Laboratory for Information and Decision Systems

The Laboratory for Information and Decision Systems (LIDS) is an interdepartmental laboratory for research and education in systems, communication, control, and signal processing.

LIDS is staffed by faculty, research scientists, postdoctoral fellows, and graduate students drawn principally from the Department of Electrical Engineering and Computer Science, as well as the Department of Aeronautics and Astronautics, the Department of Mechanical Engineering, and the Sloan School of Management. Every year, many research scientists from various parts of the world visit the laboratory to participate in its research programs.

Twenty faculty, several research staff, and approximately 130 graduate students are presently associated with the laboratory. Undergraduate students also participate in research and thesis activities through the Undergraduate Research Opportunities Program. A number of postdoctoral and visiting appointments are made annually.

The lab has had a very exciting and successful year in AY2004. It is poised for greater achievements in a state-of-the-art research facility—the Ray and Maria Stata Center for Computer, Information, and Intelligence Sciences—in the upcoming year.

View from the window of the LIDS student lounge, on the 6th floor of the Stata Center's Dreyfoos Tower

Highlights

In AY2004, LIDS maintained its visibility in teaching and research. Research volume continued to remain strong at just over $6.5 million.

View of the LIDS student lounge, 6th floor

During the spring term, LIDS moved into its new location in the Stata Center. The new space has significantly enhanced the quality of life in the lab. One of its most popular spots is the double-height student lounge located on the 6th floor. The lounge's open and flexible design has made this area an ideal location for the lab's social events, casual conversations, and interactions between faculty, students, and staff.

Through the generosity of Mr. Patrick Wang, the Wang Shui Chung Laboratory that houses the Athena computing cluster and unmanned air vehicles (helicopter) research program provides valuable computing and essential research space.

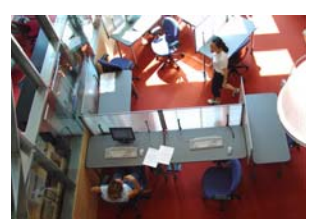

The Wang Shui Chung Laboratory, 6th floor

Several LIDS personnel received recognition and honors for their work over the past year.

Faculty

Professor David Forney, elected a member of the National Academy of Sciences in 2003, served as chair of the Electronics Section of the National Academy of Engineering.

Professor Lizhong Zheng received the Institute of Electrical and Electronics Engineers (IEEE) Information Theory Society 2003 Paper Award and the NSF CAREER Award.

In 2004, assistant professor Moe Win was named Young Aerospace Engineer of the Year by the American Institute of Aeronautics and Astronautics (AIAA), awarded the Fulbright Foundation Senior Scholar Lecturing and Research Fellowship, and elected fellow of the IEEE "for contributions to wideband wireless transmission."

Professor Alan Willsky received the 2004 IEEE Donald G. Fink Prize Paper Award for the paper "Multiresolution Markov Models for Signal and Image Processing," published in the *Proceedings of the IEEE* 90(8), August 2002.

Assistant professor Asuman E. Ozdaglar received the Graduate Student Council Award for Teaching in the School of Engineering.

Associate professor Muriel Médard was the 2004 winner of the Harold E. Edgerton Faculty Achievement Award.

Professor Sanjoy K. Mitter held the Russell-Severance-Springer chair in the Department of Electrical Engineering and Computer Science, University of California–Berkeley in the fall of 2003.

Students

Guy Weichenberg received the Natural Sciences and Engineering Research Council Award from Canada.

G. Weichenberg, along with professors Vincent W. S. Chan and Muriel Médard, was awarded the Best Paper Prize for "A Reliable Architecture for Networks under Stress" at the 4th International Workshop on the Design of Reliable Communication Networks, October 2003, Banff, Alberta, Canada.

Tom Schouwenaars received the AIAA Unmanned Aerial Vehicles Award.

Watcharapan Suwansantesuk received the Morris Joseph Levin Masterworks Award for best thesis presentation. His MS thesis is entitled "Multipath Aided Acquisition."

Aurelie Thiele won First Prize in the annual George Nicholson student paper competition at the INFORMS 2003 annual meeting in Atlanta, GA. The paper is entitled "A Robust Optimization Approach to Supply Chain Management."

Walter Sun received the Outstanding Student Paper Award at the American Geophysical Union 2003 fall meeting. The paper is entitled "Localization of Oceanic Fronts and Feature Boundaries Using a Variational Technique" and is coauthored by Dr. Mujdat Cetin, Professor W. C. Thacker (University of Miami), Dr. T. M. Chin (University of Miami and Jet Propulsion Laboratory), and Professor Alan Willsky.

Alex Ihler received the Best Student Paper Prize at the 2004 International Symposium on Information Processing in Sensor Networks. The paper is entitled "Nonparametric Belief Propagation for Self-Calibration in Sensor Networks" and is coauthored by Dr. John Fisher, Professor Alan Willsky, and Professor Randy Moses of Ohio State University.

Erik Sudderth received an Intel Doctoral Fellowship to complete his doctoral studies over AY2005.

Ramesh Johari won first place in the INFORMS Nicholson student paper competition for his work on game-theoretic aspects of network resource allocation.

Research Overview

Research at LIDS falls into four main areas, which share common intellectual bases. The Laboratory explicitly recognizes the interdependence of these fields and the fundamental role that mathematics, computers, and computation play in the research.

Communication and Networks. Research in this area includes fundamental work on networks, information theory, and communication theory. The work extends to applications in satellite, wireless and optical communications, and data networks. The objective is to develop the scientific base needed to design data communication networks that are efficient, robust, and architecturally clean. Wide area and local area networks, high-speed networks, and point-to-point and broadcast communication channels are of concern. Topics of current interest include network architectures at all layers; power control; multiple antenna techniques; network coding; media access control protocols; routing in optical, wireless, and satellite networks; quality of service control; failure recovery; topological design; and the use of pricing as a mechanism for efficient resource allocation.

