http://www.cds.caltech.edu/~murray/papers/2003l_mur03-ejc.html
Submitted, European Journal of Control

Future Directions in Control, Dynamics, and Systems: Overview, Grand Challenges, and New Courses

Richard M. Murray
Control and Dynamical Systems
California Institute of Technology
Pasadena, CA USA
murray@cds.caltech.edu

Submitted, *European Journal of Control*
31 March 2003

Abstract

This paper summarizes the findings and recommendations of a recent panel on Future Directions in Control, Dynamics, and Systems, sponsored by the US Air Force Office of Scientific Research. A set of grand challenges that illustrate some of the recommendations and opportunities are provided. Finally, the paper describes two new courses being developed at Caltech that are aligned with the recommendations of the report.

1 Introduction

In April 2000, a Panel on Future Directions in Control and Dynamical Systems was formed to prepare a report 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 report was completed in April 2002 and is being published by SIAM [8]. 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 which are important for continued vitality of the field.

Panel Composition and Activities

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

The Panel held a meeting on 16-17 July 2000 at the University of Maryland, College Park to discuss the state of the field and its future opportunities. The meeting was attended by members of the Panel and invited participants from 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 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 the report, links to other sources of information, and presentation materials from the Panel workshop and other meetings can be found there. A final printed copy of the report will be published by SIAM under the title "Control in an Information Rich World: Report of the Panel on Future Directions in Control, Dynamics and Systems" [8].

Overview of the Panel Report

The field of *control* provides the principles and methods used to design engineering 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. Control systems now play critical roles in many fields, including manufacturing, electronics, communications, transportation, computers and networks, and many military systems.

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

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

Based on an analysis of the opportunities in these application areas, the Panel recommended that the following actions be taken [8]:

1. Substantially increase research aimed at the *integration* of control, computer science, communications, and networking. This includes principles, methods and tools for modeling and control of high level, networked, distributed systems, and rigorous techniques for reliable, embedded, real-time software.

2. Substantially increase research in control at higher levels of decision making, moving toward enterprise level systems. This includes work in dynamic resource allocation in the presence of uncertainty, learning and adaptation, and artificial intelligence for dynamic systems.

Figure 1: 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.

3. Explore high-risk, long-range applications of control to new domains such as nanotechnology, quantum mechanics, electromagnetics, biology, and environmental science. Dual investigator, interdisciplinary funding might be a particularly useful mechanism in this context.
4. Maintain support for theory and interaction with mathematics, broadly interpreted. The strength of the field relies on its close contact with rigorous mathematics, and this will be increasingly important in the future.
5. Invest in new approaches to education and outreach for the dissemination of control concepts and tools to non-traditional audiences. The community must do a better job of educating a broader range of scientists and engineers on the principles of feedback and the use of control to alter the dynamics of systems and manage uncertainty.

In the remainder of this paper, we provide a brief overview of some of the findings of the Panel in selected application areas, provide some grand challenge problems to motivate future research, and give a summary of some educational activities at Caltech that are aligned with the Panel's recommendations in that area.

2 Applications, Opportunities and Challenges

The Panel organized its discussions into five main research areas and the findings are summarized here. Considerably more detail is available in the full report [8] and in a recent article in the *IEEE Control Systems Magazine* [9], from which the material in this section is drawn.

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

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 which provide logical regulation of operating modes, vehicle configurations, payload configurations, health status, etc [1, 10]. 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 laws of physics, and no possibility to impose centralized control actions. Examples of this trend include the national airspace management problem, automated highway and traffic management, and the command and control for future battlefields.

2.2 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

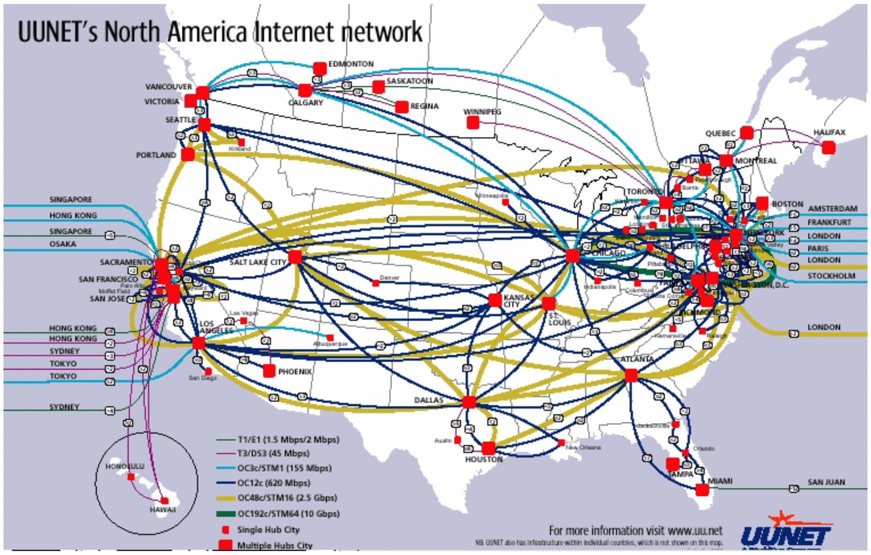

Figure 2: UUNET network backbone for North America. Figure courtesy WorldCom.

