

COVER SHEET FOR PROPOSAL TO THE NATIONAL SCIENCE FOUNDATION

PROGRAM ANNOUNCEMENT/SOLICITATION NO./CLOSING DATE <i>if not in response to a program announcement/solicitation enter NSF 03-2</i>					FOR NSF USE ONLY	
NSF 02-168			02/12/03		NSF PROPOSAL NUMBER	
FOR CONSIDERATION BY NSF ORGANIZATION UNIT(S) <small>(Indicate the most specific unit known, i.e. program, division, etc.)</small>					0326554	
CCR - ITR MEDIUM (GROUP) GRANTS						
DATE RECEIVED	NUMBER OF COPIES	DIVISION ASSIGNED	FUND CODE	DUNS# <small>(Data Universal Numbering System)</small>	FILE LOCATION	
				009584210		
EMPLOYER IDENTIFICATION NUMBER (EIN) OR TAXPAYER IDENTIFICATION NUMBER (TIN)		SHOW PREVIOUS AWARD NO. IF THIS IS <input type="checkbox"/> A RENEWAL <input type="checkbox"/> AN ACCOMPLISHMENT-BASED RENEWAL		IS THIS PROPOSAL BEING SUBMITTED TO ANOTHER FEDERAL AGENCY? YES <input type="checkbox"/> NO <input checked="" type="checkbox"/> IF YES, LIST ACRONYM(S)		
951643307						
NAME OF ORGANIZATION TO WHICH AWARD SHOULD BE MADE			ADDRESS OF AWARDEE ORGANIZATION, INCLUDING 9 DIGIT ZIP CODE			
California Institute of Technology			California Institute of Technology			
AWARDEE ORGANIZATION CODE (IF KNOWN)			1201 E California Blvd			
0011312000			Pasadena, CA. 911250001			
NAME OF PERFORMING ORGANIZATION, IF DIFFERENT FROM ABOVE			ADDRESS OF PERFORMING ORGANIZATION, IF DIFFERENT, INCLUDING 9 DIGIT ZIP CODE			
PERFORMING ORGANIZATION CODE (IF KNOWN)						
IS AWARDEE ORGANIZATION (Check All That Apply) <small>(See GPG II.C For Definitions)</small>			<input type="checkbox"/> SMALL BUSINESS <input type="checkbox"/> FOR-PROFIT ORGANIZATION		<input type="checkbox"/> MINORITY BUSINESS <input type="checkbox"/> WOMAN-OWNED BUSINESS	
					<input type="checkbox"/> IF THIS IS A PRELIMINARY PROPOSAL THEN CHECK HERE	
TITLE OF PROPOSED PROJECT ITR: Information Dynamics for Networked Feedback Systems						
REQUESTED AMOUNT \$ 3,999,943		PROPOSED DURATION (1-60 MONTHS) 48 months		REQUESTED STARTING DATE 10/01/03		SHOW RELATED PRELIMINARY PROPOSAL NO. IF APPLICABLE
CHECK APPROPRIATE BOX(ES) IF THIS PROPOSAL INCLUDES ANY OF THE ITEMS LISTED BELOW						
<input type="checkbox"/> BEGINNING INVESTIGATOR (GPG I.A)			<input type="checkbox"/> HUMAN SUBJECTS (GPG II.C.11)			
<input type="checkbox"/> DISCLOSURE OF LOBBYING ACTIVITIES (GPG II.C)			Exemption Subsection _____ or IRB App. Date _____			
<input type="checkbox"/> PROPRIETARY & PRIVILEGED INFORMATION (GPG I.B, II.C.6)			<input type="checkbox"/> INTERNATIONAL COOPERATIVE ACTIVITIES: COUNTRY/COUNTRIES INVOLVED (GPG II.C.9)			
<input type="checkbox"/> HISTORIC PLACES (GPG II.C.9)						
<input type="checkbox"/> SMALL GRANT FOR EXPLOR. RESEARCH (SGER) (GPG II.C.11)						
<input type="checkbox"/> VERTEBRATE ANIMALS (GPG II.C.11) IACUC App. Date _____			<input type="checkbox"/> HIGH RESOLUTION GRAPHICS/OTHER GRAPHICS WHERE EXACT COLOR REPRESENTATION IS REQUIRED FOR PROPER INTERPRETATION (GPG I.E.1)			
PI/PD DEPARTMENT Control and Dynamical Systems			PI/PD POSTAL ADDRESS 1200 E. California Blvd			
PI/PD FAX NUMBER 626-796-8914			Mail Code 107-81			
			Pasadena, CA 91125			
			United States			
NAMES (TYPED)	High Degree	Yr of Degree	Telephone Number	Electronic Mail Address		
PI/PD NAME Richard M Murray	Ph.D.	1991	626-395-6460	murray@caltech.edu		
CO-PI/PD Michelle Effros	PhD	1994	626-395-3729	effros@caltech.edu		
CO-PI/PD Babak Hassibi	PhD	1996	626-395-4810	hassibi@caltech.edu		
CO-PI/PD Steven Low	PhD	1992	626-395-6767	slow@caltech.edu		
CO-PI/PD Leonard J Schulman	PhD	1992	626-395-6839	schulman@caltech.edu		

Project Summary

We propose to develop theory, algorithms and experimental demonstrations for investigating the *dynamics* of information in complex, interconnected systems. The key technical thrusts of our proposed activities are:

1. *Real-time information theory* for understanding fundamental limits of information flow in the presence of timing constraints and building a framework for the study of information dynamics.
2. Theory for *robust control of networks* to provide stable, high throughput data flow using source coding and feedback within and across protocol levels.
3. *Packet-based control theory* to allow feedback control of physical and information systems across networks, including issues of latency, multi-description coding, varying channel capacity, and feedback instabilities.
4. A framework for *computational complexity of networked systems* to understand the tradeoffs between computation, communication, and uncertainty in networked information processing systems.

These thrusts connect together existing strengths of the investigators but also require new techniques that lie at the intersection of traditional disciplines. While diverse in application, they represent a common core of intellectual thrusts that integrate computer science, control, and communications.