Statistical Signal Processing. This group analyzes complex systems, phenomena, and data subject to uncertainty and statistical variability. Research spans the spectrum from broadly applicable basic theory, methodologies, and algorithms to challenging applications in a broad array of fields. Recent applications for this research include multisensor data assimilation for oceanography, hydrology, and meteorology; biomedical image analysis; object recognition and computer vision; and coordinated sensing and processing of large, distributed arrays of microsensors.

Optimization. Work in this area looks at analytical and computational methods for solving broad classes of optimization problems arising in engineering and operations research. It has applications in communication networks, control theory, power systems, and computer-aided manufacturing. In addition to traditional subjects in linear, nonlinear, dynamic, convex, and network programming, there is an emphasis on the solution of large-scale problems, including the application of neurodynamic programming methods.

Control and System Theory. The control systems group deals with problems related to complete systems analysis design. These include learning and system identification, controller design and optimization, and basic analysis of distributed systems involving the interaction of information and control. Theoretical research quantifies the fundamental limitations and capabilities of learning and feedback control for various classes of systems in the presence of dynamic uncertainty. Application-oriented work includes control architectures for single and multiple unmanned aerial vehicles and controllers for piloting epitaxy in semiconductor manufacturing. The control group is also involved in a research effort focusing on modeling aspects of the nervous system, conducted in collaboration with other laboratories.

Research Areas

Communication and Networks

Optical Networks

Professors Chan, Gallager, Médard, and Modiano continue to work on future optical network architectures. The focus of the program is to design and prototype the next-generation local and metropolitan area access network (MAN) with an increase in data rate of up to four orders of magnitude, but at the same time to decrease the cost of delivery per bit by approximately the same amount.

The network will use multiple wavelengths (colors) to increase capacity and optical devices for routing and switching. Professor Chan and his students obtained new results on the use of mesh topology for MANs. These results indicate that the cost structure of MANs is heavily dependent on the optical cross-connect technology at the switching nodes. Efficient topologies that nearly achieve fundamental bounds on performance are found, and these architectures are very different from previously used and accepted architectures.

Professor Modiano continues to work on an NSF grant to study mechanisms for providing optical bypass in the internet of the future. To this end, a number of techniques have been developed. Traffic grooming algorithms have been developed to selectively multiplex multiple low-rate traffic streams onto wavelengths; by this method, the number of wavelengths that must be processed at each node is minimized. Topology reconfiguration algorithms have been developed to reduce the load on the electronic switches and routers via dynamic load balancing. In addition, restoration algorithms are being developed that are optimized across the optical and electronic layers.

Professors Chan and Modiano continued a program sponsored by the Defense Advanced Research Projects Agency (DARPA) on all-optical, local, and metro area networks with ultrahigh reliability and performance. The objective of this research is to use optical network technology to build a highly reliable network that services high-end applications, such as aircraft control and coherent collaborative sensing. This program emphasizes cross-layer network optimization and nontraditional networks that provide a faster response than current networks, along with arbitrarily high delivery reliability. In the last year, we have found new transmission and detection schemes that make use of reliable physical topologies that lend themselves to ultrahigh reliability applications such as aircraft controls and cost-effective physical architectures for local and metro optical networks.

Professor Médard, in collaboration with her students, is working on issues of reliability and robustness of backbone and access networks. Her first project is in the area of probabilistic analysis of optical network robustness as part of an Air Force Office of Scientific Research (AFOSR) university research initiative with Stanford University, the University of Illinois, and Caltech. Recent results in this area include results for robustness for both regular and nonregular random graphs.

The Birkhoff-von Neumann (BVN) strategy for switch scheduling does not support multicast capabilities of optical switches, as it considers only permutation-based switch configurations. Professor Médard and her postdoc Dr. Supratim Deb, with their student, have extended the BVN strategy to support multicast flows in addition to unicast.

Professor Médard and her students are also working on reliability of access networks. She is the MIT member of a recent NSF Information Technology Research (ITR) project with the University of Illinois in the area of robust optical local and metropolitan area networks. This project, conducted in collaboration with Professor Chan and the DARPA CSWDM (chip-set, wavelength division multiplexing) project, considers the use of coarse unit of measure and limited signal-to-noise ratios in architecting robust networks. Results in this area concern the optimal design of local area networks with a limited number of nodes. The results have extended the understanding of circulant graphs for low probability of failure. Moreover, they have shown that, under extreme stress conditions for which failure probability rises, the characteristics of graphs required to ensure robustness are very different than in the low-failure probability mode. The results have also provided analytical results for networks with correlated failures.

Satellite Communications and Networking

The overall goal of this research addresses architecture designs for efficient data communications over satellites interconnected with terrestrial fiber and wireless systems to form a heterogeneous global internet. There are five main components to this research:

- Satellite constellation interconnection topology
- Satellite node architecture designs
- Adaptive power and rate control techniques for satellite communication systems over time-varying satellite channels to achieve greatly improved (an order of magnitude or more) data throughputs
- Efficient routing algorithms over a time-varying integrated and heterogeneous global network for maximum resource utilization, especially the space segments
- Efficient congestion control algorithms at the transport and network layers for an integrated satellite/terrestrial network

Professors Chan, Médard, Modiano, Tsitsiklis, and Win, doctors Chapin and Finn, and their students are conducting this research.

Professor Chan and his students completed a study on the preferred constellation topology of the space backbone. Based on providing coverage for spaceborne users alone, the minimum number of satellites in the backbone constellation can be as small as 12 for low-earth orbiting satellites (LEOs), six for medium-earth orbiting satellites (MEOs), and the usual three for geostationary-earth orbiting satellites (GEOs). When the cost of the space relay network is considered with crosslink, switching, and traffic demands as contributing factors, a GEO backbone of four to six satellites has the lowest cost and makes the most sense architecturally.