is probably the largest feedback control system man has ever built. Another is the decentralized nature of the control problem: local decisions must be made quickly, and based only on local information. Stability is complicated by the presence of varying time lags, as information about the network state can only be observed or relayed to controllers after a time delay, and the effect of a local control action can be felt throughout the network after substantial delay. Uncertainty and variation in the network, through network topology, transmission channel characteristics, traffic demand, available resources, etc. may change constantly and unpredictably. Another complicating 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 which controls our physical environment. Examples include automobiles, smart homes, large manufacturing systems, intelligent highways and networked city services, and enterprise-wide supply and logistics chains. Thus, this next phase of the information technology revolution is the convergence of communication, computing and control.

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

2.3 Robotics and Intelligent Machines

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

The goal of cybernetic engineering, already articulated in the 1940s and even before, has been to implement systems capable of exhibiting highly flexible or "intelligent" responses to changing circumstances. In 1948, the MIT mathematician Norbert Wiener gave a widely read, albeit completely non-mathematical, account of cybernetics [13]. 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 [11]. 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 Entertainment Robot, shown in Figure 3. Sojourner successfully maneuvered on the surface of Mars for 83 days starting in July 1997 and sent back live pictures of its environment. The Sony AIBO robot debuted in June of 1999 and was the first "entertainment" robot that was mass marketed by a major international corporation. It was particularly noteworthy because of its use of AI technologies that allowed it to act in response to external stimulation and its own judgment.

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

Figure 3: The Mars Sojourner and Sony AIBO Entertainment Robot. Photographs courtesy of Jet Propulsion Laboratory and Sony Electronics.

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.

2.4 Biology and Medicine

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

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

The Panel also identified potential roles for control in medicine and biomedical research. These included intelligent operating rooms and hospitals, from raw data to decisions; image guided surgery and therapy; hardware and soft tissue integration; fluid flow control for medicine and biological

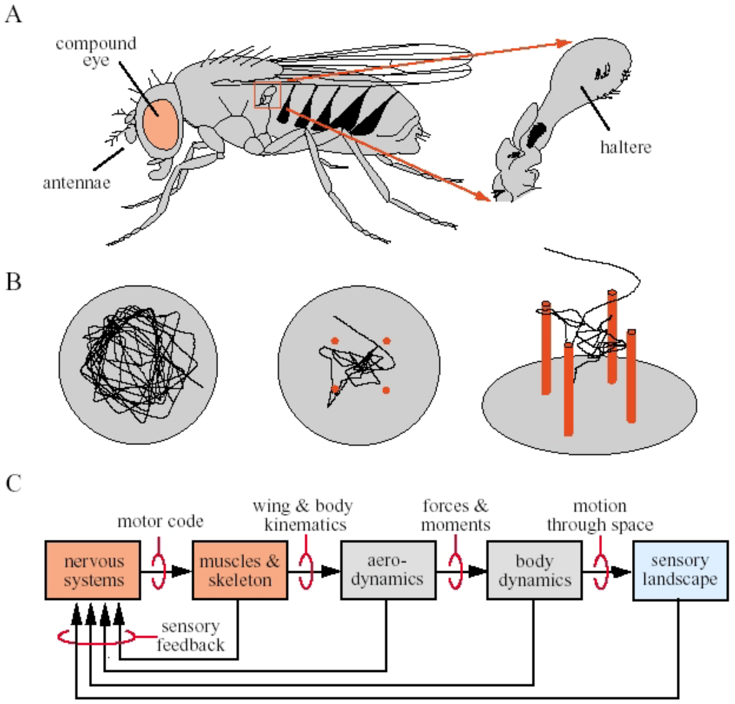

Figure 4: Overview of flight behavior in a fruit fly, *Drosophila*. (a) Cartoon of the adult fruit fly showing the three major sensor structures used in flight: eyes, antennae, and halteres (detect angular rotations). (b) Example flight trajectories over a 1 meter circular arena, with and without internal targets. (c) A schematic control model of the flight system. Figure and description courtesy of Michael Dickinson.

assays; and the development of physical and neural prosthesis. 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.

