sontag, Kevin Wise

Contents

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 for how to insure 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 Michigan), 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 help 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.

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 basis and context 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, November 2001

(a) (b) (c)

Figure 1: Applications of Control: (a) the Watt governor, invented in 1789 (b) flight control on the X-36 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 that 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(a). 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.

Control was an essential part in the development of technologies such as electricity, communication, transportation, and manufacturing. Examples include auto pilots in military and commercial aircraft (Figure 1(b)), regulation and control of the power grid, and high accuracy control of read/write heads in disk drives (Figure 1(c)). Feedback was in many cases an enabling technology in all application areas and it was reinvented and patented many times in different contexts.

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 and encompasses a number of different implementations and uses. These include control of electromechanical systems, where computer-control 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 the economy, 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.

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. The interaction with these different fields is an important part of the history and strength of the field.

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. Computer controlled machines provide 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, that maintain high product quality by monitoring thousands 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 software systems begin to interact with 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 Control (in its most general sense) plays a central role, since it

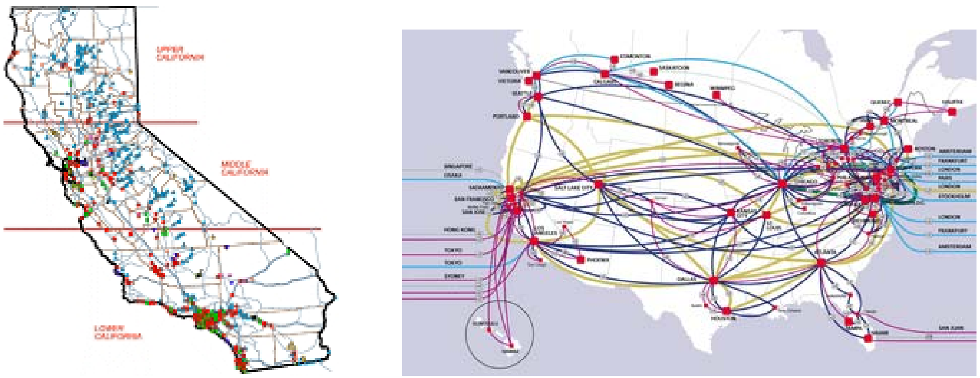

(a) (b)

Figure 2: Modern networked systems: (a) the California power network and (b) UUNET's North American backbone. [Note] *Work on improved figures; add ATC picture?*

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 inexpensive and ubiquitous computation, communication and sensing—and the correspondingly increased role of information-based systems—an important trend in Control is the move from low-level control to higher levels of abstraction. This includes such advances 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 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 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 non-traditional audiences. For Control to realize

its potential, we must do a better job of educating a broader range of scientists and engineers 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 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 3. The basic elements of of sensing, computation, and actuation are clearly seen. In modern control systems, computation is typically implemented on a digital computer, requiring the use of analog-to-digital (A/D) and digital-to-analog (D/A) converters. Uncertainty enters the system through noise in sensing and actuation subsystems, external disturbances that affect the underlying system physics, and uncertain dynamics in the physical system (parameter errors, unmodeled effects, etc).

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

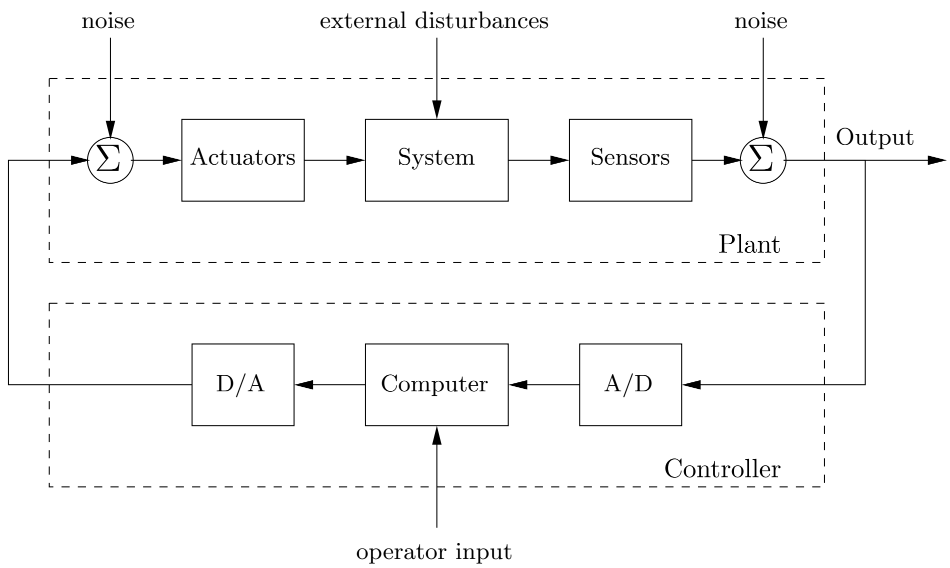

Figure 3: Components of a modern control system.

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 are most in need of new ideas and new 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 guarantees in 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 insure 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. State space methods make use of (linear) ordinarily differential equations to study the response of systems and control is achieved by placing the eigenvalues of the closed loop system to insure 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 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.

Central to all control systems is the role of dynamics, 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 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 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. 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 dy-

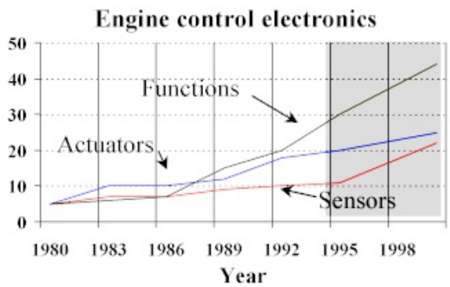

(a)

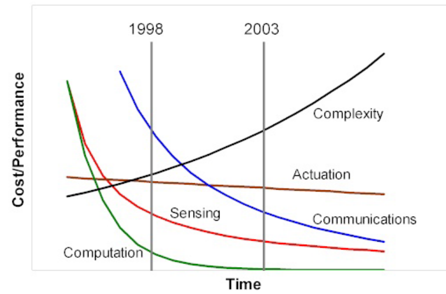

(b)

Figure 4: Trends in Control Technology.

namics of the resulting system is determined not only by the dynamics of the components, but 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).

Control Technology

Control *technology* includes sensing, actuation and computation, used together to produce a working system. Figure 4(a) 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 4(b), is also insightful. The costs of electronics technologies, such as sensing, computation, and communications, is decreasing dramatically, enabling more information processing.

