

Future Directions in Control in an Information-Rich World

Richard M. Murray
Control and Dynamical Systems
California Institute of Technology
Pasadena, CA 91125

Karl J. Åström
Department of Automatic Control
Lund Institute of Technology
Lund, Sweden

Stephen P. Boyd
Electrical Engineering
Stanford University
Stanford, CA

Roger W. Brockett
Engineering and Applied Sciences
Harvard University
Cambridge, MA

Gunter Stein
Honeywell Laboratories
Honeywell International
Minneapolis, MN

January 4, 2003

Introduction

The Panel on Future Directions in Control, Dynamics, and Systems was formed in April 2000 to provide a renewed vision of future challenges and opportunities in the field, along with recommendations to government agencies, universities, and research organizations to ensure continued progress in areas of importance to the industrial and defense base. Its report was released in April 2002 and will be published by SIAM [1]. The intent of the 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 Prof. Roger Brockett (Harvard), Prof. John Burns (VPI), Prof. John Doyle (Caltech) and Dr. Gunter Stein (Honeywell). The remaining panel members are Karl Åström (Lund Institute of Technology), Siva Banda (Air Force Research Lab), Stephen Boyd (Stanford), Munzer Dahleh (MIT), John Guckenheimer (Cornell), Charles Holland (DDR&E), Pramod Khargonekar (University of Florida), P.R. Kumar (University of Illinois), P.S. Krishnaprasad (University of Maryland), Greg McRae (MIT), Jerrold Marsden (Caltech), George Meyer (NASA), William Powers (Ford), and Pravin Varaiya (UC Berkeley). A writing subcommittee consisting of Karl Åström, Stephen Boyd, Roger Brockett, John Doyle, Richard Murray, and Gunter Stein helped coordinate the generation of the report.

The Panel held a meeting on 16-17 July 2000 at the University of Maryland, College Park, to discuss the state of the field and its future opportunities. The meeting was attended by members of the panel and invited participants from 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 among panel members and within the community at workshops and conferences, form the main basis for the findings and recommendations of the 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 the report, links to other sources of information, and presentation materials from the Panel workshop and other meetings can be found there.

Several similar reports and papers highlighting future directions in control came to the Panel's attention during the development of the report. Many members of the panel and participants in the June 2000 workshop were involved in the generation of the 1988 Fleming report [2] and a 1987 *IEEE Transactions on Automatic Control* article [3], both of which provided a road map for many of the activities of the last decade and continue to be relevant. More recently, the European Commission sponsored a workshop on future control systems [4], and several other more focused workshops have been held over the last several years [5], [6], [7], [8]. Several recent papers and reports highlighted successes of control [9] and new vistas in control [10], [11]. The Panel also made extensive use of a recent NSF/CSS report on future directions in control engineering education [6].

The bulk of the report was written before the tragic events of September 11, 2001, but control will clearly play a major role in the world's effort to counter terrorism. From new methods for command and control of unmanned vehicles to robust networks linking businesses, transportation systems, and energy infrastructure and to improved techniques for sensing and detection of biological and chemical agents, the techniques and insights from control will enable new methods for protecting human life and safeguarding our society.

Report 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. The report presents the findings and recommendations of a panel of experts chartered to examine these opportunities. It presents an overview of the field, reviews its successes and impact, and describes the new challenges ahead. The report does not attempt to cover the entire field. Rather, it focuses on those areas that are undergoing the most rapid change and that require new approaches to meet the challenges and opportunities that face the community.

Overview of Control

Control as defined in the report refers to the use of algorithms and feedback in engineered systems. At its simplest, a control system is a device in which a sensed quantity is used to modify the behavior of a system through computation and actuation. Control systems engineering traces its roots to the industrial revolution, to devices such as the flyball governor, shown in Fig. 1. This device used a flyball mechanism to sense the rotational speed of a steam turbine and adjust the flow of steam into the machine using a series of linkages. By thus regulating the turbine's speed, it provided the safe, reliable, consistent operation that was required to enable the rapid spread of steam-powered factories.

Control played an essential role in the development of technologies such as power, communications, transportation, and manufacturing. Examples include autopilots in military and commercial aircraft (Fig. 2a), regulation and control of the electrical power grid, and high-accuracy positioning of read/write heads in disk drives (Fig. 2b). Feedback is an enabling technology in a variety of application areas and has been reinvented and patented many times in different contexts.