2.5 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 forty consecutive years, it is the country's premier exporting industry: chemical industry exports totaled \$72.5 billion in 2000, accounting for more than 10% of all U.S. exports, and generated a record trade surplus in excess of \$20 billion in 1997.

Process manufacturing operations will require a continual infusion of advanced information and process control technologies if the chemical industry is to maintain its global ability to deliver products that best serve the customer reliably at the lowest cost. In addition, a number of new technology areas are being explored that will require new approaches to control in order to be successful. These range from nanotechnology in areas such as electronics, chemistry, and biomaterials, to thin film processing and design of integrated microsystems, to supply chain management and

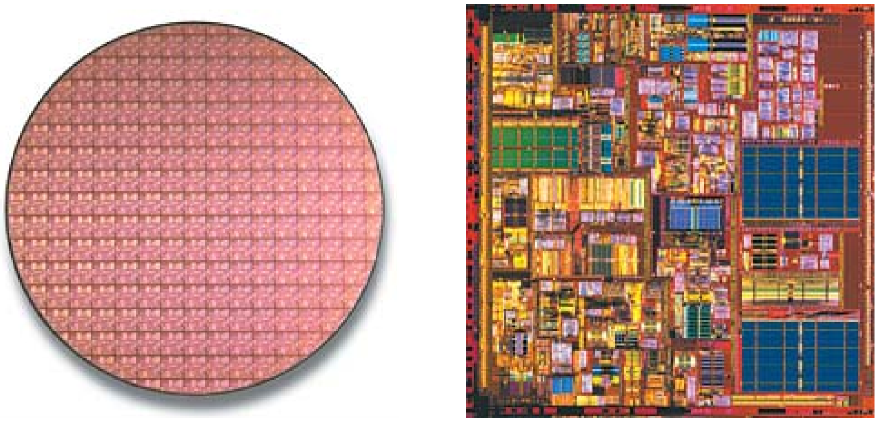

Figure 5: Intel Pentium IV wafer and die. Photographs courtesy of Intel.

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

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

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

3 Grand Challenges

To better motivate some of the research directions, the author has prepared a short list of grand challenge problems for consideration by the control community. These problems arose out of discussions at a recent NSF-sponsored workshop on modeling, dynamics, modeling and control of complex engineering systems [4] and subsequent discussions with members of the control community. Some additional ideas are given on the authors web site [7].

The intent of these grand challenges is to provide conceptual applications that can drive new thinking and new approaches to control. It is likely that none of these applications will ever come to fruition, but progress toward them will have many important uses in other areas and new grand challenges will arise from them.

Grand Challenge #1: World Cup Robotic Soccer Team

Design a robotic soccer team that is good enough that it can compete against humans in the World Cup (and win!). The robots should have the mass and volume of a human being. Successful techniques for this problem would also be of use in developing robotic search and rescue teams, security forces (police, firefighters, combat teams), and other collections of robots that perform cooperative tasks in unstructured environments.

The technical challenges in building such a system include agile motion control, distributed control and decision making, and team-based control strategies. There are many efforts underway to develop the underlying technologies, driven in part by the RoboCup project, which has the goal of solving this grand challenge problem by 2050[1]. Consistent with the findings of the Panel, progress in this area will require contributions not only from control, but also from the computer science, artificial intelligence and robotics communities. The relevance of control in this application domain is described nicely in a recent paper by D'Andrea [3].

Grand Challenge #2: InternetRT

Redesign the Internet so that it can be used to provide real-time (RT) connections between sensors, actuators, and computation that have arbitrary geographic locations. This could revolutionize the way we do control, perhaps even allowing control laws to reside anywhere on the network. It might also be useful for global service and supply chains and defense systems (e.g., missile defense).

The computer networking community has spent many years working on congestion control and related technologies, as described above and in the Panel report. For the most part, this work has been motivated by increasing the capacity of the network and, to a limited extent, increasing the real-time performance (for applications such as telephony and video conferencing). However, the use of networks for true closed loop control produces even tighter constraints on timing since delays can lead to instabilities that are not present in "one way" applications.

While this grand challenge is targeted at control of networks, it is clear that it will have to be combined with new research in control *over* networks as well. This is the topic of the next grand challenge.

Grand Challenge #3: Packet-Based Control Theory

Develop a theory for control in which the basic input/output singles are data packets that may arrive at variable times, not necessarily in order, and sometimes not at all. Related problems

---

[1] `http://www.robocup.org`

including figuring out how to do the source coding to support such networked control systems ("real-time information theory", in the words of Sanjoy Mitter [6]).