Perhaps 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 very 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 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, 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 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, toys, and manufacturing processes,

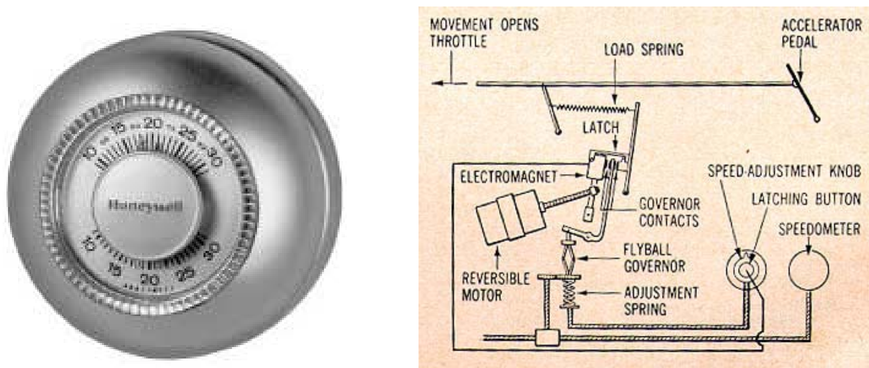

Figure 5: Early control devices: (a) Honeywell T86 thermostat, originally introduced in 1953, (b) Chrysler cruise control system, introduced in the 1958 Chrysler Imperial (note the flyball governor).

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

This modification illustrates that, even in simple cases, good control system design it not entirely trivial. It must take into account the dynamic behavior of the object being controlled in order to do a good job. The more complex the dynamic behavior 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 1958. 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, 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, 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 on wireless communication systems.

Aerospace Applications

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 infrared homing seekers, satellite navigation updates from the global positioning system, and very sophisticated processing of the "feed-

Figure 6: Flight systems: (a) 1903 Wright Flyer, (b) X-29 forward swept wing aircraft, in 1987.

back 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, 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. 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 [1]. These adjustments were critical because the Wright Flyer itself was unstable, and could not maintain steady flight on its own. As Wilbur Wright said when lecturing to the Western Society of Engineers in 1901 [2]:

> "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."

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 (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 ubiquitous 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 often can be interpreted as managing a trade-off between competing goals, or limits on the actions we can choose from. "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.

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 has created a situation that couples vast amounts of information—everything from the location of the plans to the individual customer itineraries—that must be managed to maintain high performance, robust, reliable operation at all times.

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, 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 long 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 these tradeoffs between these will be the primary design challenge in this space.

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 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, 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 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 this area, 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 of potentially very deep, 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, along 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 down time can be tolerated.

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

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

Vision for the Future

This new environment for Control presents many challenges, but also many opportunities for impact across a broad variety of application areas. The future directions in Control, Dynamics, and Systems must continue 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 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 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 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 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 contributers 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 compactification of current knowledge 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 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 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 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 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 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.

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.

[Note] *These sections are still in early draft form. Comments and suggestions are appreciated.*

```
A. Introduction
    1. Organization by application
    2. Source for data in report
    3. Source for recommendations
    4. Organization of each section
```

3.1 Aerospace and Transportation

Aerospace and transportation encompasses a collection of critically important application areas where Control is a key enabling technology. These application areas represent a very large part of the developed 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 key 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].

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, into 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. [see Emission Control Vignette]. 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.

As a historical note, the cruise control option introduced in the mid-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 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 and 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 where 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 design 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. 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 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 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. 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 vs. 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 Controls 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 Controls 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 components. 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 air space management problem, the automated highway/traffic management problem, and the problem of managing future battlefields.

The national air space problem is particularly significant today, with eventual gridlock and congestion threatening the integrity of the existing air transportation system. Many studies are under way attempting to modernize the way traffic is managed, the way individual aircraft schedules and flight paths are established, and the way the system adjusts to upsets due to local weather, local equipment failures, and various other disturbances. General solutions being explored are called "free flight". They involve distributed calculations of flight plans and trajectories aboard individual aircraft, free of established air corridors, flight plan coordination via negotiations and ground based assistance, and automated collision avoidance technology. This is yet another application where the strong system-theoretic background and the tradition of rigor held by the Controls community can make substantial contributions.

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

[Note] *Add an introductory paragraph and summarize any other items that came up in the panel meeting.*

Propulsion Systems [Note] *1 paragraph on propulsion systems, along the lines of smart engines, VAATE, etc.*

Space Systems 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 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 the areas of 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. If serious consideration is given to these approaches, it will take a conscientious, integrated, cross-agency/industry/academia partnership with a carefully planned implementation of the three-pronged approach defined earlier in this section if success is to be achieved.

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 100's 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.

Grand Challenges

[Note] *Add in some grand challenge problems that illustrate the types of things we think we could accomplish by 2020. Should be aligned with main recommendations.*

```
(1) ATC, free flight
(2) Personal aviation
```

(3) Nanometer estimation and control across kilometer distances

3.2 Information and Networks

The rapid growth of communication networks provides several major opportunities and challenges for systems and 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 below.

In routing control, the flow of packets through the network is controlled. In the simplest form each router must decide the order, and which output link(s) the packets should be sent to on their way to their final destination(s). 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. Current decisions clearly affect the future state of the network, such as the future traffic on links, future buffer levels, queuing delays, etc. Resources that must be managed include router resources like buffer limits, and link resources, such as capacities. Performance is judged in many ways: latency, delay, loss rates, bandwidth, for various streams and other types of traffic. Other measures include how well and how fast the network adapts to changing network congestion, changing traffic patterns, etc.

Several features of these control problems make them very challenging. One is the extremely large scale of the system. Another is the variation in network topology that the routing control must be able handle. Yet another is the decentralized nature of the control: local decisions must be made fast, and so have to be based on locally available information. Of course, information about link congestion, router queues, and traffic demands can be sent over the network between routers. This takes away capacity from real network traffic, and also provides delayed information, which may not be as useful for control. Another complicating issue is the different types and classes of network traffic, all with different requirements for quality of service, in terms of delay, bandwidth, loss probability, etc.

Like many control problems, routing control 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 we have optimal network planning, in which new links and nodes are proposed to meet (predicted) rising traffic demand.

Another example of a challenging dynamic resource allocation problem arising in networks is optimal forward caching. The task 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 (in network topology sense) server, thereby reducing network traffic and delay. If the source document changes, the changes (at least) must be transmitted to the servers, which consumes some network bandwidth that otherwise would have been available to real network traffic. The problem is to devise a scheme for how often to update, and where to cache copies of documents, based on predictions of access patterns and network congestion.

Several other important control problems arise in wireless networks, including for example the important question of power control. Here the transmitters must decide on an appropriate transmit power, which guarantees a large enough signal to noise and interference ratio (SINR) at the receiver for accurate reception. The subtlety is that the transmitted signal also appears as an interference term for other wireless links, so increasing the power of a transmitter can affect the SINR of many wireless links. Here too, we require fast, possibly asynchronous, somewhat decentralized decisions for a very large, coupled system.

