Laboratory for Information and Decision Systems

The Laboratory for Information and Decision Systems (LIDS) is an interdepartmental laboratory for research and education in systems, control, optimization, communication, networks, and statistical inference. These disciplines, which span the domain of the analytical information and decision sciences, play a critical and pervasive role in science, engineering, and society more broadly. LIDS provides a focal point and melting pot for disciplines that share a common approach to problems and a common mathematical base—the laboratory provides an energized environment that fosters the research needed for the future, and instills in its students the disciplinary depth and interdisciplinary understanding required of research and engineering leaders of today and tomorrow.

Faculty within LIDS are principally drawn from the Department of Electrical Engineering and Computer Science (EECS) and the Department of Aeronautics and Astronautics. However, because the disciplines in which LIDS is involved are also of great interest across the Institute, we have built and continue to build collaborations and interactions with many other units, including the Operations Research Center; the Computer Science and Artificial Intelligence Laboratory; the Research Laboratory of Electronics; the MIT Energy Initiative; the MIT Transportation Initiative; the Department of Civil and Environmental Engineering; the Department of Mechanical Engineering; the Department of Earth, Atmospheric, and Planetary Sciences; the Department of Brain and Cognitive Sciences; the Department of Economics; the MIT Sloan School of Management; the Media Lab; and the Harvard-MIT Division of Health Sciences and Technology.

LIDS leadership continues to grow within MIT as well as nationally and internationally, and LIDS faculty and students continue to receive substantial recognition for their contributions with numerous national and international awards. The LIDS website contains details on LIDS activities, people, awards, and research accomplishments and directions. In addition, LIDS and its faculty continue in leadership positions in major initiatives, including the Future Urban Mobility project within the Singapore-MIT Alliance for Research and Technology (SMART), and the Virtual Center on Connection Science and Engineering, and LIDS personnel are active participants in ongoing MIT-wide discussions of the future of statistics and of engineering systems.

LIDS researchers continue to have great success in obtaining funding for our broad and deep research agenda, and continue to develop our relationships with industrial organizations and national laboratories, including Draper Laboratory, Lincoln Laboratory, Los Alamos National Laboratory, Siemens, Shell Oil Company, Honeywell, Ford Motor Company, Société Nationale des Chemins de fer Français, Aurora Flight Sciences, and Microsoft Research. Also, thanks to a rich history of research excellence and leadership, LIDS remains a magnet for the very best, attracting not only outstanding students and faculty but also a continuing stream of world-leading researchers as visitors and collaborators.

LIDS continues to host major workshops and symposia in emerging and important areas, and in fall 2012, hosted the Second International Workshop on Information and Decision in Social Networks, following up on the very successful first workshop, held in May/June 2011.

Intellectual Vision and Research Areas

The domain of research in LIDS can be described along several different dimensions:

- A common set of mathematical disciplines, including probability and statistics, dynamical systems, and optimization and decision theory
- A set of core engineering disciplines—namely, inference and statistical data processing, transmission of information, networks, and systems and control
- A set of broad challenges, emerging applications and critical national or international needs, and intellectual opportunities

Research at LIDS involves activities within and across all these dimensions. The convergence of issues that arise in the challenges of the present and future has led to research that cuts across mathematical and engineering disciplines in new, exciting, and important ways. Research flows in both directions across these dimensions: work in each of the mathematical disciplines leads to new methodologies that enable advances in core disciplines and in interdisciplinary investigations; work in attacking those emerging interdisciplinary challenges provides direction and drivers for fundamental disciplinary activities and has led to the charting out of emerging new disciplines.

In particular, the availability of increasingly capable sensing, communication, and computation systems enables the collection and transfer of large amounts of data pertaining to complex and heterogeneous interconnected systems. The need for an intellectual platform to simultaneously address questions of data fusion, distributed learning, information transfer, and distributed decision making is stronger than ever as existing techniques fall short in addressing issues of scalability, robustness, and performance limits. Examples of areas in which LIDS research has and will continue to contribute include the following:

- Intelligence, surveillance, and reconnaissance systems
- Coordination of unmanned autonomous systems
- Energy information systems
- Biological systems and biomedical data analysis
- Large-scale data assimilation for the geosciences
- Network scheduling and routing
- Transportation network analysis, control, and design
- Machine learning for recommendation systems and social media
- Social network analysis and characterization
- Ultra-wideband and other emerging communications technologies

Furthermore, traditional paradigms of sensing, communication, and control are not adequate to address many emerging challenges. As a result, LIDS has initiated a set of fundamental research themes that cut across disciplinary boundaries and involve considerable interaction and collaboration with colleagues in other units at MIT and in other disciplines. These themes include:

- Foundations of network science, including network algorithms, approximation, and information
- Foundations of decision theory for teams involving cooperation and competition, including dynamic mechanism design in game theory, learning in stochastic games, and the study of rational decisions for large interacting networks of agents
- Foundations and a theoretical framework of systemic risk
- Foundations of cyber-physical systems, including architectural design, cross-layer algorithms, and tools for analysis, verification, and performance guarantees
- Foundational theory for multiscale/multi-granularity modeling, including methods for describing complex phenomena at multiple granularities, learning of such models from complex and heterogeneous data, and reducing/simplifying such models to levels appropriate for particular questions of analysis or design
- Foundations of scalable analysis and inference for problems involving "big data" in which the boundaries between inference, learning, data representation and access, and massively parallel algorithms are essentially nonexistent

Faculty Activities

Much of the major research activities of LIDS faculty not only cuts across the disciplines, applications, and emerging areas mentioned previously but also is collaborative with others within LIDS and elsewhere at MIT.