As described in Section 2.2, the current abstraction for signals in control theory assumes synchronized, guaranteed data delivery. One example (of many) where this is not true is in multi-vehicle systems, such as the Caltech Multi-Vehicle Wireless Testbed [2]. In this system, it is common to lose as many as 10% of the packets being sent across the network. While retransmission is possible, the delays associated with this (at least until InternetRT is developed) are prohibitively long.

Grand Challenge #4: Dynamically Reconfigurable Air Traffic Control

Design the air traffic control system so that passengers always get to their destination on time, with a plane that is always 90% or more full, and with no delays due to weather anyplace in the country except your departure or arrival city. Supply chains and the power grid could probably benefit from the resulting technology as well.

This challenge problem is special case of the more general class of resource allocation systems. In this context, control can be described as the science and engineering of optimal dynamic resource allocation under uncertainty [8]:

> We start with a mathematical model, of a system that describes how current actions or decisions can affect the future behavior of the system, including our uncertainty in that behavior. "Resource allocation" means that our decisions can be interpreted as managing a tradeoff between competing goals, or choosing from a limited set of possible actions. "Uncertainty" is critical: there is some possible variation in the system's behavior, so that decisions have to be made taking different possibilities into account.

Grand Challenge #5: Redesign the Feedback Control System of a Bacteria

Scientists are now able to genetically modify microbiological organisms so that they produce certain desired chemicals or change their behavior. Can we redesign the control systems in bacteria (including implementation!) so that we can program their behaviors in response to external stimuli? Possible applications include new types of medical treatments and in vivo sensing systems.

While this problem is out of the domain of expertise of many control researchers, it represents an enormous area of opportunity for the field. The design of the circuitry of a cell is fundamentally a control design problem, but using a very novel computational and signaling substrate. One of the challenges facing the field is to understand how it can help develop the understanding of control concepts that can be used by biologists and bioengineers to tackle and solve such problems. This is especially difficult in the United States, where undergraduate programs in biology often lack the mathematical background required to understand and use many of the tools and concepts of control.

Taken together, these five grand challenges attempt to describe some of the enormous opportunities for the field in areas ranging from aerospace, to biology, to communications. They are all very much aligned with the findings and recommendations of the report and, while sometimes a bit fanciful, they serve to provide a set of possible directions for researchers interested in control and its many applications.

4 CDS 101: Principles of Feedback and Control

As described in the Panel's recommendations, one of the challenges facing the field is making control more accessible to a broad range of scientists and engineers. This is a critical step in exploring new applications since the paradigm of communicating with domain experts through simplified models (usually represented as ODEs, "thrown over the wall") is not likely to be sufficient to solve the problems at hand. The difficulty is that most of the course materials developed over the past many decades have been tuned for students with a traditional engineering background in mathematics and physics. As we explore new applications areas such as biology and software systems, the backgrounds of the students are likely to be very different.

CDS 101 is an experimental course being developed in the Control and Dynamical Systems (CDS) Department at Caltech to begin to address this problem. This section provides an overview of the course, including the pedagogical approach we have taken. More details on the course, including PowerPoint lectures, are available on the course homepage:

`http://www.cds.caltech.edu/~murray/cds101`

The work in this section is joint work with Hideo Mabuchi at Caltech and Karl Åström at Lund.

4.1 Course Structure

CDS 101 is intended for advanced students in science and engineering who are interested in the basic concepts and applications of feedback control, but not the analytical techniques for design and synthesis of control systems. Special attention is payed to insuring that the course is accessible to students from biological, physical, and information sciences. These students have varying levels of mathematical sophistication, especially with regards to continuous mathematics.

The goal of the course is to enable students to use the principles and tools of feedback and control in their research activities. In particular, after taking this course, students should be able to build control-oriented models of physical, biological, or information systems and simulate those models in the time-domain; analyze stability, performance and robustness of the models; and design rudimentary feedback control systems in the time and frequency domain. Special emphasis is given to state space methods for analysis and synthesis since these techniques are needed for systems that are nonlinear and asynchronous.

CDS 101 is taught as a one hour/week main lecture, a weekly homework set, and a tutorial/application lecture. The course is taught in 10 weeks, with 8 homework sets, a midterm, and a final.

The main lecture (Mondays, 2–3 pm) is a prepared presentation that covers the main topics for the week. Hardware demonstrations are used to convey concepts when possible, along with simulations and interactive MATLAB sessions. Each lecture introduces the main MATLAB commands needed to implement the concepts. A printed copy of the lecture presentation is handed out at the beginning of each lecture so that students can take notes. For many lectures, a printout of the MATLAB code used for the examples used in the lecture is also included. A sample slide, showing the overview of Lecture 3, is shown in Figure 6.