A modern view of control sees feedback as a tool for uncertainty management. By measuring the operation of a system, comparing it to a reference, and adjusting available control variables,




Figure 1: The centrifugal governor (a), developed in the 1780s, was an enabler of the successful Watt steam engine (b), which fueled the industrial revolution. Figures courtesy of Cambridge University.

we can cause the system to respond properly even if its dynamic behavior is not exactly known or if external disturbances tend to cause it to respond incorrectly. This is an essential feature in engineering systems since they must operate reliably and efficiently under a variety of conditions. It is precisely this aspect of control as a means of ensuring robustness to uncertainty that explains why feedback control systems are all around us in the modern technological world. They are in our homes, cars and consumer electronics, in our factories and communication systems, and in our transportation, military and space systems.

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

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

Successes and Impact

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

- Guidance and control systems for aerospace vehicles, including commercial aircraft, guided missiles, advanced fighter aircraft, launch vehicles, and satellites. These control systems provide stability and tracking in the presence of large environmental and system uncertainties.
- Control systems in the manufacturing industries, from automotive to integrated circuits.




Figure 2: Applications of control: (a) the Boeing 777 fly-by-wire aircraft and (b) the Seagate Barracuda 36ES2 disk drive. Photographs courtesy of the Boeing Company and Seagate Technology.

Computer-controlled machines provide the precise positioning and assembly required for high-quality, high-yield fabrication of components and products.

- Industrial process control systems, particularly in the hydrocarbon and chemical processing industries. These maintain high product quality by monitoring thousands of sensor signals and making corresponding adjustments to hundreds of valves, heaters, pumps, and other actuators.
- Control of communication systems, including the telephone system, cell phones, and the Internet. Control systems regulate the signal power levels in transmitters and repeaters, manage packet buffers in network routing equipment, and provide adaptive noise cancellation to respond to varying transmission line characteristics.

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

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

Future Opportunities and Challenges

As we look forward, the opportunities for new applications that will build on advances in control expand dramatically. The advent of ubiquitous, distributed computation, communication, and sensing systems has begun to create an environment in which we have access to enormous amounts of data and the ability to process and communicate that data in ways that were unimagined 20 years ago [12]. This will have a profound effect on military, commercial, and scientific applications, especially as software systems begin to interact with physical systems in increasingly integrated ways. Fig. 3 illustrates two systems where these trends are already evident. Control will be an increasingly essential element of building such interconnected systems, providing high performance, high confidence, and reconfigurable operation in the presence of uncertainties.

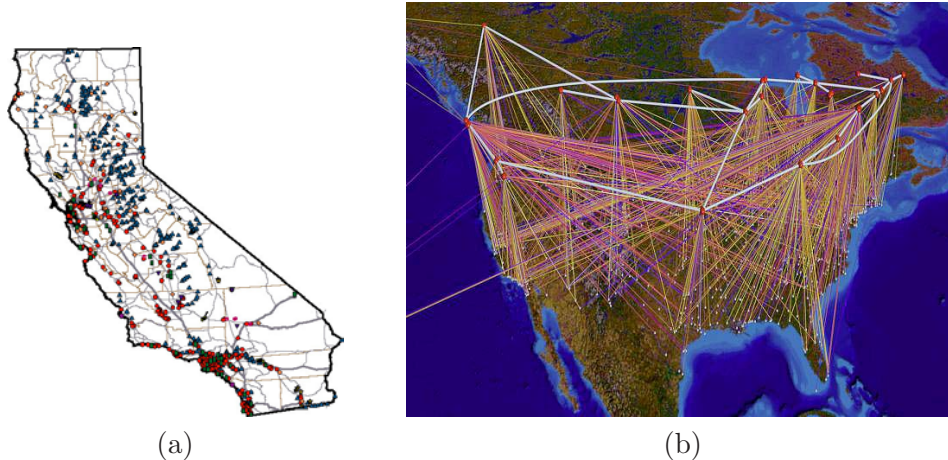


Figure 3: Modern networked systems: (a) the California power network and (b) the NSFNET Internet backbone. Figures courtesy of the state of California and the National Center for Super-computer Applications (NCSA).

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

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

Control is a critical technology in defense systems and is increasingly important in the fight against terrorism and asymmetric threats. Control allows the operation of autonomous and semi-autonomous unmanned systems that keep people out of harm's way, as well as sophisticated command and control systems that enable robust, reconfigurable decision-making systems. The use of control in microsystems and sensor webs will improve our ability to detect threats before they cause damage. And new uses of feedback in communication systems will provide reliable, flexible, and secure networks for operation in dynamic, uncertain, and adversarial environments.