In ad hoc wireless networks, the power control problem is coupled with the routing problem. In the simplest case each of n nodes can transmit to any other node. When nodes are far apart, higher transmit power is required, which in turn has the adverse affect of increasing the interference for other links (thus requiring other transmitters to increase their power). A more sophisticated model connects the transmit powers, or really the receiver SINR, to the maximum possible bit rate (i.e., capacity) over the link. In an ad hoc wireless network, packets can be routed from source to destination over a sequence of nodes. We seek joint power and routing control algorithms that adapt to rapidly changing network conditions, connectivity, and traffic conditions to satisfy constraints on bandwidth, delay, packet loss, and efficiency (as measured, for example, by total transmit power or total spare capacity). Once again we seek algorithms that are mostly decentralized, or efficiently use internode communication for control coordination.

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

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 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, robustness 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 varying amounts, 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 though, 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 synchronous clocked systems and networks that guarantee arrival and bound 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 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 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.

Other Trends in Information and Networks

```
(1) Sensing + vision as a sensor
(2) Signal processing and control
(3) Control challenges in space applications (?)
(4) Vigilant, high confidence systems (formal methods, etc)
```

Grand Challenges

```
(1) Real-time, supply change management
(2) Unified theory for computation, communications and
    control (performance + robustness)
```

3.3 Robotics and Intelligent Machines

Background and History

The goal of cybernetic engineering, already articulated in the 1940s and even before, has been that of implementing 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 [**?**]. A more mathematical treatment of the elements of engineering cybernetics was present by Tsien in 1954, driven probem problems related to control of missiles [**?**].

The early work leading up to today's robotic systems begain after World War II in the development of remotely controlled mechanical manipulators, which used a master-slave mechanims. 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 the robot to respond to changes in its enviroment and by the 1960s many new robots were capable of grasping, walking, seeing (through binary vision) and responding to voice commands.

The 1970s and 80s saw the advent of computer controlled robots and the field of robotics became a fertile ground for research at the intersection of Control and Computer Science. Manufacturing robots became commonplace (led by Japanese companies) and a variety of tasks, ranging from mundane to high precision, were undertaken with mechanical machines. Artificial intelligence (AI) techniques were also developed to allow higher level reasoning, including 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 accomplishements that demonstrate the successes of the field are the Mars Sojourner robot and the Sony Aibo robot, shown in Figure 7. Sojourner successfully manuevered on the surface of Mars in July 1997 and sent back live pictures of its environment. The Sony Aibo robot debued 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 judgement.

The vision for robotics and intelligent machines remains much as it did in the time of Wiener: to develop systems that are capable of human-like behavior and interact seemlessly with humans and society. Robotics provides

Figure 7: The Mars Sojourner and Sony Aibo robots.

an excellent application area for Control, since it has a history of strong interaction between the control and computer science communities, and future advances will require further integration of computing, communication and control, higher level reasoning and decision making, and strong theory and interaction with mathematics.

Challenges and Future Needs

The basic electromechanical engineering and the computing capabilities required to build practical systems of this type have evolved over the last half-century to the point where today there exists 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 demonstrations 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, 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'etre*, it is logical to regard motion control as being the 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.

A major area of study is adaptability and learning. As robots become more common place, they will need to become more sophisticated in "self-programming" by reasoning about their environment and 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 commonplace in our society, major advances in machine learning and cognitive systems will be required. Robotics provides an ideal testbed for such advances and applications in remote survellience, search and rescue, entertainment robots, and personal assistants are all fertile applications 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 US military is considering the use of multiple vehicles operating in a coordinated fashion for surveillance, logistics 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. 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, 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 science, 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 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; perhaps this work can 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 unstructored 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 and that will drive new research in Control.

Medical Robotics.

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. Yet 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. 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.

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 understnading of the *interactions* of humans with machines and computer controlled systems.

One current program underway in the United States is exploring the use of “variable autonomy” systems, in which machines controlled by humans are giving differing levels of command authority as the task evolves. Such systems involve humans that are integrated with a computer-controlled system in 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 having 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.

Grand Challenges

```
(1) Autonomous exploration in remote environments
(2) Personal assistants: apprentice (being taught) + nurse/butler
(3) Effective use of robots in medicine
```

3.4 Biology and Medicine

Vision: Systems Biology and Biological Engineering

Molecular Biology

The life sciences are in the midst of a major revolution, which will have fundamental implications in biological knowledge and medicine. Work in genomics has as its objective the complete decoding of DNA sequences, providing what one may call a "parts list" for the proteins potentially present in every cell of the organism being studied. The elucidation of the three-dimensional structure of the proteins so described is the goal of the area of proteomics. The shape of a protein, in turn, determines its function: proteins interact with each other through "lego-like" fitting 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 (UV radiation, temperature) as well as chemical (drugs, hormones, nutrients) signals. Its outputs include chemicals which may in turn affect other cells. Each cell can be though of, in turn, as composed of a large number of subsystems, involved in cell growth and maintenance, division, and death. A typical diagram found in the biological literature is shown in Figure **??**.[1].

The study of cell networks leads to the formulation of a large number of questions, such as the following ones:

- What is special about the information-processing capabilities, or input/output behaviors, of such networks? Can one characterize these behaviors in terms familiar to control theory (e.g., Volterra series)?

- What "modules" appear repeatedly in cellular signaling cascades, and what are their system-theoretic properties? (An example of such modules is provided by the ubiquitous "MAPK (mitogen activated kinase) cascades".)

- Inverse or "reverse engineering" issues include the estimation of system parameters (such as reaction constants) as well as estimation of state variables (concentration of protein, RNA, and other chemical substances) from input/output experiments. Generically, these ques-

[1] figure cancernet.jpg enclosed; need copyright permit from journal "Cell"

tions may be viewed respectively as the identification and observer (or filtering) problems which are at the center of much of control theory.

- What are the stability properties of the various cascades and feedback loops which 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 realistic level, one wishes to study the possibility of using of control strategies (both open and closed-loop) for therapeutic purposes, such as drug dosage scheduling.

The need for mathematical models in cellular biology has been long 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 dealing with the 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, or gene arrays).

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, but, on the other hand, large variations in levels of certain chemicals, (such as transcriptional regulators) may be lethal to the cell. Feedback loops are omnipresent in the cell. It is estimated, for example, that in *E.coli* about 40% of transcription factors self-regulate. One may ask whether these feedbacks' role is indeed that of reducing variability, as expected from principles of feedback theory. Recent work[2] tested this hypothesis in the context of tetracycline repressor protein (TetR). An experiment was designed, in which feedback loops in TetR production where modified by genetic engineering techniques, and the increase in variability of gene expression was correlated with lower feedback "gains". 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.

[2] Becskei and Serrano, Engineering Stability in Gene Networks by Autoregulation, *Nature* 2000

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 ("stop tumbling, run towards food") signal, 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.[3]

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.[4]

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

[3] Barkai and Leibler (*Nature*, 1997; Yi, Huang, Simon, Doyle *PNAS* 2000

[4] McAdams and Shapiro, Circuit Simulation of Genetic Networks, *Science* 1995

which deals with the synthesis of dynamical systems which implement a specified behavior.

Integrative Biology

Medical Systems

Imaging.