The homework set for CDS 101 consists of 2 problems per week, with at least one of these being a computer exercise. The computer exercises use MATLAB and SIMULINK, and consider examples that are moderate complexity, to allow the power of the tools to be demonstrated. The homeworks are designed to require approximately 2–3 hours to solve. Students were asked to report the number of hours spent on the homework on the first page of the homework, so that the amount of time used could be tracked and the difficulty of the homeworks could be adjusted appropriately.

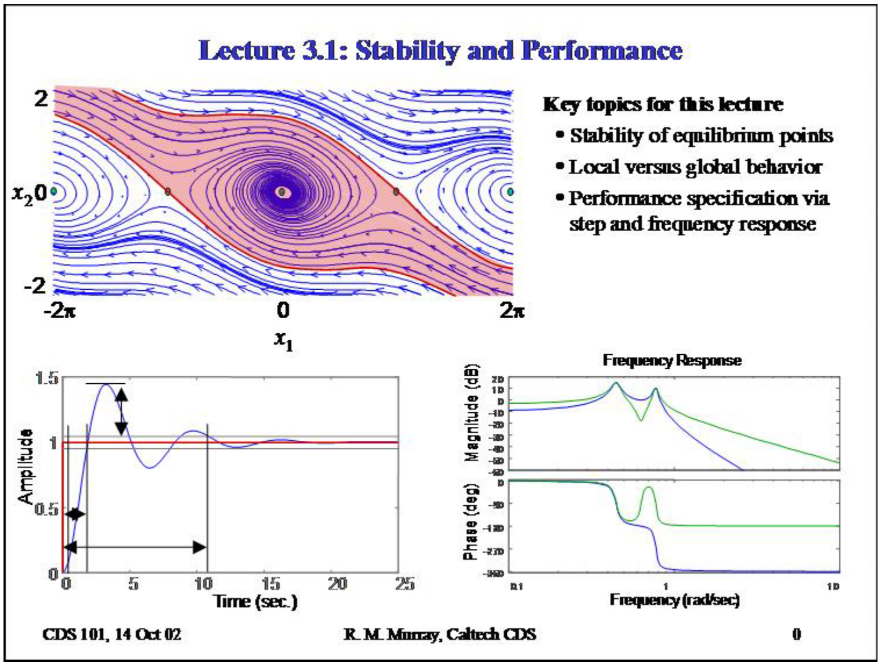

Lecture 3.1: Stability and Performance

Key topics for this lecture

- Stability of equilibrium points
- Local versus global behavior
- Performance specification via step and frequency response

CDS 101, 14 Oct 02

R. M. Murray, Caltech CDS

0

Figure 6: Summary slide for CDS 101, Lecture 3.1.

Research-oriented lectures are given each week (Friday, 2–3 pm) to provide application examples and review selected material in advance of the midterm and final. The Friday lectures are given by Caltech faculty or postdocs, and are intended to provide a bridge between the topics given in classes and the research activities at Caltech. For 2002, the lectures included:

- Insect flight modeling and control, Michael Dickinson (BE)
- Congestion control of the Internet, Eric Klavins (CS)
- Quantum feedback control, Hideo Mabuchi (Ph)
- Control of the Keck and CELT telescopes, Doug MacMartin (CDS)

These lectures are advertised to the campus so that anyone can attend.

In addition, some Friday lecture slots are used for reviews by the head TA. The first Friday lecture of the term is a tutorial on the use of MATLAB. The lectures on the fifth and tenth weeks of the term are reserved for a review of the material that will be covered on the midterm and final.

CDS 101 is co-taught with the first term of a more traditional controls course, CDS 110ab. CDS 110a shares the Monday lecture with CDS 101, but has an additional two hours of lectures on Wednesday (1–3 pm) that goes into more detail on the selected topics. The lectures are designed such that very little information is repeated in the Wednesday lectures. This structure for the course requires that the Monday lectures be self-contained, yet serve as an effective introduction to the more detailed concepts provided in the Wednesday lectures.

4.2 Teaching Pedagogy

In addition to covering a somewhat non-traditional set of topics, CDS 101 makes use of a number teaching methods to increase the effectiveness of the course.

Pre-course Since CDS 101 is intended to be taken by students with a diverse mathematics background, a special *pre-course* is offered in the week before the course begins. This pre-course gives a concise introduction to three main topics: linear algebra, ordinary differential equations, and modeling of physical systems. The pre-course spans two days, with two one-hour lectures given on each topic, along with homework sets and computer exercises (intended to be done that same day).

In 2002, the pre-course had 25 students attend the lectures. The pre-course was advertised through flyers, e-mail to faculty, and e-mail to mailing lists for various groups, centers, and departments. The most effective mechanism for announcing the course was via established mailing lists that went to students.