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

- *Control of systems with both symbolic and continuous dynamics.* Next-generation systems will combine logical operations (such as symbolic reasoning and decision making) with continuous quantities (such as voltages, positions, and concentrations). The current theory is not well-tuned for dealing with such systems, especially as we scale to very large systems.

- *Control in distributed, asynchronous, networked environments.* Control distributed across multiple computational units, interconnected through packet-based communications, will require new formalisms for ensuring stability, performance, and robustness. This is especially true in applications where one cannot ignore computational and communication constraints in performing control operations.
- *High-level coordination and autonomy.* Increasingly, feedback is being designed into enterprise-wide decision systems, including supply chain management and logistics, airspace management and air traffic control, and military command and control (C²) systems. The advances of the last few decades in analysis and design of robust control systems must be extended to these higher level decision-making systems if they are to perform reliably in realistic settings.
- *Automatic synthesis of control algorithms, with integrated validation and verification.* Future engineering systems will require the ability to rapidly design, redesign, and implement control software. Researchers need to develop much more powerful design tools that automate the entire control design process from model development to hardware-in-the-loop simulation, including system-level software verification and validation.
- *Building very reliable systems from unreliable parts.* Most large engineering systems must continue to operate even when individual components fail. Increasingly, this requires designs that allow the system to automatically reconfigure itself so that its performance degrades gradually rather than abruptly.

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

Application Areas

The Panel decided to organize the treatment of applications around five main areas to identify the overarching themes that would guide its recommendations. These five areas are aerospace and transportation, information and networks, robotics and intelligent machines, biology and medicine, and materials and processing. In addition, several additional areas arose over the course of the Panel's deliberations, including environmental science and engineering, economics and finance, and molecular and quantum systems. Taken together, the areas represent an enormous collection of applications and demonstrate the breadth of applicability of ideas from control.

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

Several overarching themes arose across all of the areas considered by the Panel. The use of systematic and rigorous tools is considered critical to future success and is an important trademark of the field. At the same time, the next generation of problems will require a paradigm shift in



Figure 4: The F-18 aircraft, one of the first production military fighters to use “fly-by-wire” technology, and the X-45 (UCAV) unmanned aerial vehicle. Photographs courtesy of NASA Dryden Flight Research Center.

control research and education. The increased information available across all application areas requires more integration with ideas from computer science and communications, as well as improved tools for modeling, analysis, and synthesis for complex decision systems that contain a mixture of symbolic and continuous dynamics. The need to continue research in the theoretical foundations that will underly future advances was also common across all of the applications.

In each subsection that follows, we briefly summarize the challenges in the subject area. More information is available in the full report.

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 significant part of the modern world’s overall technological capability. They are also a major part of its economic strength, and they contribute greatly to the well being of its people. The historical role of control in these application areas, the current challenges in these areas, and the projected future needs all strongly support the recommendations of the report.

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.

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




Figure 5: UUNET network backbone for North America. Figure courtesy of WorldCom.

controlled operation of all major functions, including fuel injection, emission control, cruise control, braking, and cabin comfort.

Despite its many successes, the control needs of some engineered systems today and those of many in the future outstrip the power of current tools and theories. 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 that provide logical regulation of operating modes, vehicle configurations, payload configurations, health status, etc. [13]. For aircraft, these functions are collectively called “vehicle management.” They have historically been performed by pilots or other human operators, but today that boundary is moving, and control systems are increasingly taking on these functions.

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

Another 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 the laws of physics, and no possibility of imposing centralized control actions. Examples of this trend include the national airspace management problem, automated highway and traffic management, and the command and control for future battlefields.

Information and Networks

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

Control of networks is a large area, spanning many topics, including congestion control, routing,

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

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

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

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 communication networks that guarantee delivery of sensor, actuator, and other signals with a known, fixed delay. Although 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 whether we can develop a theory and practice for control systems that operate in a distributed, asynchronous, packet-based environment.



Figure 6: The Mars Sojourner and Sony AIBO robots. Photographs courtesy of Jet Propulsion Laboratory and Sony.

Robotics and Intelligent Machines

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

The goal of cybernetic engineering, already articulated in the 1940s and even before, has been to implement systems capable of exhibiting highly flexible or “intelligent” responses to changing circumstances. In 1948, the MIT mathematician Norbert Wiener gave a widely read, albeit completely nonmathematical, account of cybernetics [14]. A more mathematical treatment of the elements of engineering cybernetics was presented by H.S. Tsien in 1954, driven by problems related to control of missiles [15]. Together, these works and others of that time form much of the intellectual basis for modern work in robotics and control.