Other Trends in Biology and Medicine

Grand Challenges

[Note] *Add in some grand challenge problems that illustrate the types of things we think we could accomplish by 2020. Should be aligned with main recommendations.*

```
(1) Systems biology
(2) Embedded control of biological systems
(3) Robotic insects
```

3.5 Materials and Processing

[Note] *Look for examples where controls has made a difference:*

- *optical fiber?*
- *purity of chemicals?*
- *magnetic materials?*
- *silicon?*

Process Control

Nanoscale Systems

Other Trends in Materials and Processing

Grand Challenges

[Note] *Add in some grand challenge problems that illustrate the types of things we think we could accomplish by 2020. Should be aligned with main recommendations.*

```
(1)
(2)
(3)
```

[Note] *The remaining sections are in the processing of being written and will be included when available.*

3.6 Other Applications

Environmental Science and Engineering

Atmospheric systems and pollution

Microbiological ecosystems [Note] *Look at Discover magazine article (??)*

Economics and Finance (Primbs, Doyle?)

Quantum Systems (Physical Sciences?)

[Note] *See article in Science, 5 May 00, 288, 824-828*

4 Education and Outreach

```
IV. Education
    A. The New Environment for Controls Education
        1. Current state: discipline-focused, engineering education
        2. Modern controls engineer = systems engineer
        3. Increased need for controls outside of Ae, ChE, EE, ME
    B. Making Controls More Accessible
        1. Develop books and courses that emphasize feedback
           concepts and requisite mathematics
        [CDS 110a vignette?]
        2. Software packages for implementing latest advances
    C. Broadening the Controls Education
        1. Broader grasp of engineering, science, and math
        2. Increased leadership and communications skills
        3. Balance between theory, applications and computation

Vignette: CDS 110a (Murray, Morgansen, Mabuchi)
Development of a one year course on controls @ Caltech
  - broad audience; three ‘‘sections’’
  - key learning objectives
      + principels of feedback: stability, performance, robustness
      + modern computational tools
  - integration of applications

Additional vignettes from Dennis Bernstein
```

4.1 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 for, limits to, 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 US 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, a two pronged effort in education can be contemplated: 1) developing curriculum for those still in the K-12 formal system, and 2) raising the level of awareness and understanding of the population that has completed its formal education. The latter will require the development of appropriate strategies for the teachers and professors—strategies that bridge prongs 1 and 2.

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. The multidisciplinary nature of Control is very much antithetical to that traditional thinking and structure in education. However there is some evidence of advances toward application and integration of mathematics with science. The Consortium for Mathematics and Its Applications (COMAP) does this. The University of Chicago Schools Math Program sets a new level in teacher preparedness in mathematics. Professor George T. Rublein of the College of William and Mary has written a book that sets forth mathematics as a solution to real world problems and situations centered on aviation [**?**]. This book is aimed at giving elementary school teachers a richer understanding of mathematics in our everyday world.

The general requirements for teacher certification and now popular ”Standards of Learning” tests are set in the scientific epoch of the 1940s and 50s. A 1990s survey in a large metropolitan school district found fewer than one third of the high school mathematics teachers with the equivalent of a mathematics major (36 hours at Calculus or above) [**?**]. Today’s certification requirements have improved but the strong economy that the country enjoys at this writing has precluded the hiring of these certified staff in many school districts. Middle school teachers have even weaker backgrounds in mathematics while elementary school teachers are still weaker on the average. As a result, any program designed to teach Control in the K-12 system must first address the development of teachers at their current level of sophistication in mathematics and science. That is, the teachers can be considered to be members of the group that has completed its formal education without knowledge of Control.

Simple experiments involving governors, thermostats, and “see-saws” can be accomplished with heuristic learnings in the elementary schools. As mathematical sophistication increases through middle school and high school, quantitative analysis can be added and experimentally verified. As the teaching of Algebra and Geometry continue to advance to the early middle school grades, students are beginning to complete Calculus in the junior and sophomore years. A post-calculus course in applied mathematics of differential equations and dynamical systems could be created bridging chemistry, physics, biology, and mathematics.

We must be careful that the dynamics and control curriculum for K-12 reflects the needs of the next 25 years as laid out in this report rather than the past 25 years. This can be accomplished by tying the curriculum development directly to the research universities, R&D agencies, and appropriate industries. One approach to curriculum development would be through a partnership between the NSF divisions responsible for Control and those responsible for K-12 education. A joint BAA could be issued for the development of the prototype curriculum and initial implementation in one or more school districts across the country. Complementary to this approach, federal R&D agencies could sponsor summer research programs in Control for mathematically qualified middle and high school teachers. These program would be designed along the lines of current programs for summer students, immersing the teachers in research projects which include the analysis of dynamical systems data and the design and analysis of control laws. Similar programs could be run by industry.

A second approach would be the development of a pathfinder curriculum in one or more of the various governor's schools for science and mathematics across the country.

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 contemporary capabilities 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 "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 the research grade hardware and software and the Internet, the participants formed partnerships with one another that promoted continued collaboration throughout the coming academic years. It was the responsibility of each school district to institutionalize the capabilities that NASA provided to the teachers.

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.

Finally, the leveraging of any of our efforts with COMAP could prove fruitful. 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.

4.2 Other Opportunities

Popular Books

Multi-media CDROM

Public Awareness

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 reach the entire age range of the public from children through adults. Appearances by top researchers who are also known as excellent speakers and communicators on television interviews at opportune times is desirable.

4.3 Other Trends in Education and Outreach

1. Software tools
2. Interaction with industry

4.4 Grand Challenges

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.

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 key 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, 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 Multi-disciplinary 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 traditional 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 role for Control.

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 abstraction, 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 many 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 exploding 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, 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 Control 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 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 the compactification of the existing knowledge base 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 insure 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, etc.

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 methodologies 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 exploding and this is placing new demands on education. The community must continue to compactify knowledge by integrating materials and frameworks from the past 40 years in a more unified approach. As important, material must be made more acceptable 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 basic ideas 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 giving 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. These 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 PhD 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 multi-media 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, acceleration of the benefits of Control through broader use of its principles and tools, and rigorous design principles that give safer systems, shorter development times, and more transparent understanding of key systems issues.

Appendices

A. Panel Meeting Summary
B. Other reports and workshoips
    1. Scientific American, September 1952
    2. ...
C. Panelist biographies (?)
D. Excerpts, Sept 52 Scientific American
    1. Intro to issue (text)
    2. Abstracts of main articles, with scans of figutes
    3. Selected ads

[Note] *Some additional possible vignettes are on the following pages. There were collected from the AFOSR Scientific Highlights series. Some of these can be used in other parts of the report. Additional vignettes are welcome.*

Vignette: Control Technology Advances Damage Detection Capability, Feedback Controllers (Jan 96)

A team of AFOSR-sponsored applied mathematicians has achieved two major advances resulting from their control theory investigations: an exciting new method for applying "smart" or adaptive materials technology for nondestructive damage detection and characterization in elastic structures such as airfoils and fuselages; and application of the technology to design extremely effective feedback controllers for actively suppressing vibrations and structure-borne noise in these structures.