The results of the proposed research will be evaluated on two testbeds already under development at Caltech. The first is the Multi-Vehicle Wireless Testbed (MVWT), which provides a distributed environment for control of 8-10 vehicles performing cooperative tasks in a real-time environment. The second is the WAN in Lab, a wide area network consisting of high speed servers, programmable routers, electronic crossconnects, and long haul fibers with associated optical amplifiers, dispersion compensation modules and optical multiplexers.

Broader Impact In addition to pursuing fundamental research issues to address the above problems, we will develop elements of a curriculum that will provide training to students in information systems that blends communications, computation, and control. We will integrate our research framework into a recently created course being developed at Caltech on the principles of feedback and control, CDS 101, as well as develop a second course, IST 201, aimed at bringing together faculty and students interested in working on problems at the boundaries of these traditional disciplines. These courses will be offered through the Departments of Electrical Engineering, Computer Science, and Control and Dynamical Systems at Caltech and will provide the necessarily educational training needed for this new research area.

Through the newly established Information Science and Technology Institute (ISTI) at Caltech and its Center for the Mathematics of Information (CMI), we will pursue outreach activities that include providing winter workshops for key technology leaders in industry, summer workshops to bring together students and researchers for focused study in selected topics, and seminar and visitor programs.

TABLE OF CONTENTS

For font size and page formatting specifications, see GPG section II.C.

Section	Total No. of Pages in Section	Page No.* (Optional)*
Cover Sheet for Proposal to the National Science Foundation		
A Project Summary (not to exceed 1 page)	1	_____
B Table of Contents	1	_____
C Project Description (Including Results from Prior NSF Support) (not to exceed 15 pages) (Exceed only if allowed by a specific program announcement/solicitation or if approved in advance by the appropriate NSF Assistant Director or designee)	15	_____
D References Cited	2	_____
E Biographical Sketches (Not to exceed 2 pages each)	10	_____
F Budget (Plus up to 3 pages of budget justification)	10	_____
G Current and Pending Support	9	_____
H Facilities, Equipment and Other Resources	2	_____
I Special Information/Supplementary Documentation	2	_____
J Appendix (List below.) (Include only if allowed by a specific program announcement/ solicitation or if approved in advance by the appropriate NSF Assistant Director or designee)	_____	_____
Appendix Items:		

*Proposers may select any numbering mechanism for the proposal. The entire proposal however, must be paginated. Complete both columns only if the proposal is numbered consecutively.

1 Introduction

Thanks to the revolution in information technology, we face an untamed abundance of riches: the coupling of sensors and computers has yielded reams of experimental data about our universe; the PC revolution has created staggering amounts of processing power; satellites and fiber optics have given us stunning worldwide connectivity; and inexpensive microprocessors have allowed us to embed actuators in the most commonplace of objects. What we lack is the ability to combine these resources effectively: sensors and actuators still work primarily in manufactured pairs despite ubiquitous wired and wireless communication channels; patterns in data go undetected while millions of processors sit idle. Existing disciplines, which each focus on a single aspect of the problem, have not yet provided the mathematical foundation for understanding integrated information systems.

We propose to conduct research aimed at developing the mathematical foundations for a comprehensive study of information. We focus here on the role of *information dynamics*, which we define as the study of the time dependent flow of information (through networks, computational nodes, and algorithms). We believe that a theory for information dynamics is a critical element in the broader agenda of developing the fundamental mathematics required to conceive, design, implement and operate large scale information systems.

Trends and Opportunities

Technological progress changes the nature of the world. These changes must be understood for us to benefit from past advances and achieve further progress. At least three trends are apparent from our vantage point.

First, the systems and organizations on which our lives depend are becoming increasingly large, complex, and interdependent, and with these changes comes greater vulnerability to widespread impact from local failure. We wish to understand ever more connected, data-rich, and complex systems: how neurons implement the mind, how small changes at individual computers will affect the global behavior of the Internet, how our activities today will change the climate 100 years from now, how Moore's law affects the economy, how millions of heterogeneous and loosely connected components can be combined to build reliable systems, how interaction between those components will affect their performance capabilities.

Second, our world is increasingly awash with data. More and faster CPUs, more bandwidth, sensors with higher resolution, and larger storage devices increase our data collection capabilities but do not enlarge the limits on the human ability to absorb information. Existing information extraction techniques require levels of human interaction (by scientists, journalists, politicians, meteorologists, and consultants) that are unsustainable as data volumes grow. For example, the Centers for Disease Control and Prevention computes statistics on data it receives from hospitals across the nation, but interactively requesting data to substantiate suspicions of disease outbreak requires human intervention. Since the volume of the data precludes close human oversight, there is little doubt that many outbreaks go undetected and that others are discovered much later than they could be.

Third, we are surrounded by inexpensive computational power that not only fails to solve all of our existing problems but also seems to create new technological hurdles. Because the amount of computation required to solve "hard" problems grows so quickly as the volume of arriving data increases, no conceivable improvement in computer speed will allow current methods to keep up with the increasing tides of data. As the number of computers we use grows, their connectivity becomes more important, yet networking the latest high-tech computer devices often involves crawling around the floor. As we push the limits of complex system design (e.g., high performance aircraft, modern

cars, “just in time” manufacturing), we require more computer involvement in running systems once run by humans but simultaneously increase our vulnerability to unforeseen disruptions.

The trends described above affect our daily lives, and further progress depends on our ability to master the unprecedented challenges to which they give rise. We need to learn how to design and control systems whose size and complexity boggles the mind and whose robustness is crucial to our comfort, productivity, and safety. We need to build tools that automate the process of recognizing and distilling the trickle of useful information from an ever-increasing tide of data. We need to develop techniques to make existing methods scale to future data set sizes and build seamless connections between humans, machines, and networks.

We expect that progress in information science and technology will be driven by key developments in the mathematical foundations of information. In this proposal we focus on those questions related to understanding the dynamics of information as it applies to large distributed systems. In particular, we seek to better understand how we trade performance, communication and computation in a principled way, and what are the appropriate abstractions for designing, analyzing, and controlling large scale, decentralized systems.