Dimitri Bertsekas

Professor Dimitri Bertsekas is interested in deterministic optimization problems and the role of convexity in solving them, possibly through the use of duality. He has written a textbook on the subject, which involves new research on the fundamental structures that guarantee the existence of optimal solutions while eliminating duality gaps. He has also written a companion textbook on convex optimization algorithms, which includes some of his research on problems whose cost function involves a sum of a large number of component functions: (1) separable large-scale convex optimization problems, known as extended monotropic programming problems, for which special duality results and algorithms are possible; (2) new polyhedral approximation algorithms for extended monotropic programming problems; (3) incremental subgradient and proximal methods; and (4) application of these methods to large-scale machine learning and energy production and distribution systems.

Professor Bertsekas also performs research on problems of sequential decision making under uncertainty, which are pervasive in communication networks, manufacturing systems, and logistics, and in the control of nonlinear dynamical systems. In theory, such problems can be addressed with dynamic programming techniques. In practice, 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 of real-world stochastic control problems. Effort continues to focus on the use and analysis of popular temporal difference methods and Q-learning algorithms, in the context of fully and partially observable Markov decision problems, on the simulation-based solution of large-scale least-squares problems, and on a number of issues relating to the central method of approximate policy iteration: convergence, rate of convergence, singularity, and susceptibility to simulation noise of policy evaluation, exploration-enhanced methods, error bounds, policy oscillation, and issues of decision making in an asynchronous multiagent environment. In addition, fundamental issues of dynamic programming were studied in an abstract setting, and a research monograph on the subject by Professor Bertsekas was published in April 2013.

Robert Berwick

Professor Robert Berwick's research during the past year made several advances regarding the biological and computational origins of human language and cognition. Professor Berwick developed new models for early infant learning of words paired with objects that question the current status quo regarding statistically based methods. He advanced a novel method for analyzing unsigned Supreme Court decisions to determine their authorship, and he published a book that gives an overview of the "state of the art" regarding the acquisition of language.

Along with professor Shigeru Miyagawa, Professor Berwick hypothesized that the origin of human language can be accounted for by the combination of two components that evolved independently: the first, analogous to the vocal learning of songbirds; and the second, the ability to name objects. The reflex of these two distinct components can be observed in modern human syntax. Regarding this model, Professor Berwick published a review of the known results on brain, language, and evolution along with professors Noam Chomsky and Johan Bolhuis (Utrecht University), in several highly cited *Nature* journal articles.

In the area of the computational modeling of child language acquisition, Professor Berwick and two visiting scientists and students developed a computer program that can use actual videotaped scene and utterance examples of mothers talking to their children starting from age three months to 24 months, and successfully extract the object-word pairings. The method does not require any computationally intensive Bayesian modeling, but rather is based on a trade-off between memorization and compression according to rules. The model is now being extended to cover the acquisition of syntax. An overview of efforts like these was published during the past year as the book *Rich Languages from Poor Inputs*, Oxford University Press, edited by Professor Berwick.

Along with professor Andrew Lo and several doctoral students, Professor Berwick developed a machine learning approach for the author attribution of unsigned Supreme

Court decisions, which was published in the *Stanford Law Review*. This approach extended classical machine learning methods to a challenging new domain, and the method has proven to be 80% accurate in identifying unsigned decisions. Applied to a case of recent interest, it determined that Chief Justice John Roberts had in fact authored part of the "minority" decision on the new Affordable Care Act, despite his vote with the majority, a matter that had been the subject of great speculation.

Munther Dahleh

On the Value and Price-responsiveness of Ramp-constrained Storage

Professor Munther Dahleh, graduate student Ali Faghih, and principal research scientist Mardavij Roozbehani are working to understand the economic value of storage in the presence of ramp constraints and exogenous electricity prices, and to understand the implications of the associated optimal storage management policy on qualitative and quantitative characteristics of storage response to real-time prices. To study the implications of the optimal policy, they first present computational experiments that suggest that optimal utilization of storage can, in expectation, induce a considerable amount of price elasticity near the average price, but little or no elasticity far from it. They then present a computational framework for understanding the behavior of storage as a function of price and the amount of stored energy, and for characterization of the buy/sell phase transition region in the price-state plane. Finally, they study the impact of market-based operation of storage on the required reserves, and show that the reserves may need to be expanded to accommodate market-based storage. The results are reported in their paper "On the Value and Price-responsiveness of Ramp-constrained Storage," which will appear in *Energy Conversion and Management*, 2013.

Robust Distributed Routing in Dynamical Networks—Locally Responsive Policies and Weak Resilence

Robustness of distributed routing policies is studied for dynamical flow networks, with respect to adversarial disturbances that reduce the link flow capacities. A dynamical flow network is modeled as a system of ordinary differential equations derived from mass conservation laws on a directed acyclic graph with a single origin-destination pair and a constant inflow at the origin. Routing policies regulate the way the inflow at a non-destination node gets split among its outgoing links as a function of the current particle density, while the outflow of a link is modeled to depend on the current particle density on that link through a flow function. These results imply that locality constraints on the information available to the routing policies do not cause loss of weak resilience. Some fundamental properties of dynamical flow networks driven by locally responsive distributed policies are analyzed in detail, including global convergence to a unique limit flow. Results are reported in Professor Dahleh's paper "Robust Distributed Routing in Dynamical Networks—Part I: Locally Responsive Policies and Weak Resilience," with professors Giacomo Como (University of Lund), Ketan Savla (University of Southern California), Daron Acemoglu, and Emilio Frazzoli.