Professor H. T. Banks, at North Carolina State University, leads the team which includes Dr. Yun Wang, at the Armstrong Laboratory (Brooks AFB, Texas), and Professor R. C. Smith, at Iowa State University. Their achievements build on a decade of AFOSR-sponsored research on use of computational methods to identify and control distributed parameter systems.

The team's work on damage detection in elastic structures shows great promise for a simple, effective means to use evolving smart material- based technologies to evaluate materials and structures for damage (caused by internal and/or external stresses) during a structure's lifetime. Their demonstration of the possibility and feasibility of using vibration analysis to detect damage resolves controversial issues in the engineering community for the past 20 years. Using physically based distributed parameter models for structures, they successfully developed and experimentally validated methods to detect and characterize certain types of damage and defects in elastic structures. These methods are based on vibration responses to characterize changes in fundamental geometry and physical parameters (mass density, stiffness, damping) using embedded piezoceramic sensors/actuators in a self-exciting, self-sensing structural environment. Their efforts produced capabilities and a level of resolution that can't be attained with classical modal-based methods that are at the heart of the controversies.

The team used the same models with the embedded sensors as the basis for a feedback control methodology to suppress vibrations in mechanical systems and structure-borne noise in structural acoustics systems. By developing necessary on-line filters and compensators – to treat partial observations of the system – the team implemented and tested the methodology on a component of a structural acoustics experimental system subject to both periodic sustained disturbances and "impulse" or transient disturbances. In an experiment with a clamped plate, their closed loop feedback design yielded outstanding stable performance in both cases: the response to sustained disturbances were typically reduced by 85 percent, while typical settling times were reduced fourfold in transient responses. These experimental verifications conclusively demonstrate for the first time that a distributed parameter feedback control methodology in the context of smart material vibrations and noise suppression is both feasible and can significantly improve current engineering capabilities.

Vignette: Scientists at Washington University Solve Long Standing Filtering Problem (1997)

Researchers at Washington University in St. Louis, led by Christopher I. Byrnes and Anders Lindquist, have solved a long-standing open problem in the modeling of random signals with important implications in signal processing, in speech synthesis and in linear prediction of time series. In many applications one starts with a finite window of statistics, e.g. a finite number of correlation coefficients of a time series, and needs to design a linear filter which "shapes" white noise into a time series with the observed finite window of correlation coefficients. Many commercial applications involving speech synthesis or signal processing use a particular solution of this problem, the maximum entropy filter, for which there exists several well-known algorithms for determining the filter parameters. Particulary important filter parameters are the filter poles, which determine stability, and the filter transmission zeros, which determine the frequencies of periodic signals which are absorbed, or attenuated, when passed through the filter. The maximum entropy filter has no transmission zeroes, and consequently other filters which are designed to absorb signals with a given frequency may, in some cases, provide a better model for the observed data. The main result, discovered by Professors C. I. Byrnes and A. Lindquist with their collaborators, asserts that, for a given window of n correlation coefficients, to each choice of n zeros inside the unit disc there corresponds one and only one filter which shapes white noise into a process with the given window of statistics. This parameterization is thus given in terms of a familiar systems theoretic quantity which can be directly related to desirable features in models used for the given data. Moreover, further research has resulted in the development of a highly efficient algorithm for determining the filter parameters from the data and the choice of zeros. Use of this result and the algorithm for correlation coefficients computed from, e.g. 30 millisecond, windows of human speech have potential applications in the improved synthesis of speech both for mobile, digital communications and for the potential development of new voice signature methodologies for secure communications.

The principal difficulty encountered in solving this parameterization problem is that the parameters arising in these shaping filters depend nonlinearly on the given correlation data. If one suppresses realizability of the filter by a circuit with finitely many active elements, this mathematical problem was posed and solved in the early 1900's by Caratheodory, Schur, and Toeplitz. If instead of realizability by a circuit one allows realizability by a analog computer then this modified problem was solved in the 1970's and 1980's by system theorists, beginning with fundamental work by Kalman on the realization of deterministic signals. However, matching the observed data with a random process generated by a finite dimensional filter renders the problem highly nonlinear. In the early 1980's, T. Georgiou used a nonlinear tool, degree theory, to prove that rational shaping filters, realizable by circuits of degree n and corresponding to a choice of zeros, do exist. The recently established complete parameterization followed from a combination of nonlinear dynamics and geometry, which enabled a considerable refinement of degree theoretic methods to prove both existence and uniqueness of such a shaping filter. In particular, considerable use was made of a nonlinear dynamics analysis of fast algorithms for Kalman filtering, an analysis spurred in part by an effort to understand the design of observers for nonlinear dynamics. Also important to the proof was a detailed study of the geometry of the space of shaping filters of given degree, which revealed that the problems of Kalman filtering and the problem of designing shaping filters matching a given window of correlation coefficients are related in a very precise geometric sense. From this geometric interpretation, a solution ot the problem of parameterizing all shaping filters with a given correlation window can be solved using a result of Hadamard.

Vignette: New Active Disturbance Rejection Technology Improves Airborne Laser Aim (Dec 1995)

A group of Air Force and university control scientists has developed a new adaptive feed-forward vibration-cancellation methodology. This new adaptive control technology could sharply improve the performance of a wide variety of Air Force systems—ranging from satellites to engines to aircraft structures—which require the active control of vibrations for effective operation.

Basic researchers Don Washburn and Rick Walter of the Air Force Phillips Laboratory (Albuquerque), Fred Boelitz of Aerospace Corp., and Steve Gibson of UCLA recently demonstrated the effectiveness of the control technology during a series of airborne experiments on a C-135 aircraft as an adjunct to an airborne laser data-collection experiment, ABLE ACE.

The Phillips Lab is already applying the concepts demonstrated in these experiments in an experimental program to cancel acoustic noise in precision systems. The direct pointing improvement and acoustic noise cancellation were inspired by and apply directly to the Air Force's high priority Airborne Laser System (ABL). Col. Richard D. Tebay, director of the ABL program office, cited the researchers' excellent technical work and vision in meeting this technical need. "This effort is an excellent example of how AFOSR basic research efforts can directly insure the success of advanced Air Force and DoD systems."

The most novel feature of the new adaptive methods was the adaptive least-squares lattice filter used for adaptive identification of the feedforward gains. This lattice filter and the new algorithm based on it were developed at UCLA under AFOSR sponsorship. The main advantages of the least-squares lattice over older methods (such as least mean squares filters) in adaptive vibration cancellation methodologies include faster convergence (and adaptation), numerical stability for large-order and multichannel filters, and inherently parallel computing architectures.