Program Goals

The research goals of this program are to

- R1 Develop a theory of information that accounts for real-time issues such as interaction with the physical environment and computational constraints. This theory will allow a better understanding of the underlying tradeoffs between performance/accuracy, bandwidth, and computational resources.
- R2 Develop tools for robust control of networks. We will build on recent results in both protocol design and network coding to provide provably correct protocols for stable, high performance operations of real-time networks. The resulting protocols will be tested in a realistic operating environment to demonstrate their effectiveness.
- R3 Develop a theory of feedback systems that uses packets as the fundamental mechanism of information transfer. New techniques that account for latency, dropped and redundant packets, and multi-description coding across networks will be developed, analyzed, and tested.
- R4 Develop a complexity theory for information systems that allows a more complete understanding of how much information is required to achieve a given task and the role of feedback in algorithms, protocols, and systems.

Each of these goals requires an integration of traditional disciplines and we plan to achieve them through joint research projects that expand on and integrate our current research activities. The unifying thread of all of these topics is the focus on the *dynamics* of information.

In addition to these research goals, we will pursue the following educational goals:

- E1 Development of course materials that present a unified framework for computation, communications, and control, enabling new research activity that better addresses dynamic issues in information systems. These materials will be integrated into existing courses at Caltech and disseminated through existing multi-university educational alliances and the web.
- E2 Development of a new course specifically aimed at bringing together faculty and students from the departments of Electrical Engineering, Control and Dynamical Systems, and Computer Science to spend a single term studying a problem at the boundary of these traditional disciplines.

- E3 Creation of workshops for industry and academia to disseminate results, build new research communities, and increase the level of understanding of the principles of information science.

These goals will be accomplished in the context of a new initiative at Caltech in Information Science and Technology, which includes a strong educational component.

Finally, we will test our results on two novel demonstration projects:

- D1 The Caltech Multi-Vehicle Wireless Testbed (MVWT) will be used to test real-time communication protocols, packet-based control algorithms, and robust protocols for ad-hoc wireless networks.
- D2 The Caltech WAN in Lab will be used to test new network protocols, new data compression protocols, and new networked control algorithms. The bulk of the WAN in Lab will be built using funding to be obtained from other sources, including substantial matching from Caltech.

2 Technical Approach

In this section we describe the technical approach that we plan to take to meet the program goals. We break our proposed activities into four primary areas, each of which overlaps the interests of all of the investigators. We also describe our education and outreach plans, which are closely linked to the proposed research.

2.1 Real-time information theory

The key results of information theory—such as the separation principle for source and channel coding, and rate distortion theory—are asymptotic results that require an infinite horizon of observation of the random processes involved as well as infinite delays in encoding and decoding. As such, they are idealizations and give fundamental limits. In many applications, such as signal processing and control, timing considerations are crucial. Control theory, in particular, deals with real-time constraints upfront, since controllers must take action and produce control signals using observations only currently available. However, control theory has by and large ignored information-theoretic considerations, ostensibly on the grounds that the controller could be assumed to have perfect access to the (possibly noisy) output of the plant and imperfections such as quantization noise could either be ignored or modeled statistically.

Today with the advent of more and more complex systems, one is confronted with decentralized control problems or problems of control over networks, where the measurement and control signals are transmitted over communication channels. As such these signals are quantized into bits, encoded and decoded, and transmitted at finite rates, as a result of which decoding errors will occur. Natural questions to ask therefore are what properties should the communication channels have and how much information on the measurement and control signals should be transmitted to allow for the system to be controlled, stabilized, etc? Unfortunately, classical notions such as channel capacity and entropy cannot do the job since they assume infinite delays, whereas in control we require real-time solutions.

How to bring together information theory and real-time constraints is a central problem of this proposal. While rate of convergence results for source coding algorithms and error exponents for channel coding arguments attempt to add some notion of time to information theory, they fall far short of a complete theory of real-time information theory. Although the goal of developing such a complete theory is overly ambitious and notoriously difficult, there has been recent progress in this

area, especially along the lines described in the preceding paragraphs, in the work of Schulman [25] and the Ph.D. theses of Tatikonda and Sahai [24, 26].

Schulman's work deals with the problem of how to interactively compute using distributed processors if they are connected by noisy communication links. Real-time constraints are central here since the messages to be sent from one processor to the other depend on the messages already received from the other processor(s). Therefore grouping the messages into blocks for conventional block-coding is out of the question. Nonetheless, it is shown that there exist good coding strategies that allow the computation to proceed with arbitrarily low probability of error. Tatikonda has studied some of the information-theoretic limitations of control systems, in particular the capacity of systems with feedback. Following the spirit of Schulman, Sahai studies the issue of tracking an unstable scalar Markov process over a noisy channel. This is important since stabilizing an unstable control system requires one to be able to do this tracking. Due to the real-time constraints, one needs a notion stronger than channel capacity (coined "anytime" capacity and closely related to the theory of error-exponents) and requires either perfect feedback from the decoder to the encoder (so that the encoder knows how well the decoder is tracking) or that the encoder and decoder have access to common randomness.

There are clearly many directions in which the above works can be extended and further pursued. One important way in which we would like to extend Schulman's work and apply it to control is by establishing a rate-distortion version of that work (which currently applies only to lossless communication complexity). We are also interested in extensions to settings where we have a network of computational and physical elements connected by communication links. Again, of interest is the interplay between communication, computation and control, and in coming up with the right information-theoretic and control-theoretic abstractions that will allow us to analyze these systems and determine their stability and performance.

One problem at the boundary of communication, computation and control is decentralized (or distributed) control, where the control functions are distributed across interconnected computational elements. Clearly, if there exists perfect communication between the different computing elements, then the full state knowledge of each subsystem could be communicated in real-time and each controller could implement the centralized control law. However, when communication is not perfect the question is how close can the performance be to the centralized one. Hassibi has recently studied problems of this form, in the context of determining the capacity and power efficiency of wireless communication networks [9, 11]. In particular, for a broadcast channel with an M -antenna transmitter and n users (receivers), it is shown that if the transmitter has full knowledge of all the fading coefficients of the different links in the network then the capacity scales as $O(M \log n)$. However, if the transmitter has no knowledge about the fading coefficients the capacity is much less and scales only as $O(\log M)$, independent of n . Hassibi has been able to determine the minimum amount link information that must be provided to the transmitter to obtain a much higher $O(M \log \log n)$ capacity. For ad-hoc and sensory wireless channels with n nodes (users) it is shown that the power efficiency of the network (the number of bits that can be reliably delivered from the transmitter to the receiver for a unit of energy consumed in the network) scales as $O(\sqrt{n})$, irrespective of whether the nodes have global or only local knowledge of the network connections. These anecdotal examples show that it is of interest to determine the minimum amount of communication required to have the performance of a decentralized controller approach that of a centralized one.