Robust Distributed Routing in Dynamical Networks—Strong Resilience, Equilibrium Selection, and Cascaded Failures

Strong resilience properties of dynamical flow networks are analyzed for distributed routing policies. The strong resilience of the network is defined as the infimum sum of link-wise flow capacity reductions under which the network cannot maintain the asymptotic total inflow to the destination node to be equal to the inflow at the origin. Professor Dahleh and his colleagues propose a simple convex optimization problem to solve for equilibrium limit flows of the unperturbed network that minimize average delay subject to strong resilience guarantees, and discuss the use of tolls to induce such an equilibrium limit flow in transportation networks. Finally, they present illustrative simulations to discuss the connection between cascaded failures and the resilience properties of the network. Results are reported in Professor Dahleh's paper "Robust Distributed Routing in Dynamical Networks—Part II: Strong Resilience, Equilibrium Selection, and Cascaded Failures," with Professors Como, Savla, Acemoglu, and Frazzoli, in *Institute of Electrical and Electronics Engineers (IEEE) Transactions on Automatic Control*, in press, January 2013.

Emilio Frazzoli

Professor Emilio Frazzoli and his research group pursued research in control theory, algorithmic robotics, and applications to autonomous vehicles, robotic networks, and transportation systems.

Sampling-based Algorithms for Anytime Control and Estimation

Building on recent breakthroughs, primarily the work on asymptotically-optimal sampling-based algorithms such as rapidly exploring random tree, Professor Frazzoli and his group continued developing the theory of sampling-based algorithms and their application in planning, control, and estimation problems.

A first research direction is aimed at extending the previous results, which are applicable to path-planning problems and to planning problems for general dynamical systems. For systems that admit an exact steering law, professor Sertac Karaman and Professor Frazzoli discovered the first algorithm that is both asymptotically optimal and computationally efficient. Graduate student Jeong Jeon, postdoctoral associate Raghvendra Cowlagi, *et al.*, showed how to extend the previous algorithm to a realistic model of a car, traveling at high speed on loose surfaces.

Another important direction of research centers on the development of planning algorithms that provably satisfy temporal/logical constraints, such as those arising from the rules of the road and rules of engagement given in high-level, human-like languages.

In collaboration with Professor Karaman and professor Daniela Rus, Professor Frazzoli's group developed an approach to combine state-of-the-art sampling-based algorithms for motion planning with decision-theoretic models and efficient model-checking techniques. The algorithms have been validated experimentally on an autonomous vehicle platform available to the principal investigator (PI) at the Singapore-MIT Alliance on Research and Technology, on the campus of the National University of Singapore.

This work constitutes a significant contribution to the state of the art in that it is, to the best of the authors' knowledge, the first time formal methods have been used to synthesize provably correct behavior for a full-size, "real-world" autonomous vehicle operating on a public road. With this work, Professor Frazzoli's group not only demonstrated the ability to plan and execute motions that provably satisfy a set of given rules of the road, but also developed a systematic way to integrate "hard" rules (e.g., "do not collide with other vehicles or pedestrians") with "soft" rules (e.g., "stay in the rightmost lane if possible"), or more in general a set of prioritized rules.

Other contributions in this area include extension of the basic algorithm (originally developed for motion planning problems) to more general classes of important problems in control theory, ranging from stochastic optimal control to nonlinear filtering with non-Gaussian noise models.

A final set of contributions centers on the development of novel, computationally efficient approaches for implementing "primitive" operations in sampling-based algorithms. For example, collision checking is generally considered to be the primary computational bottleneck in sampling-based motion planning algorithms. Professor Frazzoli and his group showed that this does not have to be the case. More specifically, they introduce a novel way of implementing collision checking in the context of sampling-based motion planning, such that the amortized complexity of collision checking is negligible with respect to that of the other components of sampling-based motion planning algorithms. As a consequence, proximity searches are identified as the main factor determining the complexity of sampling-based motion planning algorithms.

Autonomous Vehicle Development

Professor Frazzoli, primarily through his involvement in SMART, and in collaboration with Professor Rus, led the development of novel, low-cost experimental testbeds for autonomous (self-driving) car technologies. The main platform developed during the past year has been a self-driving golf cart, which has been used as a prototype for automation-enabled mobility-on-demand (MOD) systems.

In terms of basic technology development, a major objective has been to make autonomous vehicles not only safe and reliable, but also affordable. The autonomous vehicle technology Professor Frazzoli and his collaborators have developed is based on much cheaper sensors than other state-of-the-art approaches, while ensuring comparable capabilities. Key ideas rely on environmental feature detections (e.g., curbs); LIDAR-vision sensor fusion; and vehicle-to-vehicle, vehicle-to-infrastructure communications. All technological developments have been tested and verified on the campus of the National University of Singapore.

During the past year, Professor Frazzoli and his collaborators have acquired two additional golf carts and a Mitsubishi iMiEV electric car for further development of the testbed.

Mobility-on-Demand Systems

One of the leading emerging paradigms for future urban mobility systems is one-way vehicle sharing, which effectively merges private and public mobility and directly targets the problems of parking spaces and current low vehicle utilization rates. However, sharing has its drawbacks: when some origins and some destinations are more popular than others, the system will inevitably become out of balance; vehicles will build up at some stations, and become depleted at others.