Mud cards "Mud cards" are a simple tool—originally brought to the authors attention through the CDS program at MIT [5]—for allowing students to get additional information on topics that they did not understand in the lecture. 3×5 cards are handed out at the beginning of each lecture and the students are instructed to write on the cards questions about the "muddiest" part of the lecture. These cards are collected at the end of the class and the teaching assistants sort them into categories and write up answers to the questions that arose. The answers are posted on the web the evening of the lecture, rapidly providing additional information for students who had questions.

The database used for storing mud card responses allows the questions and responses to be posted on the web page for each lecture, so that students could see what other questions had been asked for each lecture. In addition, the database allows responses to frequently asked questions on the homeworks, which could be posted on the web for students who were working on the homework sets.

Mud cards have turned out to be very heavily used by the students. In 2002, there were an average of 13 mud card responses for each Monday lecture, with a peak of 22 (these numbers are for both CDS 101 and 110 students). The questions ranged from clarification of details in the lecture to conceptual questions regarding the material to administrative questions/comments about the course.

Course web page and videos The course web page contains links to all of the lectures (in PowerPoint format), the reading assignments, streaming video for the lectures (posted w/in one week after the lecture), copies of the homeworks and solutions. The head TA is responsible for maintaining the web page.

All lectures (including the optional lectures) are videotaped and made available to students in the course. The videos are available as standalone tapes that can be checked out, and as streaming video from the course web page.

Surveys The course uses a variety of surveys to judge the effectiveness of the teaching methods. Results of the surveys are available on the course homepage (see the presentations in the "About CDS 101" topic on the course homepage).

A background survey is distributed on the first day of class to determine the level and preparation of the students in the class. The survey asks for the year and major of the student, as well as the courses that the student has previously taken (chosen from a list). In addition, a list of topics are given and the students are asked about their familiarity with each. The topics range from those that are prerequisites for the course, to the topics that will be covered in the course, to advanced topics in control.

Surveys are also distributed as part of the midterm and final exams. These surveys ask about the utility of the various elements of the course and ask for specific suggestions on how to improve the course. On the midterm survey, the number of hours spent per week is also requested, along with a list of the aspects of the course that take the most time. In both surveys, the same list of topics given in the background survey are also included, so that the (self-assessed) progress of the students can be tracked.

4.3 Course Syllabus

The approach taken in CDS 101 is to focus on two main principles of feedback and control:

1. Robustness to uncertainty through feedback
2. Design of dynamics through feedback

These particular principles are emphasized because of their broad relevance to many different domains, including biology, computer science, and engineering.

The course makes use of a number of tools for modeling and analysis. MATLAB is used as the primary software for the course, due to its broad use in science and engineering. SIMULINK models are given to the students starting in the first weeks of the course, to allow them to explore the concepts being described in lectures using representative examples.

Some of the topics in the course are a bit non-traditional for a first course in control and are listed below (based on the 2002 version of the course). The first half of the course focuses on issues related to dynamics, modeling, stability and performance, with only passing reference to feedback systems. Although control is not formally introduced until the second half of the course, feedback is used in the examples and homeworks starting in the first lecture. This allows the concepts of control to be studied throughout the course, building intuition for the eventual formal analysis of feedback systems.

Lecture 1.1: Introduction to Feedback and Control

- Define what a control system is and learn how to recognize its main features
- Describe what control systems do and the primary principles of control
- Give an overview of CDS 101/110; describe course structure and administration

Lecture 2.1: System Modeling

- Describe what a model is and what types of questions it can be used to answer
- Introduce the concepts of state, dynamics, inputs and outputs
- Provide examples of common modeling techniques: finite state automata, difference equations, differential equations, Markov chains
- Describe some common modeling tradeoffs

Lecture 3.1: Stability and Performance

- Describe different types of local stability of an equilibrium point
- Explain the difference between local stability, global stability, and related concepts
- Describe performance measures for (controlled) systems, including transients and steady state response

Lecture 4.1: Linear Input/Output Systems

- Describe linear system models: properties, examples, and tools
- Characterize the stability and performance of a linear system in terms of eigenvalues
- Compute linearization of a nonlinear systems around an equilibrium point

Lecture 5.1: Controllability and State Feedback

- Define controllability of a control system
- Give tests for controllability of linear systems and apply to examples
- Describe the design of state feedback controllers for linear systems

Lecture 6.1: Transfer Functions

- Define the input/output transfer function of a linear system
- Derive the transfer function corresponding to a system in state space form
- Build transfer functions for interconnected systems, via block diagram algebra

Lecture 7.1: Loop Analysis

- Show how to compute closed loop stability from open loop properties
- Describe the Nyquist stability criterion for stability of feedback systems
- Define gain and phase margin and determine it from Nyquist and Bode plots