We also plan to consider problems related to real-time data compression, a technology that is central to understanding the dynamics of information. While most of the lossless data compression algorithms used in practice bear strong resemblance to algorithms analyzed in the information theory literature, lossy compression algorithms used in theory bear little resemblance to the codes that we can analyze. Conversely, most of the codes that we can analyze are computationally

intractable at the limits of large coding dimension where their performance is well understood. The entropy-constrained dithered quantization (ECDQ) algorithm of Ziv [30] is one notable exception to the above trend. While ECDQ is not commonly used in practice, this seems to be mostly a matter of historical oversight. ECDQ bucks the above trend by being both extremely practical and provably good. An ECDQ requires only uniform scalar quantization followed by entropy coding, both of which can be achieved with very low complexity. Further, for very general sources ECDQ performance is guaranteed to be within an additive constant of the rate-distortion bound, which describes the optimal performance theoretically achievable by the best infinite-dimensional code. Recent work generalizes ECDQ from the traditional data compression model to a small number of network source coding scenarios [6, 7, 29]. In future work, we hope to develop corresponding ECDQ algorithms and performance bounds for more general network configurations.

Another interesting question arising from the investigation of ECDQ relates to the performance (measured in rate and distortion) versus complexity tradeoffs suggested by these codes and their bounds. Low complexity codes achieving the rate-distortion bound for arbitrary sources are not known and may not exist in general. Yet, the ECDQ algorithm demonstrates that guaranteeing performance within a small, constant distance of the rate-distortion bound is computationally feasible. Since information theory traditionally ignores computational constraints, any natural tradeoff between rate-distortion performance and computational complexity is currently unknown.

Together, these proposed activities represent a broad attack on developing the tools needed to understand the time dependent flow of information through networks and computational nodes. They will be integrated through their application to robust control of networks and real-time control (described in the next two sections) and through the integration of technologies on the MVWT and WAN in Lab testbeds.

2.2 Robust control of networks

Robust control of networks is a large area, spanning many topics. The basic problems in control of networks include controlling congestion across network links, routing the flow of packets through the network, caching and updating data at multiple locations, and managing power levels for wireless networks. The dominant feature in this class of problems 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.

We propose to focus our efforts on the dynamics and stability of the devices and protocols that control the flow of information across the network, tying it to real-time information theory and packet-based control theory.

We start with the problem of congestion control and routing. The flow of information between two nodes is controlled by the Transmission Control Protocol (TCP), which regulates the rate at which packets are sent from a server to a client. The path that packets take through the Internet is controlled by the Internet Protocol (IP). An important issue is to understand overall stability, robustness, and performance of the network when both TCP and IP are adapting to alleviate congestion.

Recent studies have shown that any TCP congestion control algorithm can be interpreted as carrying out a distributed primal-dual algorithm over the Internet to maximize aggregate utility, and a user's utility function is (often implicitly) defined by its TCP algorithm; see, e.g., [14, 15, 17, 18, 19, 20, 22] for unicast and [1, 13] for multicast. All of these papers assume that routing

is given and fixed at the time scale of interest, and TCP attempts to maximize aggregate utility over source rates. Consider the problem of maximizing aggregate utility over *both* routes and rates, subject to link capacity constraints. The striking feature of the associated dual problem is that the maximization over routes takes the form of shortest-path routing with prices as link costs. This raises the tantalizing possibility that TCP/IP might turn out to be a distributed primal-dual algorithm over the Internet to maximize utility, with proper choice of link costs. We show in [27], however, that the primal problem is NP-hard and hence in general cannot be solved by shortest-path routing. This suggests treating TCP/IP as a distributed approximation algorithm to maximize utility and evaluating its performance in terms of its competitive ratio. We will explore the connection between control and complexity in the specific context of routing and congestion control.

Another important issue is how TCP affects the routing stability of IP at slow timescales and how IP routing affects the rate stability of TCP at fast timescales. TCP affects routing stability by generating “prices” (congestion measures) at network links. Suppose the link costs used in shortest-path routing are weighted sums of static costs, which are traffic independent, and congestion prices generated by TCP, which are traffic dependent. For a special ring network, it is proved in [27] that there is an inevitable tradeoff between maximum achievable utility and routing stability. In this case, shortest-path routing based purely on congestion prices is unstable. Adding a sufficiently large static component to link cost stabilizes it, but the maximum utility achievable by shortest-path routing decreases with the weight on the static component. We conjecture that these conclusions hold in general networks as well and plan to explore them in the proposed work.

In addition, IP routing affects the rate stability of TCP by determining the interconnection among the local control laws that represent TCP algorithms at sources and active queue management (AQM) algorithms at links. We plan to investigate how the topology between a number of dynamic agents affects the overall stability of the system by building on recent results by Fax and Murray on stability of interconnected dynamical systems [3, 4, 5]. They consider a system consisting of N interconnected systems with identical dynamics that are trying to maintain relative stability to each other. They provide a set of necessary and sufficient stability conditions that allow the topology of the information flow (represented by the adjacency graph and its corresponding graph Laplacian matrix) to be separated from the dynamics of the agents (represented by a set of linear ordinary differential equations).

The conditions have a very natural interpretation in terms of the Nyquist plot of dynamical system. In the standard Nyquist criterion, one checks for stability of a feedback system by plotting the open loop frequency response of the system in the complex plane and checking for net encirclements of the -1 point. The conditions given by Fax and Murray correspond to replacing the -1 point with $-1/\lambda_i$ for each eigenvalue λ_i of the graph Laplacian L . This interpretation is illustrated in Figure 1. The results can easily be extended to consider weightings that are nonuniform and nonlinear extensions are being developed.