Professor Frazzoli and his group, in collaboration with Professor Rus and professors Marco Pavone (Stanford University) and Stephen Smith (Waterloo University), developed the first systematic, mathematically rigorous approach to the design of a robotic solution for vehicle rebalancing in MOD (i.e., one-way car-sharing) systems, whereby the shared vehicles autonomously drive from a delivery location to the next pickup (or possibly charging, or parking) location. Rebalancing through autonomously driving vehicles has the clear potential of eliminating imbalances within the transportation network, and effectively adds another dimension to MOD systems by introducing autonomy in the design space.

Specific contributions in this area include the first rigorous bound on the minimum number of vehicles to ensure system stability, both considering automated and human-driven rebalancing, and corresponding stabilizing algorithms. Several studies on general problems of on-demand fleet management (e.g., for taxi fleets) have also been part of this effort.

Stability and Control of Transportation Networks

Using modern technologies for analyzing in real-time massive amounts of traffic data and, correspondingly, control in real-time drivers' behavior (e.g., through traditional means such as traffic light or more innovative concepts such as congestion charging and dynamic pricing), Professor Frazzoli and his group developed a novel distributed algorithm for traffic light control based on the well-known backpressure (or differential backlog) routing algorithm, which ensures maximal throughput for a signalized transportation network. The algorithm is currently under evaluation by the Singapore Land Transport Authority for a possible field trial.

Other work includes the analysis of stability and robustness of equilibria in transportation networks, based on multiple time scales (slow "learning" of network-wide traffic conditions, and fast driver reactions to observed traffic); this work was done in collaboration with Professors Dahleh, Acemoglu, Savla, and Como.

Control Theory and Applications

Other research done in the past year encompasses several topics in the general areas of control theory and applications. Examples include the complete characterization of minimum-time trajectories for vehicles with position-dependent curvature constraints, new algorithms for simultaneous input and state estimation for stochastic systems, distributed event-triggered control for multi-agent systems, and competitive search algorithms.

Jonathan How

Professor Jonathan How leads research efforts focused on the control of multiple autonomous agents, with an emphasis on distributed decision making with uncertainty; path planning, activity, and task assignment; mission planning for unmanned aerial vehicles and unmanned ground vehicles; sensor network design; and robust, adaptive, and nonlinear control. Professor How is also the principle investigator of the Aerospace Control Laboratory. Recent research includes the following:

Real Time Multi-agent Decision Making

Professor How, with student Luke Johnson, developed a distributed task-planning algorithm that provides provably good, conflict-free, approximate solutions for heterogeneous multi-agent/multi-task allocation problems. The consensus-based bundle algorithm (CBBA) had been previously extended to include task time-windows and coupled agent constraints, and to robustly handle parameter uncertainty. Recent work has introduced the hybrid information and planning consensus, an extension of CBBA that incorporates all information available in the network to implicitly coordinate with network neighbors, while still retaining performance and convergence guarantees. This enables much higher levels of coordination and cooperation than were previously available with state-of-the-art algorithms.

Markov Decision Processes (MDPs) are a framework for formulating many decision problems of interest. Professor How and students Trevor Campbell and Luke Johnson have extended CBBA to tasks specified as an MDP while preserving the same guarantees as CBBA, and provide a precise analytic trade-off between solution quality and computation time. Professor How and students Josh Redding and Kemal Ure previously developed Group Aggregated Decentralized Multiagent MDP solving technique for efficient online planning. This work is being continued by students Yufan Chen and Kemal Ure who, with Professor How, have developed hierarchical decomposition techniques, including the Decentralized Incremental Feature Dependency Discovery, which decomposes the learning problem to individual agents and allows agents to share the relevant models/features under communication constraints.

Mission Technologies

Professor How and students Mark Cutler and Bernard Michini developed a lightweight helicopter autopilot (20 grams) that utilizes an infrared camera to track specific markers. The autopilot and camera calculate range and bearing relative to a fixed set of markers, enabling the helicopter to navigate and fly relative to these markers. This fixed set of markers can be statically mounted or mounted on moving vehicles. The capability will allow for close-proximity multi-vehicle indoor flight utilizing only low-cost, low-weight hardware.

Information-gathering Networks

Professor How and student Daniel Levine developed the concept of "focused" information measures to quantify informativeness in the presence of nuisance variables and optimize the observation selection process in a class of multivariate distributions, including Gaussians. By representing the distribution as a Gaussian Markov random field, they have proven that certain measures of informativeness can be efficiently computed in a parallel or distributed manner. These advancements allow for faster identification of information-rich sources in complex domains.

Integrated Learning and Planning

Many planning algorithms require a model of the mission environment, but such models are often approximated and/or wrong. Hence a learning algorithm should be integrated in the planning loop to improve the planner's model over time. Student Kemal Ure, former postdoctoral associate Girish Chowdhary, and Professor How developed a multi-level health-aware planning architecture that supports complex missions; the main highlight of the architecture is the tight coupling between the vehicle-level learning focused adaptive controller and the mission planner, enabling the planner to re-assess mission resources and take preventive actions for future failures.

Postdoctoral associates Alborz Geramifard and Thomas Walsh and Professor How extended the previous work on adaptive representations by providing theoretical work on the convergence of incremental Feature Dependency Discovery (iFDD) algorithm and introducing the $iFDD^{+}$ algorithm as a new batch technique. The new method outperformed the state-of-the-art adaptive representations across three domains, including a network administration task with 20 computers. Furthermore, together with student Robert Klein, Dr. Geramifard released an open-source reinforcement learning platform for education and learning named RLPy, which was featured on *MIT News,* May 29, 2013.