The team also completed a study to reach a baseline understanding of current traditional space-qualified processors and commercial processors. Given that radiation-hardened processors are at least one order of magnitude slower (approximately 7–10 years slower) than commercial processors, the team's vision of using commercial processors as replenishable, networked spaceborne processors may suggest a pathway to significant improvement of computation power in space. Many overhead intelligence functions can benefit from such an approach by reducing the data volume to be downlinked for processing. Subsequently, Professor Modiano's group has begun to analyze the architecture of a networked processing system in space and the associated problem of task scheduling.

Professors Chan and Modiano's groups are working on resource allocation schemes for satellite channels. This past year, they developed a novel auction algorithm for the distributed and fair allocation of a satellite fading (time-varying) channel. They also developed energy-efficient transmission scheduling schemes that take into account channel conditions in deciding when and how fast to transmit data in multiple-beam downlink antenna systems. They have also developed scheduling algorithms for multiple-beam satellite fading channels that can be used to provide quality-of-service guarantees on throughput and delay.

On transport protocols for hybrid networks, Professor Modiano and his students have explored the interaction between protocols at different layers. In particular, they developed a model for the interaction between transmission control protocol (TCP) and the Aloha multiple-access protocol and showed that channel "collisions" due to the Aloha protocol result in TCP window closures that significantly degrade end-to-end network performance. Using their models, they are able to optimize protocol parameters to achieve significant improvement in network performance.

In addition, Professor Modiano continues to work with NASA on developing intelligent architecture and protocols for NASA's future space internet. A particular focus of this project will be the development of communication architectures for supporting NASA's planned space exploration initiative.

Wireless Communications and Networking

Communication for Challenging Environments

Professor Win and his graduate students are working on the application of mathematical and statistical theories to communication, detection, and estimation problems. Specific current research topics include measurement and modeling of time-varying channels, design and analysis of multiple antenna systems, ultrawide bandwidth (UWB) communications systems, optical communications systems, and space communications systems. The group

- Developed an analytical framework for transmitted-reference signaling for UWB communications;
- Developed an analytical framework for evaluating the performance of practical diversity systems with nonideal channel estimates;
- Analytically derived tight upper and lower bounds on the inverse bit error probability for diversity reception in fading. The new bounds enable the derivation of the bit error outage (BEO) and normalized throughput for slow adaptive M-ary quadrature amplitude modulation with diversity reception. Their results show that the SAM technique can provide substantial increase in throughput with respect to fixed schemes while maintaining an acceptably low BEO.

Professor Forney has been doing research in the areas of syndrome realizations of linear codes and systems, constraint complexity of cycle-free graph realizations of linear codes, and the role of minimum mean-square error estimation in approaching the information-theoretic limits of linear Gaussian channels.

Wireless Wideband Communications

Professors Zheng and Médard made progress in the study of wideband communication over wireless fading channels. A novel approach is developed to understand the interactions among a number of system parameters, including the signal-to-noise ratio, the energy and spectral efficiency, constraints on the peak power for the signaling, and different channel modeling assumptions.

MIMO Coding and Relay Networks

Professor Zheng has also started a new research project on the noise accumulation over successive transmissions. This work is useful in the study of multihop communication schemes over relay networks, especially in a wireless fading environment. This research makes connections between the existing studies on network information theory and designs of structured codes used in multiple-input-multiple-output (MIMO) channels. The goal of this research is to build fundamental frameworks to study the costs of performance for using simple and distributed algorithms.

Multiple-Access Wireless Channels

Professor Médard is working in the area of capacity and stability of coded packetized multiple-access channels with students at MIT and with Professor Sean Meyn of the University of Illinois–Urbana-Champaign and Professor Andrea Goldsmith of Stanford University. This research establishes the capacity of such channels and examines trade-offs between energy and delay. It allows the uncoordinated access in satellite networks of multiple users without requiring total performance in the event of a packet collision.

In the area of energy-efficient channels, Professors Zheng and Médard, with their students, have considered nodes in which energy is used for purposes other than just radiated communications signaling. For such nodes, they have shown that capacity-achieving transmissions may need to use bursty transmissions rather than the continuous approaches dictated by traditional information-theoretic frameworks. This research has significant impact for the deployment of wireless sensor nodes, since it suggests that interference may not be a significant hindrance in setting up wireless sensor networks.

Communication Under Channel Uncertainty

Professor Médard has been investigating several issues in the area of wireless communications over uncertain channels. In collaboration with Professor R. Srikant at the University of Illinois, she has investigated the effect of unequal channel knowledge at the sender and receiver. In collaboration with Professor Goldsmith of Stanford University, Professor Médard has investigated the capacity of time-varying channels with sender- and receiver-side information, in particular channels with perfect-side information but significant intersymbol interference, for which no capacity formulas existed. In collaboration with Professor Ibrahim Abou-Faycal of the American University of Beirut and Professor Madhow of the University of California–Santa Barbara, she is working on the use of an adaptive modulator without feedback in which the sender adapts to the quality of the receiving channel measurement as well as the channel strength.

Professors Médard and Zheng, with their students, have established a new, practical way of transmission over ultrawideband channels. Their work has discovered a significant family of signals that can achieve capacity under infinite bandwidth limits. They have also recently shown, with Professor Tse of UC–Berkeley, that capacity of fading channels does not necessarily approach the wideband limit as slowly as previously thought but is instead very significantly affected by the relative values of SNR and coherence in time and frequency.

Professors Médard and Abou-Faycal and their students have established new results relating bandwidth and error probability for ultrawideband fading channels. Their results show that error probability decreases very slowly with bandwidth and therefore, unlike nonfading channels, infinite-bandwidth performance cannot be achieved in the finite bandwidth regime.