This condition illustrates how the dynamics of the system, as represented by a set of linear ordinary differential equations, interacts with the topology of the information flow in the system, as represented by the graph Laplacian. In particular, we see that it is the eigenvalues of the Laplacian that are critical for determining stability of the overall system. Additional results in this framework allow tuning of the information flow (considered as both sensed and communicated signals) to improve the transient response of the system [5]. Recent extensions in a stochastic setting [8] allow analysis of interconnected systems whose dynamics are not identical and where the graph topology changes over time.

These tools will be useful in understanding how IP routing affects the rate stability of TCP. The stability of the overall system depends on both the local algorithms and the graph structure of the interconnection among them. The graph is determined by the routing matrix, which also imposes a

computation must be taken into account, and algorithms must address the tradeoff between accuracy and computation time. Progress will require significantly more interaction between information theory, computer science and control than ever before.

As an example of the type of problem that we wish to treat, consider the problem of controlling a system where the computation, sensing, and actuation are all separated by a network. We assume that the network has multiple routes between the nodes, but that we must packetize our data and some packets may be lost or delayed. For simplicity, we will assume that we can model the network as a collection of N channels with some description of the channel characteristics that describes the bandwidth, latency, and probability that each channel will fail. Our challenge is to decide which information to send across which channel to maintain stability and achieve the optimal performance.

As a starting point, one could imagine breaking the data into frequency bands and prioritizing the data in each band. For example, we might suppose that information near the crossover frequency for the closed loop system might be more important (for stability) than very high or very low frequency data. Thus we might want to encode that data redundantly so that we are assured that it is received. Furthermore, we need to insure that the data is received with low latency, to minimize the phase delay near crossover. Understanding how to jointly optimize the compression and error control with the feedback architectures and algorithms in this context is an open problem, although some initial attempts have been made. In particular, there are initial efforts in understanding the effects of quantization [12, 28] and variable latency [16].

We plan to develop new approaches to studying control theory in the context of packetized data that flows across networks. Issues such as the congestion control, routing dynamics in unicast and multicast environments, and multi-description network coding will all be relevant and this activity is expected to require close coupling of concepts from control and communications. Conversely, we will also consider the impact of this sort of feedback algorithm on data compression, routing, congestion control, and computational complexity. The real-time nature of feedback systems requires that the dynamics of the information flow be taken into account. A concrete application of these results will be its use in control for the multi-vehicle wireless testbed, described in more detail below.

2.4 Computational complexity of feedback controlled systems

A major element in all of the problems described above is understanding computational complexity in a real-time setting. The theory of computational complexity has been constructed primarily around the following diagram, which we refer to as the *static* scenario:

$$input \longrightarrow computation \longrightarrow output$$


Frequently there is an implicit loop in which a human evaluates the results of the computation and provides revised inputs. However, since this part of the process is not automated, the theory has properly focused on studying the speed with which each “static” computation is performed.

This model has dictated the primary outlines of computational theory, since the most demanding applications of computers have been in intensive, “static” computations. In particular, the polynomial-time, log-space and parallel-computation hierarchies, and their associated algorithmic theories, have been extensively developed to analyze the difficulties of computational tasks in this “static” framework.

However, thanks to hardware advances, engineers can now put increasing autonomy in the control circuits of machines operating in real-world environments. The computer is now in a loop in which the input to each new computation is the result, in part, of previous actions taken by the computer itself. Auxiliary to the principal control task (stabilizing an aircraft, staying out of sight, etc.) are the secondary tasks of: (a) Choosing actions, consistent with the primary task, which

secondarily help keep the next round of computation manageable. (b) Orienting sensors so as to collect the most useful information for the next round of computation. (c) Choosing how long to spend on each computation: there is generally a tradeoff between, on the one hand, collecting a lot of data and computing a highly-optimized response, and on the other hand, taking early action.

We are driven therefore to develop a theory of complexity and algorithms revolving around the computational difficulty of *dynamic* tasks—those in which a single task is specified, in pursuit of which the computer undergoes many cycles in which it orients sensors, collects and processes data until it decides on an action, and repeat. We think of the “task” as being optimization of some parameter of the environment (the environment of the computer may include the craft housing it). Information dynamics is concerned with the computational complexity of the computations in the following diagram:



An especially important part of the proposed research concerns the computer’s choice of what data to gather. The squash coach’s weary admonition to “keep your eyes on the ball” is a reminder that even for a creature that is intelligent, and experienced at a particular task, this choice can be difficult in the face of the choices available—as well as critical to performance. It is therefore essential to model “attention” (or sensor orientation) as part of the automated control task.

A engineering example where this theory could play a role is in the area of optimization-based control theory [2]. Increasingly, control engineers are using online, real-time optimization as part of the feedback loop for stabilizing dynamical systems. These techniques are particularly useful when the dynamics are nonlinear and the system must operate in the presence of constraints on the inputs and states of the system. A common approach is to use so-called “receding horizon” techniques, where an optimization is performed over a finite time horizon to compute the input for the next sample period. The computation is then repeated immediately, updating the control input. A primary tradeoff is the length of the optimization horizon (which affects computation time) and the refresh rate for the controller (which depends on the computation time). Current experimental work at Caltech has explored these tradeoffs on flight control applications [21], but there is no theory to guide the design choices that must be made.

An even more challenging example is one in which the motion of a set of sensor platforms affects the information that is collected and the difficulty of the computation. This occurs, for example, in multi-vehicle control problems where one wishes to explore how to control the motion of the individual vehicles such that they maximize the information content available and minimize the amount of computation or communication required to reason about their environment. This is precisely the sort of problem that we expect to encounter in the scenarios we are exploring with the Caltech multi-vehicle wireless testbed, particularly related to cooperative control tasks.

2.5 Demonstrations

To demonstrate the theory and algorithms we will develop as part of this effort, we will make use of two hardware testbeds available at Caltech.

Caltech Multi-Vehicle Wireless Testbed (MVWT) Caltech has recently built a testbed consisting of 8 mobile vehicles (soon to be expanded to 18) with embedded computing and communications capability for use in testing new approaches for command and control across dynamic




Figure 2: The Caltech Multi-Vehicle Wireless Testbed. The leftmost figure shows the layout of the testbed area, including overhead cameras and fixed communication nodes (crosses and hexagons). The middle picture is the current laboratory, with two vehicles shown. The right picture is the current vehicle design being used.

networks. The system, shown in Figure 2, allows testing of a variety of communications-related technologies, including distributed command and control algorithms, dynamically reconfigurable network topologies, source coding for real-time transmission of data in lossy environments, and multi-network communications. A unique feature of the testbed is the use of vehicles that have second order dynamics, requiring real-time feedback algorithms to stabilize the system while performing cooperative tasks.