Two accomplishments that demonstrate the successes of the field are the Mars Sojourner robot and the Sony AIBO robot, shown in Fig. 6. Sojourner successfully maneuvered on the surface of Mars for 83 days starting in July 1997 and sent back live pictures of its environment. The Sony AIBO robot debuted in June of 1999 and was the first “entertainment” robot to be mass marketed by a major international corporation. It was particularly noteworthy because of its use of AI technologies that allowed it to act in response to external stimulation and its own judgment.

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

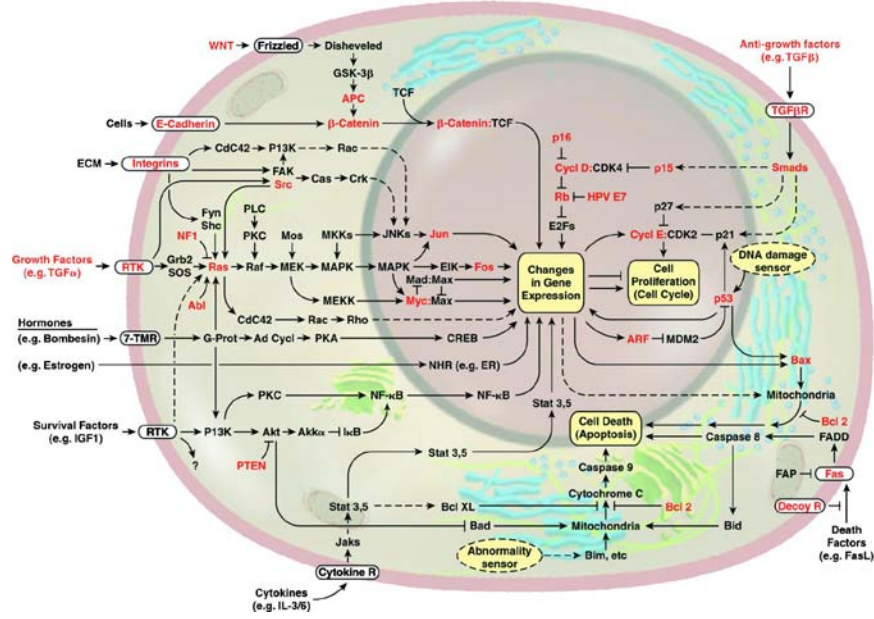


Figure 7: The wiring diagram of the growth signaling circuitry of the mammalian cell [16].

Despite the enormous progress in robotics over the last half century, the field is very much in its infancy. Today’s robots still exhibit extremely simple behaviors compared with humans, and their ability to locomote, interpret complex sensory inputs, perform higher level reasoning, and cooperate together in teams is limited. Indeed, much of Wiener’s vision for robotics and intelligent machines remains unrealized. While advances are needed in many fields to achieve this vision—including advances in sensing, actuation, and energy storage—the opportunity to combine the advances of the AI community in planning, adaptation, and learning with the techniques in the control community for modeling, analysis, and design of feedback systems presents a renewed path for progress. This application area is strongly linked with the Panel’s recommendations on the integration of computing, communication and control, development of tools for higher level reasoning and decision making, and maintaining a strong theory base and interaction with mathematics.

Biology and Medicine

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

A major theme identified by the Panel was the science of reverse (and eventually forward) engineering of biological control networks. There are a wide variety of biological phenomena that provide a rich source of examples for control, including gene regulation and signal transduction; hormonal, immunological, and cardiovascular feedback mechanisms; muscular control and locomotion; active sensing, vision, and proprioception; attention and consciousness; and population dynamics




Figure 8: Intel Pentium IV wafer and die.

and epidemics. Each of these (and many more) provide opportunities to figure out what works, how it works, and what we can do to affect it.

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

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

Materials and Processing

The chemical industry is among the most successful industries in the United States, producing \$400 billion of products annually and providing over one million U.S. jobs. Having recorded a trade surplus for 40 consecutive years, it is the country's premier exporting industry: chemical industry exports totaled \$72.5 billion in 2000, accounting for more than 10% of all U.S. exports, and generated a record trade surplus in excess of \$20 billion in 1997.

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

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

phenomena.