Lecture 8.1: Loop Shaping

- Describe the use of frequency domain performance specification
- Show how to use loop shaping to achieve a performance specification
- Work through a detailed example of a control design problem

Lecture 9.1: PID and Root Locus

- Define PID controllers and describe how to use them
- Introduce the root locus technique and describe how to use it to choose loop gain
- Show some of the limitations of feedback due to RHP poles and zeros

Lecture 10.1: Uncertainty Management

- Describe how feedback and control are used as tools for uncertainty management
- Summarize the main principles and tools for the course

One of the features of the course is that it is taught entirely without the use of Laplace transforms (they are introduced in the CDS 110a variant, but not in CDS 101). This is in keeping with the desired audience for the course, who often will have little or no background in complex variables or transform techniques. The material is presented such that it is entirely consistent with the use of Laplace transforms, but only the term “transfer function” is used.

In future years, we expect to modify these topics somewhat. We are likely to remove root locus techniques completely and insert more material on observability (to balance what is already included on controllability). In general more effort is also needed to focus on state space representations for concepts, since these are often the more natural form for problems in many of the sciences. We are currently in the process of writing a textbook to go along with the course, in collaboration with Karl Åström.

5 CDS 270: Frontiers in Control and Dynamical Systems

In addition to CDS 101, which is aimed at making control more accessible, we have also developed a course at Caltech to explore the application of control and dynamical systems (CDS) tools to new domains. This course, CDS 270, was developed jointly with Hideo Mabuchi starting in 2001. This section gives some information about how the course is run and our experiences with teaching. More details are available on the course homepage at

`http://www.cds.caltech.edu/~murray/courses/cds270`

5.1 Course organization

The course is organized around small teams consisting of CDS and non-CDS students who work on projects of mutual interest in some faculty member's research area. The main goals are for the participating CDS and science/engineering faculty to become more familiar with each other's work and expertise, and to get our graduate students from different groups interacting with each another. The initial output of the course is a paper that could be submitted to a conference (either in control or the application domain). In addition, we hope to explore new research directions that can lead to collaborations and projects between CDS faculty members and other groups around the campus.

The course is "taught" by two faculty who share the responsibilities for identifying topics and running the course. There are no formal lectures for the course, so this course is taught on top of the normal teaching load. Students receive six units for the course (of a 48 unit typical load; equivalent to 2 credits for universities with a 16 credit/term load).

Planning for the course begins in the winter term, when the instructors seek out Caltech faculty who are interested in "sponsoring" a project for the course. We typically identify faculty that we have interacted with before and that we feel would be open to increasing that interaction, but who do often not have a strong background in control or dynamical systems. We ask them to provide an idea for an area that is of interest to them and for which they CDS tools could apply, plus two students from their group to participate in the course. In addition, we ask the CDS faculty to encourage their students to participate, so that we have enough students to form 4-5 person teams in each topic area. We usually shoot for 6-8 topics each year.

The course is given in the spring term (April through mid-June), with the one 90 minute class meeting scheduled per week for 9 weeks. The first class is an organizational meeting where each topic is described briefly. Students are asked to submit a form listing their top three choices for a project. The course instructors use this information to form the teams. The domain experts are typically pre-assigned to the topic that their advisor is sponsoring while the CDS students are assigned based on their knowledge of the types of tools that are likely to be required. The second class of the term is used to announce the groups and to get each group together to start talking about their topic and to pick a time for their next meeting. The instructors go around the classroom to make sure that each group is making progress.

There are no formal classes for the rest of the term. At midterms (week 5), each group gives a short (10 minute) presentation of their progress to date. Sometimes this extends across two weeks. Similarly, at the end of the course, each group gives a 15 minute presentation of their results. The final paper from the group is due at the end of finals week.

5.2 Experience to Date and Lessons Learned

So far, the class has been very successful. Eight different Caltech faculty have participated in the first two years of the course and several joint projects have developed from these initial interactions (see next section). Several faculty have sponsored projects in multiple years, although we try to rotate around to increase the exposure to new research groups. The CDS students have reacted very positively to the course as a mechanism for increasing their breadth of knowledge about control applications.

A key factor in making the class work is identifying faculty who are willing to sponsor a topic. It is best to meet with the faculty member to talk about the course ahead of time and to work with them in formulating the outline of the topic (the students will fill in the details). Also, the sponsoring faculty member should be willing to spend time during the course to meet with the students and provide some guidance. This keeps the students from heading off in a direction which is too ambitious for a 10 week, 3 unit course.