The recent experiments represent several advances in active-vibration suppression and noise cancellation. Adaptive noise cancellation methods have been applied to stationary disturbance rejection, but the ABLE ACE experiments show the new methods improve precision tracking in a difficult but realistic airborne environment dominated by nonstationary disturbances. The combination of nonstationary and large track-error components uncorrelated with noise references makes adaptive identification of optimal filter gains particularly challenging. In these experiments, the aeroelastic interaction of atmospheric turbulence and the C-135 airframe produced disturbances that were highly nonstationary.

Vignette: New Mathematical Tool Will Improve the Design, Performance of DoD Systems (Jul 96)

A team of AFOSR-supported applied mathematicians at Virginia Tech's DoD- URI Center for Optimal Design And Control has developed a new, fast algorithm for improving the design and performance of existing and future military systems. The method has already been successfully tested on a wide variety of Air Force problems, including the design of jet engine inlets, nozzle design, and aeroelastic tailoring to enhance aircraft lifespans.

This new mathematical tool has a multitude of potential DOD applications, such as wing-body design and optimization, flow tailoring for improved performance, propulsion system design, and multidisciplinary design optimization of aircraft. Many non-military applications, such as improved design and function of inkjet printers, are also possible.

Professor John Burns, the center director, leads a multidisciplinary team of specialists in distributed parameter control, guidance, fluid mechanics, optimization, and computational mathematics. Their new methodology builds on a solid foundation (10 years) of AFOSR-sponsored research in distributed parameter control and computational mathematics.

The new approach shows great promise as an effective means to eliminate the need to compute mesh sensitivities for gradient-based aerodynamic design applications. Meshing can account for as much as 90design cycle and computational time in 3-D flows. The new algorithm is a Sensitivity Equation Method (SEM) based upon partial differential equations that define flow sensitivities. The sensitivity equations are used as a guide to develop and select efficient computational methods for approximating sensitivities and the corresponding gradients. The algorithm combines robust optimization with carefully chosen numerical schemes. This approach has reduced design cycle times by as much as an order of magnitude in certain nozzle design applications.

The team has also developed a mathematical framework to construct and analyze new numerical algorithms that can be applied to a wide spectrum of aerospace systems. Using this framework they have developed a methodology that, for the first time ever, allows a designer to select in advance those combinations of discretization schemes that guarantee convergence of the resulting optimal design algorithm. This result is important for verifying the reliability of final designs.

Vignette: New Control Concepts will Enhance Aggressive Flight of Uninhibited Combat Aerial Vehicles (Mar 98)

Researchers at the California Institute of Technology have developed a new approach to nonlinear control that promises to radically improve the performance of future missiles and uninhabited combat aerial vehicles (UCAVs) engaged in aggressive maneuvers. Their development of new theoretical and computational tools is providing an extremely effective methodology for computing aggressive trajectories for certain classes of flight systems in real-time: a critical need for greatly enhancing the maneuvering capabilities of future uninhabited combat aircraft.

Professor Richard M. Murray leads the research team conducting this work, part of the AFOSR-sponsored Partnership for Research, Excellence and Transition (PRET) Center for Robust Nonlinear Control of Air Vehicles. Murray also directs the Center, which was established in 1995 at Caltech. Aggressive trajectory tracking in flight systems requires fast and accurate calculation of the desired attitude and control surface motions needed to complete a maneuver. Traditional flight control systems either rely on pilots for this task—which exploit a human's unparalleled ability to learn the behavior and limits of a system and to operate near the boundaries of achievable performance—or make use of precomputed trajectories and maneuvers. However, in the future, aerial vehicles are destined to be uninhabited, though they will still require high performance in unstructured environments. Pilots will have to remotely control one or several (heterogeneous) aircraft. Future flight control systems must perform commanded maneuvers that push the edge of the achievable operating envelope while respecting the aircraft's dynamical limitations.

These problems motivated Caltech's new work in nonlinear control theory, which is based on a mathematical property known as differential flatness. Differentially flat systems share an important property: the generation of the system's trajectories can be reduced from a dynamic problem - requiring techniques that are computationally demanding and extremely difficult to implement in real-time—to an algebraic problem—for which tractable computational algorithms can be generated.

Murray and his group are now developing new techniques for characterizing differentially flat systems as well as new algorithms for exploiting flatness to generate trajectories for the system in real time. This is an important step in understanding how to design vehicles with properties that can be exploited by the new nonlinear control techniques.

Differential flatness is beginning to play a role in the development of control laws for some air-to- missile systems, where high performance in rapidly changing conditions is essential.

The Caltech group has implemented and tested their algorithms for real-time trajectory generation on a small, flight-control experiment which mimics the longitudinal dynamics of a thrust-vectored aircraft, demonstrating substan-tial improvement over conventional algorithms. Murray's PRET partners—Honeywell, Boeing, Northrop, and Hughes—are currently investigating the use of this technology in more realistic systems.

Vignette: Scientists Develop New Method for Automatic Control System Design (Jan 95)

Professor Stephen Boyd of Stanford University, in collaboration with several other AFOSR-sponsored researchers including John Doyle (Cal Tech) and Michael Safonov (USC), has developed a completely new approach to control system analysis and design. The method will allow the application of advanced and powerful techniques to a vast number of practical control problems involving nonlinearity and model uncertainty. For example, changing mission requirements often cause the Air Force to make substantial reconfigurations of aircraft and missiles which in turn require major changes in control system design. When implemented in automatic control system design software, the new method will provide an efficient and reliable means to shorten the labor-intensive control system design cycle for these reconfigurations.

One of the researchers' innovations was to show how a number of difficult fundamental control problems can be formulated as mathematical optimization problems involving matrix inequalities. These nonlinear, large-scale optimization problems can appear difficult and generally do not have analytical solutions. In a major breakthrough, the researchers succeeded in developing efficient interior-point algorithms to solve these problems numerically. Boyd and his colleagues were able to show that by exploiting the structure of the control problems, great efficiencies could be obtained. Their methods can solve problems involving thousands of variables and tens of thousands of constraints in a few minutes on individual workstations.

Several commercial computer-assisted control design packages based on this method are under development. The method has already been used with great success in several challenging, experimental preliminary design studies including the control of highly-maneuverable air-to-air missiles and the control of advanced, semiconductor manufacturing equipment. Because of its great computational efficiency and the wide variety of practical problems that can be addressed, the method is especially well-suited for computer-aided control system design. The matrix inequality solution algorithms can be readily embedded in computer-aided design tools that incorporate a natural user-interface for problem specification. The control problem is automatically translated into the appropriate matrix inequality problem which is then solved using efficient interior point methods.

Vignette: New Nonlinear Control Concept Expands Jet Engine Operating Envelope (Mar 95)

Professor Eyad H. Abed and his research group at the University of Maryland, College Park, have introduced a new nonlinear control design which will permit the operation of an axial flow compressor up to and beyond the stall limit. This research provides a key step toward the integrated control of jet engine aerodynamic and combustion instabilities which will lead to lighter engines with greatly improved performance characteristics for use in future military and civilian aircraft.