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

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

Other Areas

The previous sections have described some of the major application areas discussed by the Panel. However, there are many more areas where ideas from control are being applied or could be applied. The report describes additional opportunities and challenges in the following areas:

- Environmental science and engineering, particularly atmospheric systems and microbiological ecosystems
- Economics and finance, including problems such as pricing and hedging options
- Electromagnetics, including active electromagnetic nulling for stealth applications
- Molecular, quantum, and nanoscale systems, including design of nanostructured materials, precision measurement, and quantum information processing
- Energy systems, including load distribution and power management for the electrical grid

Education and Outreach

Control education is an integral part of the community's activities and one of its most important mechanisms for transition and impact. In 1998, the National Science Foundation (NSF) and the IEEE Control Systems Society (CSS) jointly sponsored a workshop in control engineering education that made a number of recommendations for improving control education (see [6]). The Panel's findings and recommendations are based on that report and on discussions between Panel members and the control community.

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

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

In addition, control is increasingly being applied outside of its traditional domains in aeronautics, chemical engineering, electrical engineering, and mechanical engineering. Biologists are using ideas from control to model and analyze cells and animals; computer scientists are applying control to the design of routers and embedded software; physicists are using control to measure and modify the behavior of quantum systems; and economists are exploring the applications of feedback to markets and commerce.

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

For the opportunities described in the report to be realized, control education must be restructured to operate in this new environment. Several universities have begun to make changes in the way that control is taught and organized, and these efforts provide some insights into how this restructuring might be done successfully. Current approaches include establishing a cross-disciplinary research center in control involving researchers from many different disciplines, establishing a shared graduate curriculum in control across multiple departments, and establishing M.S. or Ph.D. programs in control (common in Europe).

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

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

very little background in continuous mathematics beyond a sophomore course in scalar ordinary differential equations and linear algebra.

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

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

Recommendations

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

Integrated Control, Computation, Communications

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

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

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

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

In the United States, the Department of Defense has already made substantial investment in this area through the Multidisciplinary University Research Initiative (MURI) program, and this trend should be continued. It will be important to create larger, multidisciplinary 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 communication networks (wired, wireless, and cellular), and enterprise-wide operations and supply networks (electric power, manufacturing, service and repair).

Control of Complex Decision Systems

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

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

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

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

The extension of control beyond its traditional roots in differential equations is an area 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 so as to understand how to formulate the problem.

A useful technique may be the development of testbeds to explore new ideas. In the military arena, these testbeds could consist of collections of unmanned vehicles (air, land, sea, and space), operating in conjunction with human partners and adversaries. In the commercial sector, service robots and personal assistants may be a fruitful area for exploration. And in a university setting, the emergence of robotic competitions is an interesting trend that control researchers should explore as a mechanism for developing new paradigms and tools. In all of these cases, stronger links with the AI community should be explored, since that community is currently at the forefront of many of these applications.

The benefits of research in this area include replacing ad hoc design methods with 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.

High-Risk, Long-Range Applications of Control

The potential application areas for control are increasing rapidly as advances in science and technology develop new understanding of the importance of feedback and new sensors and actuators allow manipulation of heretofore unimagined detail. To discover and exploit opportunities in these

new domains, experts in control must actively participate in new areas of research outside of their traditional roots. At the same time, we must find ways to educate domain experts about control, to allow a fuller dialog, and to accelerate the uses of control across the enormous number of possible applications.

In addition, many applications will require new paradigms for thinking about control. For example, our traditional notions of signals that encode information through amplitude and phase relationships may need to be extended to allow the study of systems where pulse trains or biochemical “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 new domains such as nanotechnology, quantum mechanics, electromagnetics, biology, and environmental science.

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

Another need is to establish “meeting places” where control researchers can join with new communities and each can develop an understanding of the principles and tools of the other. This could include focused workshops of a week or more to explore control applications in new domains or four- to six-week short courses on control that are tuned to a specific application area, with tutorials in that application area as well.

At universities, new materials are needed to teach nonexperts 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 Center for Control Engineering and Computation at Santa Barbara) and programs in control (such as the Control and Dynamical Systems 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 previous attempts have not always been successful, but the opportunities are great, and we must continue to strive to move forward.

Support for Theory and Interaction with Mathematics

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

An ongoing need is to make the existing knowledge base more compact so that the field can continue to grow. Integrating previous results and providing a more unified structure for understanding and applying those results is necessary in any field and has occurred 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. Control experts also need to expand the applications base by having the appropriate level of abstraction to identify new applications of existing theory.

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

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

Some possible areas of interaction include dynamical systems, graph theory, combinatorics, complexity theory, queuing theory, and statistics. Additional perspectives on the interaction of control and mathematics can be found in a recent survey article by Brockett [10].