Another important aspect of the course is getting students with a background in control and dynamical systems to participate. At Caltech, the faculty in CDS strongly encourage all of their students to participate in the class. For PhD students, this means that they will take the class multiple times during their studies. They can either choose related topics, to build up some background in a specific area, or very different topics, to build some insight into multiple areas.

The format of the course has evolved to the current format to provide some driving events for the course. Early in the course, some groups did not meet regularly, so we set up the second class session to get the groups together and force them to choose a time for their next meeting. Similarly, the midterm presentation provides a clock for the groups so that they have some incentive to meet and come up with some ideas before presenting their ideas to the rest of the class. The final presentation gives a starting point for the report that is due at the end of the term.

5.3 Success Stories

To illustrate some of the success of the course, we present here three specific examples of outputs from CDS 270.

Software Control Theory Jason Hickey, a faculty member in CS, has sponsored projects for two years in robust software. The second year, the team working on the project explored the idea of using feedback as an integral element of software systems, motivated by an example of a sorting algorithm that used feedback to provide robustness to machine loading and software bugs. The group developed models for sorting using Markov chains that were predicted the performance of the algorithm and allowed analysis of its dynamics [12]. In late 2002, a proposal was submitted to NSF based on the work done as part of the class project.

Microbial Ecosystems Dianne Newman, a faculty member in Geobiology, sponsored a project on the dynamics of microbial ecosystems that are important in corrosion of metals. The students developed some preliminary models of a simplified system using MATLAB and explored the dynamics of the system. This project is expected to lead to the establishment of a reading group in microbial ecosystems that includes approximately 8 faculty from CDS, Environmental Science and Engineering, Biology, and Chemical Engineering.

CDS 101 One of the motivations for creating CDS 101 was the interest around the campus in applying CDS tools to new application domains. Since CDS 101 is taught in the fall term, CDS

270 acts as advertising for the course: non-CDS students who want to learn more about some of the tools that they are seeing in their projects can sign up for the course in the following academic year.

6 Summary

In this paper we have given a brief summary of the findings and recommendations of the Panel on Future Directions in Control, Dynamics, and Systems. This has been augmented with a set of grand challenge problems and a description of some ongoing educational initiatives at Caltech that are aligned with the recommendations of the Panel.

The author would like to thank the CDS Panel and particularly the writing committee for the numerous discussions and meetings that led to the Panel report and to the additional thoughts described here. The author also thanks Marc Jacobs and Belinda King at AFOSR for their support on this project, both intellectual and financial.

References

[1] 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`.

[2] T. Chung, L. Cremean, W. B. Dunbar, Z. Jin, E. Klavins, D. Moore, A. Tiwari, D. van Gogh, and S. Waydo. A platform for cooperative and coordinated control of multiple vehicles: The Caltech multi-vehicle wireless testbed. In *Workshoip on Cooperative Control and Optimization*, 2002.

[3] R. D'Andrea. Robot soccer: A platform for systems engineering. *Computers in Education Journal*, 10(1):57–61, 2000.

[4] C. Georgakis and W. Levine. Workshop on the research needs in modeling, dynamics, modeling and control of complex engineering systems. `http://www.isr.umd.edu/ISR/Complex/`, May 2002.

[5] S. R. Hall, I. Waitz, D. R. Brodeur, D. H. Soderholm, and R. Nasr. Adoption of active learning in a lecture-based engineering class. In *ASEE/IEEE Frontiers in Education Conference*, 2002. Available at `http://www.cdio.org/papers/papers.html`.

[6] S. Mitter. Real-time information theory. In *IEEE Information Theory Workshop on Detection, Estimation, Classification and Imaging*, 1999.

[7] R. M. Murray. Top ten (challenge) problems in control. `http://www.cds.caltech.edu/~murray/topten`.

[8] R. M. Murray, editor. *Control in an Information Rich World: Report of the Panel on Future Direcitons in Control, Dynamics and Systems.* SIAM, 2003. To Appear. Available at `http://www.cds.caltech.edu/~murray/cdspanel`.

[9] R. M. Murray, K. J. Åström, S. P. Boyd, R. W. Brockett, and G. Stein. Future directions in control in an information-rich world. *Control Systems Magazine*, April 2003.

[10] National Research Council. *Uninhabited Air Vehicles: Enabling Science for Military Systems.* National Acadamy Press, 2000.

[11] H. S. Tsien. *Engineering Cybernetics.* McGraw-Hill, 1954.

[12] S. Waydo, W. B. Dunbar, and E. Klavins. Feedback stabilization of faulty software systems: A case study. In *Proc. IEEE Control and Decision Conference*, 2003. Submitted.

[13] N. Weiner. *Cybernetics: Or Control and Communication in the Animal and the Machine.* John Wiley, 1948.