Nonparametric Adaptive Control for Health-aware Operation

Traditional approaches in adaptive control have focused on adaptive elements where the number and parameters of the regression basis functions are fixed a priori. However, it is often difficult to choose the right set of bases without prior knowledge of the operating domain. Professor How's work with Dr. Chowdhary has shown that nonparametric model reference adaptive control (MRAC) approaches can dynamically choose these basis functions. In addition, Professor How, student Robert Grande, and Dr. Chowdhary developed an MRAC framework that uses Bayesian nonparametric elements to learn the model error online, allowing both the basis functions and hyperparameters to be learned online. Work in this area will enable control of nonlinear systems such as airplanes in atypical flight regimes. Professor How, Dr. Chowdhary, and student John Quindlen extended the MRAC framework to situations where states are not fully observed (for instance, in flexible aircraft) by eliminating approximate model inversion assumptions and extending a nonparametric concurrent learning framework.

Value of Information-based Sensing for Decentralized Data Fusion

In many distributed applications, agents have limited on-board resources and their information is not equally valuable. Therefore, it is important to balance resources between exploring and exploiting important information. Student Beipeng Mu, Dr. Chowdhary, and Professor How developed a Value of Information decision system for networked agents with heterogeneous abilities, such as targeting or searching. In the newly developed algorithm, the value of exploration is defined in terms of its potential improvement on mission rewards, coupling exploration directly into mission activities.

Bayesian Nonparametric Models for Planning and Control

Classical approaches for learning models to improve mission planning typically assume the environment has a certain parameterized structure. However, finding the right parameterization a priori can be challenging. Professor How's work on Bayesian Nonparametric Models (BNP) leverages powerful nonparametric learning techniques to build models for control without requiring such background knowledge.

Students Sarah Ferguson and Brandon Luders and Professor How, along with former students Georges Aoude and Sameera Ponda, developed an efficient framework for long timescale prediction in autonomous traversal problems that is suitable for real-time applications, while achieving high accuracy. An offline unsupervised learning algorithm uses unlabeled trajectory data to create mixture models of dynamic obstacle mobility patterns. This algorithm leverages Dirichlet processes to classify trajectories and Gaussian processes to provide probabilistic motion patterns. An online risk-aware motion planner, developed by student Luders and Professor How, ensures safe traversal subject to efficiently computed risk thresholds. Professor How and student Christopher Lowe are currently extending this technology to collision avoidance and path planning problems in three dimensions, with particular application to autonomous aircraft operations in civilian airspace.

Student Rob Grande, visiting student Miao Liu, Dr. Chowdhary, Dr. Walsh, and Professor How developed a BNP reinforcement learning algorithm called GPQ for control situations. Unlike previous reinforcement learning methods for continuous domains, GPQ chooses its own basis points to represent the task value function and also can be used in an off-policy setting where data was collected using a safe strategy. The algorithm has guaranteed convergence, is computationally efficient, and explicitly models uncertainty.

Student Buddy Michini and Professor How have developed a Bayesian Nonparametric Inverse Reinforcement Learning technique (BNIRL) for learning task specifications, in the form of subgoals, based on user demonstrations of the task. Crucially, BNIRL is tractable for large domains through extensions that allow it to scale only with the number of data points collected and the use of Gaussian Processes to model and plan in continuous-state environments. Student Michini and Professor How have demonstrated the practical use of BNIRL on a Quadrotor and a small robotic car.

In multiagent systems, Professor How and student Trevor Campbell, along with external collaborators from Duke University and Ohio State University, have developed the Dynamic Means algorithm, an asymptotic version of the dependent Dirichlet process that provides the flexibility of Bayesian nonparametrics, the long-term stability of the dependent Dirichlet process, and the tractability of classical methods in a single algorithm. Combining this method with iFDD, students Campbell and Klein, Dr. Geramifard, and Professor How developed a simultaneous clustering/representation expansion model learning algorithm.

Patrick Jaillet

Professor Patrick Jaillet is co-director of the Operations Research Center (ORC), and his main research continues to be concentrated on formulating and analyzing online, dynamic, and real-time versions of classical network and combinatorial optimization problems, such as the shortest path problem, the traveling salesman problem, the assignment/matching problem, the secretary problem, as well as some of their generalizations. The research deals with provable results (algorithmic design and analysis) on how to solve such problems under uncertainty, with or without explicit stochastic modeling of the uncertainty. Methodological tools include those from online optimization (competitive analysis), stochastic optimization (robust analysis), online learning (min-max regret analysis and Bayesian updates), and their integrations.

Motivating applications include various routing problems that arise from transportation and logistics networks, data communication and sensor networks, autonomous multi-agent systems, as well as dynamic resource allocation problems in various applications, such as the internet (search engines and online auctions) and healthcare (kidney exchange programs).

Professor Jaillet's research group this past year included postdoctoral associates Vahideh Manshadi (EECS/ORC) and Thibaut Vidal (EECS); ORC graduate students Iain Dunning, Chong Yang Goh, Swati Gupta, Nikita Korolko, Maokai Lin, and Xin Lu; EECS graduate students Dawsen Huang, Andrew Mastin, and Jin Hao Wan; Manufacturing Systems and Technology graduate student Setareh Borjan; and EECS undergraduate student Nicoli Ludvigsen. The research group in Singapore included postdoctoral associates Yossiri Adulyasak and Ali Oran (SMART) and several graduate students from the National University of Singapore, Nanyang Technological University, and Singapore Management University.