Linear methods cannot be used to stabilize the system on the stall boundary using practical actuation methods. The new design concept discovered by Abed's team uses bifurcation theory and other nonlinear analysis tools to give a practical characterization of the family of smooth nonlinear feedback control laws which result in a supercritical bifurcation (in simplified terms this means the new robust feedback controller causes a nearby stable operating condition to emerge). This work takes advantage of the compression system model developed and refined earlier by Edward Greitzer at MIT (AFOSR-sponsored) and Frank Moore at Cornell. Abed's new nonlinear control research motivated related work by Dr. Carl Nett and his group at Georgia Tech. Nett's group provided experimental validation of the concept at low speeds and pointed out key deficiencies at high speeds representative of industrial compressor designs. Dr. Nett subsequently moved to the United Technologies Research Center (UTRC) where the work is being continued. UTRC researchers have succeeded in modifying the original concept for use at high speeds and have filed a patent application on the modified approach.

Axial flow compressors control the pressure ratio and mass flow of most large aircraft engines. Operation of these compressors near their maximum pressure ratios can cause two types of aerodynamic instability: surge and rotating stall. The desirability of operating at maximum pressure ratio has stimulated several investigations into active control schemes that "quench" these instabilities. The combined efforts of the Maryland, MIT, Georgia Tech and UTRC groups has resulted in a new nonlinear approach that is proving effective in theory and practice. In addition to jet engines, the concept has potential uses in gas turbines in chemical plants and electrical power plants.

Vignette: Control Theory Used To Reduce Costs For Air Force Weapon Systems (Nov 98)

Air Force use of new flight control system design tools are significantly cutting design and development costs for current operational and test munition weapon systems. AFOSR-sponsored researchers at the University of Illinois Urbana-Champaign (UIUC), working with engineers at The Boeing Phantom Works in St. Louis, have developed and transitioned the new design tools for the Joint Direct Attack Munition (JDAM) and Miniaturized Munitions Technology Demonstration (MMTD) smart weapons.

The JDAM (operational) and MMTD (in test) gravity bombs use Global Positioning System navigation to accurately guide the weapon to target. Affordability is the primary objective in these programs. The new flight control design technology played a key role in meeting this directive by enabling Boeing to design low-order, simple flight controllers which:

- minimized the number of sensors,
- reduced parts count, and
- reduced weight, while meeting restrictive packaging and volume constraints.

The new MMTD smart bomb has the same military effectiveness (in tests) against fixed hard and soft targets as the currently fielded JDAM. However, the bombs weight and volume decreased approximately 10-fold. These features will:

- significantly reduce logistics operations,
- enable new, low-observable aircraft such as the F-22 to carry many more bombs in their internal bays, and
- still maintain their stealth characteristics.

This is a stunning example of the impact of basic research. The size of the munitions was dramatically reduced because of the new controller design capability. Precise automatic control of the smart bombs angle-of-attack gives it the same penetration capabilities against hardened targets as current munitions that are 10 times larger.

In 1988 and 89, under AFOSR support, Dr. Juri Medanic and Dr. William Perkins, professors at UIUC, developed several new control design methodologies that simultaneously satisfied a diverse set of control requirements including transient performance, disturbance rejection, robustness and reliability. One method they developed, now known as ¡I¿projective control theory¡/I¿, is based on finding suitable projections from high-order complex optimal designs to much simpler, low-order controller designs, which achieve near optimal performance while meeting the design constraints.

Dr. Kevin Wise, at the Boeing Phantom Works in St. Louis, leveraged the AFOSR-sponsored research. He developed a revolutionary flight control design tool called AUTOGAIN that completely automates autopilot design over the entire flight envelope. The enabling technology in AUTOGAIN was the projective control theory, used to project optimal state feedback designs into output feedback designs, thus eliminating hardware sensors in the implementation. AUTOGAINs success has made it a ”best practice” in the Boeing Phantom Works. On many projects it has cut controller design effort by at least 50 percent. Besides the JDAM and MMTD, several other programs use the AUTOGAIN system including the Tomahawk cruise missile, the BQM-74 (Navy target drone), and the $4th$ Generation Escape System program.

Under AFOSR basic research support, Dr. Wise augmented these design tools in an important way. He developed new control system robustness analysis algorithms that have proved instrumental in analyzing a flight control systems dependency on knowing uncertain aerodynamic parameters. This analysis problem was of critical importance on the $4th$ Generation Escape System program where the mass properties and aerodynamic characteristics significantly vary between the 95 percent male pilot and 5 percent female pilot population. Wises algorithms were used to accurately determine that the ejections seat flight control system would perform over this wide range of parameters, culminating in the first ever supersonic ejection seat flight test at Holloman AFB, N.M.

Vignette: Uncertainty Management for Complex Systems

Researchers at Caltech, with their colleagues at UCSB and UCLA, have made dramatic progress in creating a fundamental new theory of uncertainty management in complex, multiscale systems, with a variety of applications from shear flow turbulence to networking protocols to global optimization. The results of this research will not only provide a rigorous basis for designing future networks of networks involving ubiquitous control, communications and computing, but is also resolving many persistent mysteries at the foundations of physics.

Two of the great abstractions of the 20th century were the separation, in both theory and applications, of 1) controls, communications, and computing from each other, and 2) the systems level from its underlying physical substrate. This horizontal and vertical isolation of systems facilitated massively parallel, wildly successful, explosive growth in both mathematical theory and technology, but left many fundamental problems unresolved and a poor foundation for future systems of systems in which these elements must be integrated. The unifying theme in this work is the new concept of Highly Optimized Tolerance (HOT) that arises when deliberate robust design aims for a specific level of tolerance to uncertainty. The resulting "robust, yet fragile" features of HOT systems are high performance and high throughput, but potentially high sensitivities to design flaws and unanticipated or rare events. HOT provides a framework in which the previously fragmented mathematical tools of robust control, communications, computation, dynamical systems, and statistical physics are unified and brought to bear on a variety of applications. For example, congestion due to bursty Internet traffic can be traced to HOT design of web layouts and protocols, a generalization of source coding that suggests novel new protocol designs. This is leading not only to better control of networks, but should facilitate distributed control of dynamical systems using networks. Similar insights have been obtained in domains as diverse as biological signal transduction and gene regulation, forest ecology, cascading failures in power grids, and financial market volatility.

What is perhaps surprising is how the HOT framework developed in this MURI Center is resolving many persistent mysteries at the foundation of physics where interconnected, multiscale systems issues arise. Promising examples with entirely new and novel theories include the ubiquity of power laws in natural and man-made systems, the nature of shear flow turbulence, the origin of dissipation and thermodynamic irreversibility, and the quantum/classical transition and quantum measurement. The most well developed HOT system in physics is a fundamentally new view of turbulence in the highly sheared flows that results from design for drag minimization. A key result is that the Navier-Stokes equation with external forcing exhibits radically different structure from the unforced equation. In particular, slight perturbations to the laminar solution are amplified by the fluid dynamics, and this perturbation energy grows as the cube of the Reynolds number. This suggests that one of the factors in transition is the transfer of energy from the mean flow to the eddy fields with fine-scale perturbations amplified to create large-scale vortical structures. Robustness analysis and model reduction can thus be used to analyze the flows and simplify computation, giving for the first time a global theoretical view of coherent structures in shear flow turbulence.