Wireless Ad Hoc Networks

Professor Modiano continues to work with Draper Laboratory on management and control of mobile ad hoc networks. Such networks are of critical importance for future combat systems, sensor networks, and autonomous systems involving mobile ground and air vehicles. These systems heavily depend on cooperative control between mobile vehicles and consequently on the availability of a communication capability between the vehicles. In a dynamic, mobile environment, one cannot assume that such communication capabilities are always present. In this effort, Professor Modiano and his students are developing architectures and protocols for providing reliable communication in this environment. Last year, the team developed cooperative routing schemes that take advantage of network layer diversity to increase network reliability over a wireless network with fading channels. They introduced a novel "outage probability" model for measuring reliability in a wireless network and devised network routing algorithms for minimizing network outage probability.

This year, Professor Modiano, together with researchers from the University of Illinois, received a grant from NSF to develop efficient resource management schemes for wireless networks of autonomous air vehicles. The team is studying the interplay between competing requirements for communication, computation, and mobility. They have also developed a fundamental relationship between network throughput and delay in a mobile ad hoc network. Using a simple node mobility model, they established that the ratio of delay to throughput must be greater than the number of nodes in the network and developed routing protocols based on redundant transmissions that achieve this bound.

In addition, Professors Modiano, Eric Feron, and Nancy Lynch, along with Dr. Jinane Abounadi, are collaborating on a multidisciplinary university research initiative with Stanford University and the University of Illinois on Cooperative Networked Control of Dynamical Peer-to-Peer Vehicle Systems. A major focus of the project is the interplay between communication and control in an environment of networked vehicles.

Professor Chan and his students developed a new model of wireless network in which processing energy is considered together with transmission energy and found that optimum routing strategies can be substantially different from previous results where processing energy is overlooked. Thus, nearest-neighbor routing is not usually the optimum strategy but one that takes into account node density, transmission environment, and bandwidth allocations/data-rates. Moreover, with the flexibility of inserting relay nodes or strategic movement of nodes, significant gain in network performance can be obtained.

Network Codes

Professor Médard, in collaboration with her students and Professor Ralf Koetter of the University of Illinois–Urbana-Champaign, is working on an algebraic description of codes on graphs for data transmission over networks. All routing over a network can be described as a code over that network. Moreover, network capacity in error-free networking can be significantly enhanced through the use of codes over these networks. This research is being extended to ensure robustness when links or nodes are permanently removed. Thus, it provides results on the fundamental network management requirements for recovery from nonintermittent failures.

Professor Médard, in collaboration with her students and with Professor David Karger of the Computer Science and Artificial Intelligence Laboratory (CSAIL), Professor Koetter of the University of Illinois–Urbana-Champaign, Professors Michelle Effros and Babak Hassibi of Caltech, and Dr. Abounadi of MIT, is working on using linear network codes as a unified framework for source, channel, and network coding. They have demonstrated that source-channel coding (or source-network) separation holds for several canonical network examples when the whole network operates over a common finite field. The researchers' simple, unifying framework for these codes not only allows them to reestablish with economy the optimality of linear codes for single transmitter channels and for Slepian-Wolf source coding; it also enables them to establish the optimality of linear codes for multiple-access and for erasure-broadcast channels. The researchers have illustrated the fact that design for individual network modules may yield poor results when such modules are concatenated, making end-to-end coding necessary. Thus, they argue, it is the lack of decomposability into canonical network modules rather than the lack of separation between source and channel coding that presents major challenges for network coding.

Sensor Web, Interference, Coding, and Statistical Mechanics

Recent research on turbo coding, decoding of low-density parity check codes, and statistical mechanisms of disordered systems has shown that there are deep connections between these subjects. Professor Sanjoy Mitter, in collaboration with Dr. Nigel Newton, has recently given an information theoretic view of maximum-likelihood decoding and nonlinear estimation of diffusion processes. In the recently completed thesis of Maurice Chu, a unified view of distributed estimation with application to the sensor web has been presented. In the completed thesis of Louay Bazzi, various aspects of coding and complexity have been investigated. Finally, Professor Mitter, with Dr. Reuben Rabi, is developing a theory of interconnections that has applications in distributed control, coding, and inference on graphs.

The Interaction of Information and Control

Professor Mitter and colleagues Professor Nicola Elia of Iowa State University, Professor Sekhar Tatikonda of Yale University, and Professor Anant Sahai of UC–Berkeley have been working on fundamental issues related to the interaction of information and control. The major goal of this research is to develop a dynamical (nonequilibrium) view of information theory that is relevant when channels are used not just for communication but also control. There appear to be deep connections with current research in nonequilibrium statistical mechanics as developed by Gallavotti, Ruelle, and others.

MIT–Cambridge Project on Quantum Information

Professor S. K. Mitter has worked on developing a theory of quantum trajectories, which is essential in understanding quantum control.

Statistical Signal Processing

Stochastic Systems Group

The Stochastic Systems Group (SSG) is led by Professor Alan S. Willsky, with the assistance of Dr. Mujdat Cetin of LIDS and Dr. John Fisher of CSAIL. In addition, the group includes 10 to 12 graduate students, visitors, and participants from other groups within LIDS and from other MIT laboratories and departments.

The general focus of research within SSG is on the development of statistically based algorithms and methodologies for complex problems of information extraction and analysis from signals, images, and other sources of data. The work in the group extends from basic mathematical theory to specific areas of application. Current applications include biomedical image analysis, oceanographic and geophysical data assimilation, and fusion of multisource (e.g., acoustic and video) information both in centralized processors and in power-limited distributed sensor networks. Funding for this research comes from a variety of sources, including the Office of Naval Research (ONR), AFOSR, the Army Research Office (ARO), Office of the Director, Defense Research and Education (through AFOSR, ARO, and ONR), the National Institutes of Health, NSF, MIT Lincoln Laboratory, and Royal Dutch Shell Corporation.