Current funded research programs come from the Operations Research program of the National Science Foundation (NSF): Online Optimization for Dynamic Resource Allocation Problems; from the Operations Research division of the Office of Naval Research (ONR): Online and Dynamic Optimization under Uncertainty, and Decentralized Online Optimization in Multi-agent Systems in Dynamic and Uncertain Environments; from the Optimization and Discrete Mathematics program of the Air Force Office of Scientific Research (AFOSR): Data-driven Online and Real-time Combinatorial Optimization; and from SMART: Future Urban Mobility, a large project with nine other MIT PIs.

Sanjoy Mitter

Professor Sanjoy Mitter's research has spanned the broad areas of systems, communication, and control. Although his primary contributions have been in the theoretical foundations of the field, he has also contributed to significant engineering applications, notably in the control of interconnected power systems and pattern recognition. His current research interests are theory of stochastic and adaptive control; mathematical physics and its relation to system theory; image analysis and computer vision; and structure, function, and organization of complex systems. In particular, he has continued to work on systems aspects of power systems (smart grids) during the past year.

Professor Mitter has also continued his long-standing collaboration with Dr. Charles Rockland (RIKEN Brain Science Institute, Tokyo) on issues of autonomy and adaptiveness in neural systems. There is renewed interest in studying the nematode from the viewpoint of understanding the structure-to-function map, a program the two proposed in the 1980s.

Professor Mitter has been investigating topological properties of large data sets using ideas from differential geometry and algebraic topology. New results on manifold learning have been obtained in joint work with postdoctoral associate Hariharan Narayanan and professor Charles Fefferman (Princeton University).

In addition, Professor Mitter has continued his collaboration with Dr. Nigel Newton (University of Essex) on the relation between statistical mechanics, statistical inference, and information theory. In recent work, they have given a proof of the noisy channel coding theorem (including error exponents) from the variational point of view of Bayesian inference. He has also been working with postdoctoral associate Saverio Bolognani on issues of congestion, control, and scheduling in smart grids with shiftable loads.

Professor Mitter has been investigating issues of systemic risk in the financial system in collaboration with Professors Dahleh and Lo. He has actively participated in meetings of the Systemic Risk Consortium, and is a member of the academic advisory board of the Consortium for Systemic Risk Analytics.

Eytan Modiano

Professor Eytan Modiano leads the Communications and Networking Research Group (CNRG), consisting of nine graduate students and three postdoctoral associates. The primary goal of CNRG is the design of architectures for aerospace networks that are cost-effective, scalable, and meet emerging needs for high data-rate and reliable communications. In order to meet emerging critical needs for military communications, space exploration, and internet access for remote and mobile users, future aerospace networks will depend upon satellite, wireless, and optical components. Satellite networks are essential for providing access to remote locations lacking in communications infrastructure, wireless networks are needed for communication between untethered nodes (such as autonomous air vehicles), and optical networks are critical to the network backbone and in high performance local area networks. The group is working on a wide range of projects in the area of data communication and networks with application to satellite, wireless, and optical networks.

During the past year, CNRG continued to work on a new project titled Optimal Control of Wireless Networks: From Theory to Practice, funded by the Office of Naval Research. This project is a collaboration between CNRG and researchers at the Naval Research Laboratory, and the goal is to design practical network control algorithms based on the theories developed by Professor Modiano and his students over the past decade. The algorithms have been shown to optimize network performance, e.g., maximize throughput and utility, but so far have been largely limited to being a theoretical framework. Through this new project, Professor Modiano and his group are developing practical network control algorithms that make efficient use of wireless resources through the joint design of network layer routing, link layer scheduling, and physical layer power and rate control. The collaboration with the Naval Research Lab will facilitate transitioning this promising technology for use in military communication systems. In a related project, funded by NSF, the group is developing distributed network control algorithms for wireless networks. Such distributed algorithms are needed in wireless networks where it is neither practical nor desirable to use centralized control mechanisms.

CNRG also continues to work on an Army Multidisciplinary University Research Initiative (MURI) project titled MAASCOM: Modeling, Analysis, and Algorithms for Stochastic Control of Multiscale Networks. The project deals with control of communication networks at multiple time scales, and is a collaboration among MIT, Ohio State University, the University of Maryland, the University of Illinois, Purdue University, and Cornell University. An important aspect of this project is to understand the impact of traffic correlation at multiple time scales (e.g., heavy-tailed traffic statistics) on the performance of network control algorithms. In particular, Professor Modiano and his group have shown that classical network control algorithms, which are widely used in both wireless and wired networks, perform very poorly in the presence of highly correlated traffic, and have developed network control mechanisms for alleviating the effects of traffic correlation.

The group continues to work on a Department of Defense–funded project toward the design of highly robust telecommunication networks that can survive a massive disruption that may result from natural disasters or intentional attack. The project examines the impact of large-scale, geographically correlated failures on network survivability and design. The impact of geographically correlated failures is an important aspect of robust network design that has received very little attention in the past. Professor Modiano and his team developed mechanisms for assessing the vulnerability of networks to geographical failures. These mechanisms can be used to identify the most vulnerable parts of the network, and give insights to the design of network architectures that are robust to natural disasters or physical attacks.