This view of turbulence is being connected with the development of a new class of subgrid scale models that uses a methodology combining volume-preserving diffeomorphism group techniques, asymptotic expansions of stochastic processes, and averaging of the variational principle. The resulting models, the Lagrangian averaged Navier-Stokes equations (LANS), have a natural closure, and provide novel numerical algorithms that remove energy content from the small subgrid scales, while maintaining the crucial features of the large scale flow using dispersive rather than dissipative mechanisms. While this work is a very new and radically different view of turbulence, computational and analytical tools from it are already competitive with traditional methods with decades of research behind them. The future prospects for both computation and control of fluids is truly revolutionary.

Military and commercial technological visions emphasize ubiquitous control, communications, and computing, with both biology and nanotechnology creating additional novel multiscale challenges. A rigorous, practical, and unified theoretical framework will be essential for this vision, but until this work, a solution has proven stubbornly elusive. This research offers not only a theoretical research direction of unprecedented promise, but one that has already proven remarkably useful in a wide variety of practical applications.

Vignette: Tunable High Resolution Spectral Estimators

The 1980's witnessed a flurry of activity in the area of signal processing with the advent of nonlinear methods for spectral analysis. These have provided tools for a range of applications from speech analysis to spectroscopy and sensor array processing. The advantage of nonlinear methods over traditional linear Fourier analysis has been their ability to beat the uncertainty time-bandwidth limit which is inherent in Fourier theory. This bound which limits resolution, is often compared to the quantum mechanical uncertainty principle. However, in the realm of signal analysis, processing does not destroy the measurements and the bound is only a limitation of linear analysis tools. Hence, the modern techniques took advantage of that niche. However, it soon became apparent that the new techniques were no panacea, and that when applied to short observation records of very noisy data, they were not sufficiently robust and reliable. Hence in many application areas, such as in synthetic aperture imaging, Fourier methods remained the workhorse of the industry.

A breakthrough and a sequel to the modern nonlinear methods came about from the combined efforts of researchers at the University of Minnesota and the Washington University in St. Louis, led by Christopher I. Byrnes, Tryphon T. Georgiou and Anders Lindquist, and supported by AFOSR grants. The mathematics underlying the new developments have deep roots in analytic function theory and opened up the field for new techniques with remarkably high resolution as well as robustness, and hence, are highly relevant to several areas of interest to the Air Force.

For example, in SAR imaging, a signal illuminates a terrain and the echo contains information on the position and velocity of objects in the field of view. A typical SAR image of a field, with an armored personnel carrier present, is shown in Figure 1, while a photograph of the APC is shown in Figure 2. Given the reflected echo, the goal is to resolve nearby scatterers which in this case include edges on the vehicle, with the maximal possible accuracy and reliability, for targeting and identification purposes.

The basis of the Byrnes-Georgiou-Lindquist idea came about with the discovery that non-traditional "tunable designerstatistics" of the echo can amplify the effect of scatterers within any specified range, and thereby, reveal more accurately their presence. With a combination of analytic function theory and optimization, they then showed how such statistics can provide models for the profile of the scattering field. For instance, Figure 3 depicts the highlighted area in Figure 1 on the left, and a reconstruction with new high resolution methods on the right. Figure 4 also compares edge detection based on these images; typically used in automatic object recognition. The new techniques represent something like a "magnifying glass" which can be pointed to any area of interested and is capable of a dramatic improvement in resolution and dynamic range.

The discovery of "designer-statistics" and the development of the pertinent analytic theory by the Byrnes-Georgiou-Lindquist group has opened up a new chapter in high resolution spectral analysis. Besides a range of applications in SAR/ISAR, Ladar, micro-Doppler, etc., which are of direct interest to the Air Force mission, these advances are expected to yield new technology in medical imaging (ultrasound, MR, PET) as well as more accurate models for speech coding in the context of mobile digital communication.

Vignette: A sound approach to imaging

It has been long known that interaction between electromagnetic fields and various material excitations can occur. Brillouin scattering results from the coupling of electromagnetic fields and acoustic waves in matter. For the first time in the history of science, as far as we know, researchers working with the Air Force Office of Scientific Research are exploiting this electromagnetic-acoustic coupling to attempt to extend Air Force imaging capabilities.

In a real material medium such as foliage, wood, soil, concrete, or plastic, the material polarizability and material conductivity are dependent on the local pressure. Material polarizability refers to the degree to which fixed charge in the material can be distorted by an electric field. A measure of material polarizability is the dielectric constant. Material conductivity refers to the tendency of free charges within the material to move in the presence of an electric field. While polarizability and conductivity are functions of local pressure, these material properties also influence electromagnetic wave speed in the medium. Thus, there is a connection between pressure and the movement of electromagnetic energy, and, it is this interaction that is being exploited to enhance imaging.

Professor Thomas Banks and colleagues at the Center for Scientific Computation at NC State University, under AFOSR sponsorship, have performed computational research that indicates that this electromagnetic-acoustic (pressure) interaction can be exploited to help identify materials. Banks and colleagues have worked at microwave frequencies and have shown that an acoustic wave front can act as a weakly reflective moving mirror within a medium such as soil. Using short microwave pulses and exploiting the vast difference between electromagnetic speeds in materials and acoustic speeds, acoustic wave fronts can be interrogated as they spread in a structure. Interrogating the acoustic or pressure fronts appears to be able to aid in the characterization of the material between the surface of the object and the "collision point" of the acoustic wave front and the microwave pulse.

This research is aiming at exciting military applications. Specifically, it is expected that underground bunkers will have air conditioning units or motorized airflow systems. These constitute a localized acoustic source that may be located using microwave scattering off of the outwardly moving acoustic waves. The research will also investigate whether vehicles with motors running have acoustic signatures that will permit identification in settings of high clutter. Finally, the earth has naturally occurring vibratory events. This research will ask whether scattering of these naturally occurring events can aid in the imaging of ground and sub-surface targets.

References

[1] T. P. Hughes. *Elmer Sperry Inventor and Engineer.* The John Hopkins University Press, London, 1993.

[2] M. W. ed. McFarland. *The Papers of Wilbur and Orville Wright.* Mc Graw Hill, New York, 1953.

```
Additional references:
  Fleming + previous IEEE report
  Kumar, Brockett, etc (future directions)
  EC report + other government reports; NSF education workshop
  1952 Scientific American issues
  Out of Control and other popular books (Wiener)
  http://www.dfrc.nasa.gov/History/Publications/f8ctf/contents.html
  http://www.theorem.net/theorem/lewis1.html
```

Index