In addition to directing these research activities, Professor Willsky is very active in supporting government and, in particular, Department of Defense organizations in assessing and planning technology investments. He previously served as a member of the Air Force Scientific Advisory Board and continues to support that organization informally. Each of the following research areas being pursued within SSG involves both theoretical development and applied studies to the application areas mentioned previously.

Multiresolution Statistical Signal/Image Processing and Graphical Models

For a number of years, SSG research in multiresolution statistical and image processing has received considerable international attention and has found application in an extremely wide range of disciplines. These disciplines extend well beyond those in which the group has been and continues to be involved (namely large-scale geophysical data assimilation, computer vision, and distributed sensor networks) to fields such as chemical engineering and biomedical imaging. Because of the wide use of and interest in these methods, Professor Willsky was invited to prepare a tutorial/survey paper on this field, which appeared in the *Proceedings of the IEEE* and for which he received the 2004 IEEE Donald G. Fink Prize Paper Award. This area remains one of the most vigorous components of SSG research activities, with significant expansions in the scope of inquiry.

The key to SSG's previous (and some of its current) research in this area is the direct statistical modeling of phenomena at multiple resolutions using graphical models on trees and other graphs, in which each level on a tree corresponds to a particular resolution. Using these models, the group has developed very efficient algorithms for estimation, data fusion, and other image analysis tasks and has also demonstrated that a wide variety of real phenomena and applications can be captured within this framework. They continue to work on expanding the domain of applicability of this methodology by pursuing additional applications and by developing tools for constructing multiresolution models needed as the basis for applying their results. For example, the group's multiresolution methods are a key component of the five-year NSF ITR program on large-scale data assimilation for geophysical processes, which involves collaboration with MIT researchers in Civil and Environmental Engineering, Earth and Planetary Sciences, and Electrical Engineering and Computer Science.

In addition, major new thrusts involve the investigation of how SSG can exploit its methodology for problems involving much more complex graphical models, such as those that arise in military command and control, or in problems of monitoring complex systems—problems of great national concern because of the need to make critical national infrastructure secure. SSG's work involves the examination of graphical structures more complex than trees—in particular, graphs containing (typically many) loops. The group's approach aims at the exploitation of embedded, acyclic tree structures in these graphs.

SSG has had several major successes in this area, including a new class of signal and image processing algorithms that expands the range of applicability of the group's method considerably and new sets of algorithms for inference on loopy graphs that have provable performance properties, outperform previously developed methods, and have led to several awards for recent SSG graduate, Dr. Martin Wainwright. In addition, the group is now heavily involved in extending these methods to problems of distributed fusion in sensor networks, in which the relevant graphical structure captures both the statistical relationships among the variables of interest as well as the network connectivity among sensing, processing, and communication nodes. Our most recent algorithms in this area—combining methods from nonparametric statistics with so-called belief propagation algorithms for graphical models—has received considerable attention and recognition, including the recent Best Student Paper Award received by current SSG doctoral student Mr. Alexander Ihler. In addition, the codeveloper of this new inference method, Mr. Erik Sudderth, who is also an SSG doctoral student, recently received an Intel doctoral fellowship to complete his doctoral dissertation on the development of new approaches to inference on graphical models and their application to problems in computer vision.

Nonlinear and Geometric Image Analysis

During the past year, SSG has continued its efforts in the area of nonlinear/non-Gaussian image analysis. These include the explicit estimation/extraction of geometric information, such as object boundaries and segmentation. The group's work continues to focus on the development of statistically based curve evolution algorithms. Such algorithms involve explicitly defining and dynamically evolving curves in ways that lead to accurate and efficient segmentation of images. Methods of this type that have been developed by others have a number of very attractive features, including the fact that they provide seamless ways in which a curve could separate into multiple curves or merge from several disjoint curves to a single curve, allowing automatic and easy segmentation of multiple regions of interest.

The group's work aims at developing first-principles statistical approaches to curve evolution that deal with noise and variability in a statistically optimal way. Recent accomplishments include methods that use training examples to build and then use statistical models for shapes of interest, learning-based methods that perform segmentation while simultaneously learning the statistical differences between the regions being segmented, and the tracking of dynamically evolving curves. This work is being pursued in the context of several significant interdisciplinary efforts. The first of these involves collaboration with clinicians and researchers at Brigham and Women's Hospital toward applications in automatic interpretation of medical imagery and image-guided therapy procedures. Another major interdisciplinary effort involves problems of dynamically estimating and tracking curves, with applications including both 4-D medical imaging of the heart and the tracking of major fronts in meteorology and oceanography (e.g., the Gulf Stream). Recent work in this area led to a paper, presented at the 2003 American Geophysical Union Fall Meeting, for which Mr. Walter Sun, a doctoral student in SSG, received the Best Student Paper Award.

Machine Learning and Information-Theoretic Methods for Multisensor Fusion

A third component of SSG's research program involves the development of statistically based algorithms for the fusion of information from multiple sensors in the presence of substantial uncertainties—for example, in the nature of the signals being sensed (such as acoustic signatures of unknown character), in the number of sources generating those signals, in the locations and calibration of the sensors themselves, and in the relationships among the signals being sensed by sensors of very different modalities.

The group's work has found application in medical image analysis, in particular in functional magnetic resonance imaging and in a variety of multisensor fusion applications. These include fusion of audio and video sensors (e.g., for the localization of acoustic sources in video scenes) and fusion of multiple acoustic sensors for the detection and localization of multiple sources in complex and highly uncertain environments that defeat standard coherent processing methods. SSG has demonstrated that information-theoretic methods can provide robust solutions to such problems without the need for *any* training (i.e., fusion is performed "on-the-fly" as data are collected). Current applications of this work include new methods for "blind" segmentation of imagery (i.e., without prior models for the statistics of foreground and background), as well as methods for detecting dependencies among multiple variables without prior knowledge of the statistical nature of possible dependencies.