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

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

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

New Approaches to Education and Outreach

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

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

invest in new approaches to education and outreach for the dissemination of control concepts and tools to nontraditional audiences.

As a first step toward implementing this recommendation, new courses and textbooks should be developed for both experts and nonexperts. 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 bioengineering. It is also important that these courses emphasize the *principles* of control rather than simply providing tools that can be used in a given domain.

An important element of education and outreach is the continued use of experiments and the development of new laboratories and software tools. These are much easier to do than ever before and also more important. Laboratories and software tools should be integrated into the curriculum, including moving beyond their current use in introductory control courses to increased use in advanced (graduate) course work. The importance of software cannot be overemphasized, both in terms of design tools (e.g., MATLAB toolboxes) and implementation (real-time algorithms).

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

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

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

Concluding Remarks

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

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

The payoffs for investment in control research are substantial. They include the successful development of systems that operate reliably, efficiently, and robustly; new materials and devices that are made possible through advanced control of manufacturing processes; and increased understanding of physical and biological systems through the use of control principles. Perhaps most important is the continued development of individuals who embrace a systems perspective and provide technical leadership in modeling, analysis, design, and testing of complex engineering systems.

Acknowledgments

The authors would like to the Panel on Future Directions in Control, Dynamics, and Systems and the many members of the control community who participated in generating the report. We would also like to thank Dr. Marc Q. Jacobs and the Air Force Office of Scientific Research (AFOSR) for supporting the activities of the Panel.

References

- [1] R. M. Murray, editor. *Control in an Information Rich World: Report of the Panel on Future Directions in Control, Dynamics, and Systems*. SIAM, 2002. To appear; available at <http://www.cds.caltech.edu/~murray/cdspanel>.
- [2] W. H. Fleming, editor. *Future Directions in Control Theory: A Mathematical Perspective*. SIAM, 1988.
- [3] A. H. Levis, S. I. Marcus, W. R. Perkins, P. Kokotovic, M. Athans, R. W. Brockett, and A. S. Willsky. Challenges to control: A collective view. *IEEE Transactions on Automatic Control*, AC-32(4):275–285, 1987.
- [4] European Commission. Workshop on future and emerging control systems, November 2000. Available at ftp://ftp.cordis.lu/pub/ist/docs/ka4/report_controlws.pdf.
- [5] M. Athans, J. J. S. Sentieiro, and M. A. Dahleh. Workshop on future directions in systems and control theory. Cascais, Portugal, June 2000.
- [6] P. Antsaklis, T. Basar, R. DeCarlo, N. H. McClamroch, M. Spong, and S. Yurkovich, editors. *NSF/CSS Workshop on New Directions in Control Engineering Education*. National Science Foundation and IEEE Control Systems Society, 1998. Available at <http://robot0.ge.uiuc.edu/~spong/workshop>.
- [7] National Science Foundation. NSF workshop for high school teachers of mathematics and science. American Control Conference, June 2000. Available at http://www.math.ukans.edu/ksacg/nsf_hswkshop.html.
- [8] National Science Foundation. NSF workshop on dynamic data-driven application systems, March 2000. Available at <http://www.cise.nsf.gov/eia/dddas/>.
- [9] H. Nijmeijer and J. M. Schumacher. Four decades of mathematical system theory. In J. W. Polderman and H. L. Trentelman, editors, *The Mathematics of Systems and Control: From Intelligent Control to Behavioral Systems*, pages 73–83. Univ. of Groningen, 1999.

- [10] R. W. Brockett. New issues in the mathematics of control. In B. Engquist and W. Schmid and W. Schmid, editors, *Mathematics Unlimited—2001 and Beyond*, pages 189–220. Springer Verlag, 2000.
- [11] P. R. Kumar. New technological vistas for systems and control: The example of wireless networks. *Control Systems Magazine*, 21(1):24–37, 2001.
- [12] National Research Council. *Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers*. National Academy Press, 2001.
- [13] S. Banda, J. C. Doyle, R. M. Murray, J. Paduano, J. Speyer, and G. Stein. Research needs in dynamics and control for uninhabited aerial vehicles. Panel Report, November 1997. <http://www.cds.caltech.edu/~murray/notes/uav-nov97.html>.
- [14] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. John Wiley, 1948.
- [15] H. S. Tsien. *Engineering Cybernetics*. McGraw-Hill, 1954.
- [16] D. Hanahan and R. A. Weinberg. The hallmarks of cancer. *Cell*, 100:57–70, 2000.