The testbed is part of the Caltech Vehicles Laboratory and consists of individual vehicles with PC-based computation and controls, and multiple communications devices (802.11 wireless Ethernet, Bluetooth, and infrared), and multiple sensor capability (vision, gyros, and sonar). We currently have two types of vehicles: wheeled mobile robots with a maximum speed of approximately 1 m/s and freely moving, castored platforms propelled by high performance ducted fans. The room contains access points for the 802.11 and Bluetooth networks, overhead visual sensing (to allow emulation of GPS signal processing), a centralized computer for emulating certain distributed computations, and network gateways to control and manipulate communications traffic.

The testbed will be used for implementing and testing ideas related to all aspects of this proposal. In particular, we plan to use the testbed to test new algorithms for multi-description source coding, routing and congestion control protocols, and distributed multi-vehicle control. More information is available on the web at <http://www.cds.caltech.edu/~murray/mvwt>.

Caltech WAN in Lab Caltech has recently proposed a new networking facility, WAN in Lab, for developing robust and stable ultrascale networking technologies. This facility is currently scheduled for a site review by the CISE Research Infrastructure (RI) program and we plan to seek additional funding as needed to construct this facility.

WAN in Lab will consist of four major building blocks: high speed servers, programmable routers, electronic crossconnects, and long haul fibers with associated optical amplifiers, dispersion compensation modules and optical multiplexers, as shown in Figure 3. It is literally a state-of-the-art WAN, in a laboratory. It will be integrated with, and complementary to, other high performance research and production networks that span the globe, such as Calren2, Abilene, the international high energy physics (HEP) network, and various Grid infrastructures. We believe that WAN in Lab will be the first such facility in any university in the world.

The WAN in lab has a number of key features that make it ideal for use in this activity. It is truly a wide area network (WAN), as opposed to a simulation or emulation, that provides both high speed and large delay. The design is highly reconfigurable and evolvable, providing flexibility in speed, delay, network topology and size. The facility is completely under our control and allows maximum




Figure 3: The Caltech WAN in Lab architecture.

freedom to conduct research that is impossible or too disruptive on research and production WANs shared by multiple groups. At the same time, WAN in Lab will be directly connected to the international high speed research and production networks, such as Calren2, Abilene, and HEP networks, expanding the scope and usefulness of the facility.

We plan to use the WAN in Lab facility for testing of network source coding algorithms, new protocols for congestion control, active queue and routing management, and crosslayer feedback.

2.6 Education and Outreach

In order to train the type of graduate students who will be able to analyze and design systems of this complexity and sophistication, a new set of courses will be required. These courses must cross traditional discipline boundaries and be accessible to students outside the field in which many of the tools were originally developed.

Caltech has a tradition of interdisciplinary education and is currently embarking on the development of several new courses that will support the educational goals of the proposal. We intend to use the funding provided by this proposal to develop course modules consisting of lecture materials, classroom examples, simulation models, and homework sets that use information systems as the driving application. The specific courses that we are developing are described in more detail below.

CDS 101. Principles of Feedback and Control. CDS 101 is an experimental course being developed in the Control and Dynamical Systems (CDS) Department at Caltech. The course is intended for advanced students in science and engineering who are interested in the principles and tools of feedback control, but not the analytical techniques for design and synthesis of control systems. Special attention is paid to insuring that the course is accessible to students from biological, social, 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 social, 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.

We anticipate that the research results carried out under this proposal will be integrated into CDS 101. In particular, CDS 101 includes a mandatory research lecture in which Caltech faculty summarize their research in a context that ties into the course objectives. Networked information systems provide an excellent opportunity to illustrate the key concepts in the course and to illustrate some of the future challenges in control.

IST 201. Communications, Computing and Control. We plan to develop a new graduate course to better educate students who will work on the types of topics described in this proposal. Initially, the course will begin as a reading group to explore the emerging literature in the (mainly pairwise) connections between computer science, communications and control. As the research in the program progresses, this will evolve into a regular course that will be part of the curriculum for graduate students in CDS, CS, and EE who are pursuing research in this area. We anticipate that this course will become a part of the core requirements in the Information Science and Technology (IST) program described in the next section.

Winter short course for industry As part of the Information Science and Technology Institute (ISTI), described in more detail below, we plan to develop a short course that will expose technology leaders from industry to the activities described in this proposal. The program will be part of a larger “winter school” on information science and technology that Caltech is developing through ISTI. The winter school is modeled on workshops and short courses offered through top business schools throughout the country, but with an exclusive focus on the science and technology of information.

Summer school for students Also as part of ISTI, we plan to develop a summer school in information dynamics, targeted at undergraduates and first year graduate students interested in complex, networked systems. This will be modeled after the Computing Beyond Silicon Summer School (CBSSS), offered in summer of 2003, in which we attracted leading research faculty and 32 highly motivated students from across the country to a four-week long program of lectures and laboratory exposure.

3 Project Team and Institutional Resources

Our team combines faculty with expertise in communications, networks, computer science, and control theory, along with a unique research and educational environment at Caltech.

3.1 Key Personnel

The following individuals will actively participate in the proposed research program.

Michelle Effros Michelle Effros is an expert in information theory and lossless and lossy compression algorithms. In recent years, her research has focused on theory and algorithms for compression of data traveling through networks.

Babak Hassibi Babak Hassibi’s expertise is in communication theory, information theory and control. He is currently actively involved in research in wireless communications systems and wireless networks. He has also worked in the areas of robust control and estimation and stochastic and adaptive signal processing.

Steven Low Steven Low is a networking expert and a leading researcher on congestion control of large scale networks. He will be responsible for the construction and operation of WAN in Lab and participate in research on the control of networks.

Richard Murray Richard Murray is an expert in nonlinear control theory and distributed control systems. He brings expertise in networked control systems, nonlinear dynamical systems, and engineering design. Murray will be responsible for the Caltech multi-vehicle wireless testbed and will participate primarily in activities related to packet-based control theory and control of networked systems.