Neurobiological Modeling

Professors Munther Dahleh and Steve Massaquoi are interested in three problems. The first is the development of a hierarchical model of the interaction between the cerebrum and cerebellum that is anatomically justified and that can explain multivariable dynamic stabilization and control. The second problem is deriving a multiscale, multiresolution model that explains electroencephalography (EEG) data, with specific interests in motor control, anesthesia, and evaluation of cortical function and dysfunction. These projects are in collaboration with various laboratories/departments at MIT as well as Massachusetts General Hospital (MGH). The third is the development of a circuit model of basal ganglia that describes the basal ganglia's function in both low-level control of movement speed and in motor programming.

Substantial progress was made in the area of developing reduced-order models for the cerebellum and its interactions with the cerebrum and spinal cord. Progress has been made in utilizing these models for interpreting speed and directional information present in actual cerebellar data. According to collaborator Dr. Timothy Ebner of the University of Minnesota, a cerebrocerebellar control model now appears to explain observed input-output behavior as well as approximate many neural signals observed in vivo.

In a parallel effort concerning modeling EEG data, professors Dahleh and Massaquoi have developed a basic circuit that constitutes a fundamental cerebral function module. The circuit describes local and global interconnections between the different layers and has been successful in simulating several important states of the brain. This development is unique, and the professors expect several interesting fundamental models to emerge. The work is done in collaboration with Professor Dahleh's postdoctoral associate, Fadi Karameh, and Dr. Emery Brown (MGH). The objective of this research is to utilize such a model to classify different sleep stages while applying anesthesia, detect structural and functional aberrations in the cerebrum, and ultimately gain insight into the mechanisms of cognition.

Progress has also been made in developing a unified model of basal ganglia function that interprets the structure as implementing a logical operator that enables programmed control of behavior ranging from cruising movements to cooperative interaction with the environment.

Perceptual Systems

Professor Sanjoy Mitter and his collaborators, Professor Stefano Soatto of UCLA and Dr. Horst Haussecker of Intel, have been working on various aspects of perception and recognition. Perception and recognition involve recovering useful information about the environment from sensed data and prior knowledge about the real world and the sensors. Artificial systems designed to carry out this task are much inferior to biological systems, largely due to the size and intricacy of the knowledge required to carry out reliable inference in unrestricted and uncertain domains. For instance, in visual perception, several factors contribute to render the problem difficult: clutter, occlusion, and variability of the objects in the scene. The basic engineering principle of decomposing a complex task into simpler and independent tasks is difficult to apply to perception and recognition due to the extremely complicated and yet unknown patterns of interdependency among the many “acts of perception” involved. For example, the recognition of an occluded chair in a cluttered office environment is highly dependent on the interpretation of its subparts, the other objects near it, and the overall scene of which it is part.

What are the components involved in perception and recognition? Into what architecture should these components be organized? How does one minimize the interdependence of these components? How should uncertainty be represented? How does one acquire and represent the knowledge about the real world and the sensors? Several projects are being undertaken to find answers to these questions.

A new computational theory for the recognition of occluded deformable templates in a cluttered scene has led to efficient algorithms with guaranteed performance in terms of localization errors and time complexity. Currently, this approach has been applied to features consisting of points in the plane and to affine deformations. Future work will seek to generalize these assumptions.

Early recognition of moving ground targets from an approaching platform is an important task for the military. To enhance the performance of existing systems, it is necessary to combine information from multiple frames, which contain the target at different resolutions. This project is still at an early stage, and initial efforts have focused on the incorporation of continuity and smoothness constraints of the relative motion of the target with respect to the camera by means of a geodesic approach.

Automotive Safety

In 2002, LIDS became involved with developing safety-enhancement mechanisms for the automotive industry under Ford sponsorship. Under the Ford–MIT alliance, Professor Eric Feron has assumed the responsibility of developing and managing the safety research program of the alliance, along with investigators in CSAIL, the Center for Transportation and Logistics, the Department of Aeronautics and Astronautics, and LIDS. His research group focuses on the development of collision alerting systems for operation onboard a single vehicle or in a networked fashion.

Optimization

Algorithms

This project focuses on analytical and computational methods for solving broad classes of optimization problems arising in engineering and operations research, as well as for applications in communication networks, control theory, power systems, computer-aided manufacturing, and other areas. Currently, in addition to traditional subjects in nonlinear and dynamic programming, there is an emphasis on the solution of large-scale problems involving network flows, as well as on the application of decomposition methods. Professors Dimitri Bertsekas and John Tsitsiklis and their students perform this work.

Neurodynamic Programming

Problems of sequential decision making under uncertainty are all-pervasive; for example, they arise in the contexts of communication networks, manufacturing systems, and logistics and in the control of nonlinear dynamical systems. In theory, such problems can be addressed using dynamic programming techniques. In practice, however, only problems with a moderately sized state space can be handled.

This research effort deals with the application of neural networks and other approximation and interpolation methodologies to overcome the curse of dimensionality in real-world stochastic control problems. The objectives driving this research are twofold: first, to develop the theoretical foundations and improve the understanding of such methods using a combination of tools from approximation theory, dynamic programming, and stochastic algorithms; and second, to use these methods for solving some large-scale problems of practical interest. Application areas currently being investigated include problems in logistics (resource scheduling and assignment), finance (pricing of high-dimensional derivative instruments, dynamic portfolio management in the presence of risk constraints), supply chain management, and communications (dynamic channel allocation). Professors Bertsekas and Tsitsiklis and their students perform this work.