Last year, the group started working on a project with the Masdar Institute of Science and Technology dealing with smart grids and the design of communication networks that would enable future smart grid applications. Future power grids will require rapid and dynamic control of the grid, in order to respond to changes in demand, supply, or failures. Such dynamic control requires the ability to monitor the power grid in real-time and rapidly respond to changes by appropriate control actions, and inevitably depends on the availability of a reliable communication infrastructure. The goal of this new project is to develop the communication network's architecture for enabling effective control of future smart grids.

Over the past year, the group started to work on a related project dealing with the robustness of the power grid to physical attacks. Much like communication networks, the power grid is vulnerable to failures of (power) transmission lines. However, the flow of power on transmission lines obeys the laws of physics, and the failure of one power line can lead other lines to overheat and fail in a cascading manner. These cascading effects make power networks very vulnerable to disasters or attacks. The goal of this project is to understand the vulnerability of the power network to geographically correlated failures that may result from a natural disaster or physical attack, and to develop methods to make the power grid robust to such failures. Moreover, an interesting interplay arises between the power grid and other networked infrastructures that depend on the grid for power (e.g., telecommunication networks), whereby failures in the power grid can lead to failures in the communication networks due to the loss of power. Thus, another goal of this project is to understand the interplay between such interdependent networked infrastructures and to develop mechanisms to mitigate the effects of such failure cascades.

CNRG's research crosses disciplinary boundaries by combining techniques from network optimization, queueing theory, graph theory, network protocols and algorithms, hardware design, and physical layer communications.

Asuman Ozdaglar

Professor Asuman Ozdaglar's research group focuses on modeling, analysis, and optimization of large-scale dynamic multi-agent networked systems. The research draws on advances in game theory, optimization theory, dynamical systems, and stochastic network analysis. It focuses on both investigating substantive issues in these areas and developing new mathematical tools and algorithms for the analysis of these systems.

A major current research area in Professor Ozdaglar's group is socio-technological networks, which include communication, decision, and power networks. The operation and failures of such networks depend not only on their engineering design but also on the human element—the intentions, information, and strategic behavior of the users and participants. Professor Ozdaglar's work has integrated the analysis of social and economic interactions into the study of networks.

One strand of her work focuses on developing models and control mechanisms for contagious processes (e.g., opinion dynamics and epidemics) over complex networks. In a series of papers, she developed new models of information exchange and opinion dynamics in the presence of prominent agents (which may represent community leaders, media outlets, or competing firms that target individuals embedded in a social network). In two recent papers, she addressed the question of influencing the opinion dynamics by targeting central agents (with and without adversarial influence) and presented computationally efficient algorithms for locating such central agents. Recent work, joint with professor John Tsitsiklis and graduate student Kimon Drakopoulos, studies the design of dynamic intervention mechanisms to influence opinion or epidemic dynamics. In contrast with the literature in this area that focuses on static (ex-ante) targeting of agents, this work allows for real-time feedback about network state. This research effort is supported by an Army Research Office (ARO) project, an AFOSR project, and the Draper Directed Research and Development Program.

Another strand of Professor Ozdaglar's work shows how interactions over economic and financial networks can lead to the amplification and propagation of small shocks into systemic cascades. A recent project, with professor Daron Acemoglu (Economics) and professor Alireza Tahbaz-Salehi (Columbia University), investigates the cascade effects that arise in a production economy with an input-output structure, whereby shocks to some sectors spread to their downstream sectors and beyond. Results show that the traditional economic insight that neglects firm level variations in aggregate economic fluctuations (using law of large numbers type arguments) fails in interconnected systems, and provides a general framework for the analysis of the relationship between the network structure of an economy and its aggregate volatility. In another joint paper, she focuses on interlinkages created by financial transactions (counterparty relations). This paper shows that contagion in financial networks exhibits a form of phase transition as interbank connections increase. In particular, as long as the magnitude and the number of negative shocks affecting financial institutions are sufficiently small, more "complete" interbank claims enhance the stability of the system. However, beyond a certain point, such interconnections start to serve as a mechanism for propagation of shocks and hence lead to a more fragile financial system. This research effort is supported by an ARO MURI.

Professor Ozdaglar's recent work investigates the implications of strategic, selfish agents on system-wide performance on a variety of problems. In a project with Professor Acemoglu and joint postdoctoral associate Azarakhsh Malekian, she developed a model of investments in security in a network of interconnected agents subject to cascading failures due to an exogenous or endogenous attack (which depends on the profile of security investments by the agents). This paper provides a systematic analysis of the equilibrium and optimal security investments in general random networks subject to an attack, and delineates conditions on the network structure and the attack model under which underinvestment or overinvestment incentives will dominate. In other projects, she develops game-theoretic models for resource allocation problems in communication networks, with a focus on pricing and investment incentives of providers and implications of competition on network performance. In a recent paper, joint with student Paul Njoroge and with professors Nicolas Stier-Moses and Gabriel Weintraub (both of Columbia University), she presents a game theoretic model based on a two-sided market framework to investigate net neutrality issues from a pricing perspective.