Leonard Schulman Leonard Schulman studies the theory of computation. He brings expertise in algorithms, complexity theory, coding theory, communication protocols, combinatorics and probability. He is interested in the computational and communication complexities of control tasks.

3.2 Caltech Information Science and Technology (IST) Initiative

Caltech has launched a major initiative in Information Science and Technology (IST), with \$20M in seed funding anticipated from the Gordon and Betty Moore Foundation. The IST initiative consists of two phases, with the first phase focusing on the creation of the research component of IST. The second phase will focus on creating the academic and outreach programs. In order to launch the research thrust of IST, as well as to create an institute-wide community related to IST, Caltech is creating an Institute-wide IST Institute housing four new research centers that will provide faculty and students with new opportunities for research initiation and collaboration in IST-related areas.

A key element of the new initiative is the creation of a Center for the Mathematics of Information (CMI), which will consist of approximately 10-12 faculty (including all of the investigators on this proposal). The CMI is charged with developing a common language of information between researchers from different fields. Mathematics has provided the foundation for virtually every major new advance of the industrial revolution and beyond. It is fitting that a dedicated community of mathematicians, engineers and scientists combine to create a new way of thinking about information. Fundamental new ideas will emerge from this effort to influence all of the activities we are pursuing.

The CMI will provide leveraged support for the proposed activity, including funding for graduate students, seed funding for faculty, and infrastructure support for equipment and facilities. The CMI will also coordinate the outreach activities described in the previous section.

3.3 The CDS Alliance

Caltech has recently been awarded a grant by the Department of Education to establish the Control and Dynamical Systems (CDS) Alliance between the US and Brazil. The CDS Alliance is developing a shared curriculum at leading research and educational institutions in the US and Brazil in the areas of Control and Dynamical Systems. These areas are well positioned to play a pivotal role over the next several years in shaping a common scientific language across a variety of disciplines, from mechanical, aerospace, and environmental engineering to bioengineering, communications, and economics. This wide range of applications makes these disciplines ideally suited for cross-disciplinary fertilization and it fosters applications attractive to a wide scientific audience. A key focus of the activity is in developing new approaches to education and outreach for the dissemination of basic ideas to non-traditional audiences.

The current members of the CDS Alliance are Caltech, Princeton and UC Santa Barbara in the US, and the University of Campinas, the National Laboratory for Scientific Computation (LNCC) and the Institute for Pure and Applied Mathematics (IMPA) in Brazil. The Lund Institute of Technology has also proposed to join the alliance, using funding provided by Sweden.

The CDS Alliance will provide a mechanism to further develop course materials in this area and to disseminate those materials to a national and international collection of partner universities.

3.4 Prior NSF Support

Michelle Effros During the period of 1 October 1999–30 September 2001, Michelle Effros was supported by NSF SGER Grant CCR-9909026, “Exploring Source Codes for Network Technologies”. Supported work included theoretical bounds, code design, and universal coding for networks. Results include rate-distortion bounds for multi-resolution and multiple description source codes, source independent bounds on the maximal rate penalty (called the *rate loss*) for a variety of network compression algorithms, and design algorithms for lossless and lossy compression algorithms for sending information over networks.

Babak Hassibi Babak Hassibi was awarded NSF CAREER grant 0133818, “Multi-antenna communications: Information theory, codes and signal processing”, for the period 2002–2007. This award supports investigations related to the information-theoretic, coding-theoretic, and signal processing aspects of multi-antenna wireless systems, as well as the impact of integrating their solutions into a multi-user wireless network. Results include the invention of efficient space-time codes for coherent and non-coherent detection, the development of polynomial-time algorithms for maximum-likelihood decoding, and the study of the asymptotic capacity and power efficiency of large wireless networks.

Steven Low Steven Low was awarded NSF ITR ANI-0113425 (co-PI: John Doyle), 2001–04, for the project “Optimal and Robust TCP Congestion Control”. We have developed a mathematical model to understand the equilibrium and stability of large scale networks under end-to-end control, based on duality theory and distributed linear control. It has led to scalable TCP algorithms, the implementation of which has broken new grounds in a recent demonstration on a high speed intercontinental network.

Richard Murray During the period of 1995–2000, Murray was supported by NSF CMS-9502224 for a project on “Two Degree of Freedom Design for Robust Nonlinear Control of Mechanical Systems”. Motivated by applications in flight control and robotics, this project was focused on the use of two degree of freedom design techniques to generate nonlinear controllers for mechanical systems performing motion control tasks. New theoretical approaches were developed that exploited the special structure of these systems and provided insights into the control mechanisms for classes of mechanical systems.

Leonard Schulman Leonard Schulman was awarded NSF CAREER grant 9876172 (renumbered 0049092 upon change of affiliation), “Computation Methods,” for the period 1999–2003. This award supported investigations related to clustering algorithms, statistics and learning theory, combinatorics, discrete probability, coding theory, interaction in computation, and quantum computation.

4 Statement of Work

To achieve the research goals stated above, we will work as a team to develop new theory, algorithms, and experimental demonstrations in the area of information dynamics. While there are already many pairwise connections between the investigators (including 2 current joint students and several joint projects and proposals), we anticipate a much more cohesive research activity in which a larger group of faculty and students will participate. In addition to standard mechanisms such as seminars, biweekly group discussions, and joint advising, the two experimental testbeds will provide a means

to integrate research ideas in common settings and build on the results of the different activities in the program.

We anticipate that the following milestones will be achieved over the course of the program. While the details of these milestones will certainly change in the course of our research, they provide our current plan for how the research activities will proceed.

Year One

- Initial results in real-time information theory and computational complexity of feedback systems, demonstrating how to put time and computation constraints into rudimentary networked feedback systems.
- Implementation of multi-description codes for real-time, networked control systems, with implementation and testing on the Caltech multi-vehicle wireless testbed.
- Application of results in stability of interconnected dynamic systems to routing and congestion control problems.
- Establishment of the WAN in Lab testbed. (The equipment for this testbed will be supported by outside funds, but some integration of the equipment will be performed by students and postdocs supported by this program.)
- Development of a prototype course (reading group) in information dynamics, to be offered in Winter 2004.