Fundamental Issues in Optimization

This research focuses on fundamental analytical and computational issues in (deterministic) optimization that are connected through the themes of convexity, Lagrange multipliers, and duality. The aim is to develop the core analytical issues of continuous optimization, duality, and minimax/saddle point theory using a handful of unifying principles that can be easily visualized and readily understood. Numerous research results on these topics are published in the new graduate-level textbook, *Convex Analysis and Optimization,* by Dimitri Bertsekas, with Angelia Nedic and Asuman Ozdaglar (Athena Scientific, April 2003).

As part of this research effort, a new development of Lagrange multiplier theory is explored based on a set of optimality conditions that are stronger and more general than the classical Karush-Kuhn-Tucker optimality conditions. These optimality conditions motivate the introduction of a new condition called *pseudonormality,* which emerges as central within the taxonomy of significant characteristics of a constraint set. This analysis also yields identification of special Lagrange multipliers, which carry significant sensitivity information regarding the constraints that directly affect the optimal cost change. Professors Bertsekas and Ozdaglar perform this work.

Network Optimization

Multicommodity network flow problems involve several types of supply/demand (or “commodities”), which simultaneously use the network and are coupled through either link capacities or through the cost function. This research considers linear/integer multicommodity flow problems for some special types of graphs, such as rings, that frequently arise in practical applications such as data communication networks. Professors Bertsekas and Ozdaglar and their students show that these problems can be polynomially solved without loss of optimality by relaxing the integer constraints and rounding the solutions.

One important context in which such problems arise is the routing and wavelength assignment (RWA) problem, which is critically important for increasing the efficiency of wavelength-routed all-optical networks. Given the physical network structure and the required connections, the RWA problem is to select a suitable path and wavelength among the many possible choices for each connection so that no two paths sharing a link are assigned the same wavelength. In work to date, this problem has been formulated as a difficult integer programming problem that does not lend itself to efficient solution or insightful analysis. In this work, the researchers propose several novel optimization problem formulations that can be addressed with highly efficient linear integer programming methods and that yield optimal or near-optimal RWA policies.

Game Theory and Communication Networks

The key characteristic of today's large-scale networks, such as the internet, is their decentralized nature. These networks have emerged from the interconnections of different networks, service providers, and users with different objectives and performance measures, and their operation relies on varying degrees of collaboration and competition among these entities. With the increasing commercialization and the imminent arrival of multiple classes of differentiated services, service provider objectives will become an integral part of the operation of networks. A combined study in which the interactions between economic incentives of service providers and users and traditional network control algorithms are understood is essential for the design of future networks. This research requires the combination of classical techniques of optimization theory with new tools from game theory, the study of multiagent decision making.

Our goal in this research is twofold: first, to provide a theoretical framework for understanding the interactions of service providers, users, and traffic engineering aspects of communication networks; and second, to develop and evaluate new computationally feasible algorithms and pricing schemes that can be implemented in practice for allocation of resources in the presence of users with variable quality of service (QoS) requirements. Toward this goal, we study different networks (treating the wire-line and wireless networks separately) and different pricing schemes (flat, usage-based, various QoS-pricing schemes). The large scale and complexity of the systems under study also necessitate the development of new tools in optimization theory and advances in game theory. Professor Ozdaglar and her students perform this research.

Game Theory in Resource Allocation

Professor Tsitsiklis, together with his students and coworkers, have been studying simple, market-based mechanisms for resource allocation in various settings. Particular contexts of interest include allocating bandwidth to users in a communication network, as well as allocating production levels in day-ahead markets for electric power generation. Different mechanisms have been evaluated in terms of the possible efficiency losses in the presence of strategic (gaming) behavior by market participants that hold significant market power. Models of network formation that capture negotiation and certain types of contractual agreements are also been studied.

Control and System Theory

Multivariable and Robust Control

The systematic design of multiple-input, multiple-output systems, using a unified time-domain and frequency-domain framework to meet accurate performance in the presence of plant and input uncertainty, is an extremely active research area in LIDS. Various theoretical and applied studies are being carried out by Professors Munther Dahleh, Eric Feron (chair of the IEEE Technical Committee on Robust Control), Steve Massaquoi, Alexandre Megretski, and their students.

Theoretical research deals with issues of robustness, aggregation, and adaptive control. The aim of the research is to derive a computer-aided design environment for the design of control systems that can address general performance objectives for various classes of uncertainty. Furthermore, new results on the robustness of nonlinear feedback systems, using feedback linearization, have been obtained for unstructured uncertainty model errors.

Recent application-oriented studies include the control of large space structures, helicopters, and submarine control systems; issues of integrated flight control; control of chemical processes and distillation columns; automotive control systems; and the modeling and analysis of biological control systems.

New applications for robust and programmed (finite state-based) control theory are now emerging at LIDS, including the real-time, agile guidance of single and multiple unmanned aerial vehicles, as well as vehicle anticollision problems arising in air traffic control. Some of these concepts are implemented and tested on small helicopter systems.

Professors Feron and Massaquoi are involved in a collaboration regarding the internal mechanisms that underlie the brain’s ability to acquire programs that manage external dynamics and communication. In a separate research track, Professor Feron has been investigating the real-time decision issues arising in highly dynamic, multiagent environments, using partner dancing as a case study. Together with a graduate student, Professor Feron has uncovered the essential role of gaming in competitive partnered dancing.

In a recent initiative, Professors Megretski and Feron have begun a new investigation about applying the tools of robust control systems analysis to real-time software analysis, with a specific focus on run-time error. They have established a collaboration with the group of Professor Patrick Cousot, Ecole Normale Supérieure, France, to facilitate this activity.

Evolutionary Control

Another new thrust regards the general question of how control systems might evolve over time to manage complex control problems. Professors Mitter, Dahleh, Massaquoi, and Berwick and postdoctoral associates Reuben Rabi and Fadi Karameh conduct this work. The hope is to understand principles common to self-optimizing control systems across multiple scales of time and space. Biology is used as the guiding example, with analysis of systems ranging from molecular biological control of metabolism to organ system interaction to ecological regulation.