Year Two

- Preliminary implementation of new protocols and networked coding/routing algorithms developed in year 1 on the WAN in lab testbed.
- Demonstration of distributed, real-time control of vehicle formations using optimization-based algorithms across wireless networks on the Caltech multi-vehicle testbed.
- Development and implementation of a teaching module in CDS 101 for Fall 2004 on control of information systems, focused on results relevant to this proposal.
- Creation of IST 201, Information Dynamics, as a formal course.
- Creation of an industry workshop on complex networked systems, to be offered as part of the ISTI winter course on information science and technology.

Years Three and Four

- Continued development of theory in computer science, communications and control motivated by theoretical and experimental work in the first two years of the proposal.
- Transition of new algorithms for robust control of networks to industry, through WAN in Lab validation and testing.
- Demonstration of distributed, cooperative real-time control on the Caltech multi-vehicle testbed, including dynamic resource allocation for communications and computing.
- Organization of ISTI Summer Course on Information Dynamics (SCID), to be offered in Summer 2005.

References

- [1] S. Deb and R. Srikant. Congestion control for fair resource allocation in networks with multicast flows. In *Proceedings of the IEEE Conference on Decision and Control*, December 2001.
- [2] R. M. Murray *et al.* Online control customization via optimization-based control. In T. Samad and G. Balas, editors, *Software-Enabled Control: Information Technology for Dynamical Systems*. IEEE Press, 2003. To appear.
- [3] J. A. Fax. *Optimal and Cooperative Control of Vehicle Formations*. PhD thesis, California Institute of Technology, Control and Dynamical Systems, 2002.
- [4] J. A. Fax and R. M. Murray. Graph laplacians and stabilization of vehicle formations. In *Proc. IFAC World Congress*, 2002.
- [5] J. A. Fax and R. M. Murray. Information flow and cooperative control of vehicle formations. In *Proc. IFAC World Congress*, 2003.
- [6] H. Feng, Q. Zhao, and M. Effros. Network source coding using entropy constrained dithered quantization. In *Proceedings of the Data Compression Conference*, Snowbird, UT, 2003. To appear.
- [7] Y. Frank-Dayana and R. Zamir. Dithered lattice-based quantizers for multiple descriptions. *IEEE Transactions on Information Theory*, 48(1):192–204, January 2002.
- [8] V. Gupta, R. M. Murray, and B. Hassibi. Stochastic analysis of information flow in multi-agent systems. In *Proc. IEEE Conf. on Decision and Control*, 2003. In preparation.
- [9] B. Hassibi and A. F. Dana. On the power efficiency of sensory and ad-hoc wireless networks. In *36th Asilomar Conference on Signals, Systems and Computers*, 2002.
- [10] B. Hassibi and B. Hochwald. How much training is needed in a multi-antenna wireless link? *IEEE T. Information Theory*, 2003. To appear.
- [11] B. Hassibi and M. Sharif. On the asymptotic capacity of multi-antenna broadcast channels with minimal channel-state-information. *IEEE T. Information Theory*, 2003. In preparation.
- [12] H. Ishii and B. A. Francis. *Limited Data Rate in Control Systems with Networks*, volume 275 of *Lecture Notes in Control and Information Sciences*. Springer-Verlag, 2002.
- [13] Koushik Kar, Saswati Sarkar, and Leandros Tassiulas. Optimization based rate control for multirate multicast sessions. In *Proceedings of IEEE Infocom*, April 2001.
- [14] Frank P. Kelly, Aman Maulloo, and David Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of Operations Research Society*, 49(3):237–252, March 1998.
- [15] Srisankar Kunnipur and R. Srikant. End-to-end congestion control schemes: utility functions, random losses and ECN marks. In *Proceedings of IEEE Infocom*, March 2000. <http://www.ieee-infocom.org/2000/papers/401.ps>.
- [16] B. Lincoln and B. Bernhardsson. Optimal control over networks with long random delays. In *Mathematical Theory of Networks and Systems*, 2000.

- [17] Steven H. Low. A duality model of TCP and queue management algorithms. *IEEE/ACM Trans. on Networking*, to appear, October 2003. <http://netlab.caltech.edu>.
- [18] Steven H. Low and David E. Lapsley. Optimization flow control, I: basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, December 1999. <http://netlab.caltech.edu>.
- [19] Steven H. Low, Larry Peterson, and Limin Wang. Understanding Vegas: a duality model. *J. of ACM*, 49(2):207–235, March 2002. <http://netlab.caltech.edu>.
- [20] L. Massoulie and J. Roberts. Bandwidth sharing: objectives and algorithms. In *Infocom'99*, March 1999. <http://www.dmi.ens.fr/~Emistral/tcpworkshop.html>.
- [21] M. B. Milam, R. Franz, and R. M. Murray. Real-time constrained trajectory generation applied to a flight control experiment. In *Proc. IFAC World Congress*, 2002.
- [22] Jeonghoon Mo and Jean Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, October 2000.
- [23] R. M. Murray, editor. *Control in an Information Rich World: Report of the Panel on Future Directions in Control, Dynamics and Systems*. SIAM, 2003. To Appear. Available at <http://www.cds.caltech.edu/~murray/cdspanel>.
- [24] A. Sahai. *Anytime Information Theory*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [25] L. J. Schulman. Coding for interactive communication. *IEEE T. Information Theory*, 42(6):1745–1756, 1996.
- [26] S. Tatikonda. *Control Under Communication Constraints*. PhD thesis, Massachusetts Institute of Technology, 2000.
- [27] Jiantao Wang, Lun Li, Steven H. Low, and John C. Doyle. Can TCP and shortest-path routing maximize utility? In *Proc. of IEEE Infocom*, April 2003. <http://netlab.caltech.edu>.
- [28] L. Xiao, M. Johansson, H. Hindi, S. Boyd, and A. Goldsmith. Joint optimization of communication rates and linear systems. *IEEE Transactions on Automatic Control*, 48(1):148–153, 2003.
- [29] R. Zamir and M. Feder. Information rates of pre/post filtered dithered quantizers. *IEEE Transactions on Information Theory*, 42:1340–1353, September 1996.
- [30] J. Ziv. On universal quantization. *IEEE Transactions on Information Theory*, IT-31(3):344–347, May 1985.