Unmanned Air Vehicles

Professors Dahleh and Feron and their students have been working on developing control architectures for unmanned vehicles. They have derived an architecture for the autonomous controller that enables the vehicle to perform agile maneuvers. The basis for this architecture is the derivation of a robust hybrid automaton. This automaton describes a rich set of controlled trajectories that can be attained by the vehicle, as well as the control necessary to transition between these trajectories. The robustness analysis of this dynamical description gives rise to a new and exciting class of robustness analysis problems that has not been looked at in the literature.

The researchers have developed a complete simulation/animation environment, and their software (based on the above architecture) is now in use at Draper Laboratory, Barron Associates, Inc., and the Air Force Research Laboratory. A recent development in this problem is deriving efficient algorithms for performing real-time motion planning (contrasted from path planning, where vehicle dynamics are not taken into account) in a cluttered environment. These algorithms are based on randomization techniques performed on the manifold on which the dynamics evolve. This research entails the development of a hierarchical control system that replaces the human pilot in order to perform agile maneuvers.

As of 2004, most of this work has been or is in the process of being transitioned to industry (Nascent Technology, Lockheed Martin, Boeing) and users (Natick Army Laboratories). Eric Feron's students have obtained several awards related to this work, including a TR100 Award in the fall of 2003 (Vladislav Gavrilets) and the AIAA Unmanned Aerial Vehicle Systems Award (Tom Schouwenaars).

With sponsorship from Draper Laboratory, Professor John Deyst and his students are developing new guidance and control methods for operation of intelligent unmanned air vehicles (UAVs). This work addresses the coordinated action of groups of UAVs that operate together to accomplish complex tasks. Such coordinated action is required to accomplish tasks that are impossible, or would take excessively long periods of time, for a single vehicle to complete. Significant issues being addressed are the safe and effective flight of UAVs near each other, including rendezvous and docking of one vehicle with another. This capability is of particular significance for resupply of one vehicle by another, so as to allow sustained operation near some desired location that might be some distance from a user. Coordinated flight is also essential for integrating various kinds of information sensed by many vehicles simultaneously. This includes the deployment of ground-based wireless sensor networks from the UAVs, especially when precise placement of sensors is an important issue. The operational needs of this class of systems pose particularly stringent requirements on various aspects of vehicle guidance and control.

Identification and Learning of Complex Systems

Professor Dahleh has led a research effort in developing a theoretical framework for learning and identification of complex systems. To accurately define such a problem, one needs to make assumptions about the generation of data; choose a model class from which a model will be selected; and choose a metric that captures the distance between the model and the actual system. One can also choose multiple model classes and derive a metric to evaluate which model class to choose. Classical approaches suffer from many pitfalls. First, they assume that the system that generated the data belongs to the model class (or one of the model classes). Second, they assume that the data record is long enough that one can accurately estimate the actual system. In fact, model quality evaluation methods assume such asymptotic convergence in computing a metric evaluating different model classes.

Professor Dahleh and his students have developed a new theoretical framework in which undermodeling is explicit in the problem formulation. Equivalently, the process that generates the data is not a member of the model classes considered. This work began in Dr. Saligrama Venkatesh's thesis several years ago. Recently, Dr. Soosan Beheshti developed a new measure of model quality evaluation, Model Description Complexity (MDC), that is computed from finite data. MDC is the correct generalization of AIC and MDL methods used heavily for evaluating model sets. Recent work connects Beheshti's results with learning theory and universal coding.

Identification of Jump Parameter Systems

Many systems are best modeled as jump systems—systems that switch between relatively simple systems. Switching is controlled by a Markovian system. A hurdle in identifying such systems is the estimation of the sequence of switching from the continuous observations at the output. Professor Dahleh and his students have developed a new framework for analyzing such systems based on Shannon's channel coding theorem and distortion theory. This work is the topic of the PhD thesis of Nuno Martins.

Control, Communication, Computation

Communication channels impose constraints on feedback systems that limit the achievable closed-loop stability and performance. Control theory has focused on characterizing the fundamental limitations and capabilities of closed-loop systems in the presence of both plant and input uncertainty. Communication constraints introduce a new class of uncertainty (e.g., quantization, average bit rate, or capacity) that existing theory deals with only indirectly.

Professor Nicola Elia visited LIDS for a semester to help in this area. Professor Dahleh and his students, in collaboration with Professor Elia, have derived new results for computing stability limitations of feedback systems in the presence of various channels, using both deterministic and probabilistic models.

Analysis and Synthesis of Hybrid Systems

Many applications involve the interaction of both discrete (logic) and continuous systems. A feedback system with bit constraints is an example of such interaction. The motion planning problems of UAVs is another example. Professor Dahleh and Professor Megretski are leading an effort to derive a formal theory for modeling, analysis, and synthesis of pure discrete systems. This work is the first step toward the derivation of a complete formal theory for hybrid systems.

Model Reduction of Discrete and Hybrid Systems

This work is pioneered by Professors Dahleh and Megretski and their students. The objective is to develop a theory for model reduction of discrete systems represented by hidden Markov chains. Such systems arise frequently in enterprise models as well as hybrid systems.

Systems Aspects of Biology

Professor B. Mishra (Courant-NYU), professors Peter Doerschuk and Saul Gelfand (Purdue University), and professors D. Bertsekas and S. K. Mitter (MIT) initiated a collaborative research project involving an integrative view of systems biology, with special emphasis on the study of viruses interacting with *E. coli*.

Vincent W. S. Chan
Director
Joan and Irwin M. Jacobs Professor of Electrical Engineering and Computer Science and Aeronautics and Astronautics

More information about the Laboratory for Information and Decision Systems can be found on the web at http://lids.mit.edu/.