Professor Ozdaglar also investigates mechanism and market design focusing on provision of incentives to selfish agents to align their performance with a system-wide objective. A recent project, joint with professor Pablo Parrilo and graduate student Ozan Candogan, focuses on iterative auctions for selling multiple items. Iterative multi-item auctions are employed widely in practice in a number of applications, including spectrum and procurement auctions. Despite their ubiquity, however, the auctions used in practice usually lack optimality guarantees. This paper focuses on valuation functions that admit a compact graphical representation, which is a natural assumption that holds in practical settings and provides efficient iterative auctions with simple pricing structures. In another project, joint with Professor Acemoglu and postdoctoral associate Mohamed Mostagir, she studies crowdsourcing where the goal is to organize a crowd of individuals for achieving tasks that a single human being and computer cannot perform alone. This work provides a mathematical framework to analyze the problem of assigning and pricing innovation tasks in a crowdsourced environment. Another recent work, joint with graduate student Ermin Wei and Dr. Malekian, focuses on the provision of dynamic incentives, while limiting the extent of fluctuations, in power networks consisting of heterogeneous strategic users. This work is supported by an NSF Cyber-Physical Systems grant, joint with the University of California, Berkeley; the University of Michigan; and Vanderbilt University.

Professor Ozdaglar's group also works on developing novel distributed methods that enable collection, storage, and processing of data using large number of agents connected through a network. Many of these problems can be formulated as a (coupled) convex optimization problem, which needs to be solved by the agents using local information and communication through the network. The standard approach for designing distributed algorithms for such problems uses a (sub)-gradient based method, which suffers from slow rate of convergence: the best known rate of convergence for subgradient based distributed methods is $O(1/\sqrt{k})$ (where k is the iteration number). Her recent work introduced new distributed optimization algorithms with much faster rate of convergence. In a project joint with graduate student Wei and professor Ali Jadbabaie (University of Pennsylvania), she develops novel Newton-type second order methods to solve network utility maximization problems in a distributed manner, which are significantly faster than the standard first-order (or gradient) approaches. In another paper, joint with graduate student Wei, she develops an algorithm based on alternating direction method of multipliers (ADMM), which is a classical method for sequentially decomposing optimization problems with coupled constraints. This paper provides a completely asynchronous and distributed ADMM algorithm for solving coupled convex optimization problems over networks and shows that this algorithm converges at the rate $O(1/k)$. Another recent work focuses on a structured version of this problem, where the local objective functions take an additive form with a differentiable and a nondifferentiable component, and develops a distributed proximal gradient method that solves this problem at rate $O(1/k)$ over a network with time-varying connectivity. This research effort is supported by an AFOSR MURI grant, joint with the Georgia Institute of Technology and the University of Maryland, and an ONR project, joint with the University of Pennsylvania.

Pablo Parrilo

Professor Pablo Parrilo (associate director of LIDS) and his research group are focused on mathematical optimization, systems theory, and control, with emphasis on development and application of computational tools based on convex optimization and algorithmic algebra.

The book *Semidefinite Optimization and Convex Algebraic Geometry*, edited by Grigoriy Blekherman (Georgia Institute of Technology), Professor Parrilo, and Rekha Thomas (University of Washington), was published by the Society for Industrial and Applied Mathematics (SIAM) in the Mathematical Programming Society–SIAM Series on Optimization, in January 2013. Convex algebraic geometry is an important and applicable research area, developed by Professor Parrilo and collaborators under an NSF Focused Research Group grant. The book features contributions from several areas of convex and algebraic geometry, semidefinite programming, and optimization. It is a self-contained introduction to the topic suitable for readers at the graduate level and can be used for a class or seminar.

In collaboration with PhD student Hamza Fawzi, Professor Parrilo developed a framework to prove lower bounds on the nonnegative rank. The nonnegative rank of an entrywise-nonnegative matrix A of size mxn is the smallest integer r such that A can be written as $A=UV$ where U is mxr and V is rxn and U and V are both nonnegative. The nonnegative rank arises in different areas such as combinatorial optimization and communication complexity. Computing this quantity is NP-hard in general and it is thus important to find efficient bounding techniques. New techniques were proposed for lower-bounding the nonnegative rank, which, unlike existing lower bounds, do not explicitly rely on the matrix sparsity pattern and applies to nonnegative matrices with arbitrary support. This lower bound is expressed as the solution of a copositive programming problem and can be relaxed to obtain polynomial-time computable lower bounds using semidefinite programming. The idea involves computing a certain nuclear norm with nonnegativity constraints, which allows to lower bound the nonnegative rank in the same way the standard nuclear norm gives lower bounds on the standard rank.

With James Saunderson (PhD student, joint with professor Alan Willsky), Professor Parrilo and his group have developed new polynomial-sized semidefinite representations of derivative relaxations of spectrahedral cones. These convex cones form a family of non-polyhedral outer approximations of the nonnegative orthant that preserve low-dimensional facial structure while successively discarding high-dimensional facial structure. These representations allow the use of semidefinite programming to efficiently solve the linear cone programs associated with these convex cones as well as their (less well understood) dual cones.

With Frank Permenter (PhD student, joint with professor Russell Tedrake), Professor Parrilo developed a new method for choosing "good" monomial basis for sums of squares programs formulated over a quotient ring. They showed that in many cases it is possible to use a reduced subset of standard monomials by combining Groebner basis techniques with the well-known Newton polytope reduction. This reduced subset of standard monomials yields a smaller semidefinite program for obtaining a certificate of non-negativity of a polynomial on an algebraic variety.

With PhD student Ozan Candogan (jointly advised with Professor Ozdaglar), Professor Parrilo has developed novel multi-item iterative auctions. These auctions are a class of mechanisms that are commonly employed in practice, for instance, in the context of spectrum and procurement auctions. In this work, they developed practical and efficient iterative auctions for multi-item settings that exhibit both value complementarity and substitutability. These auctions feature a natural class of valuation functions that admit a compact representation, referred to as graphical valuations. For special classes of graphical valuations, such as tree valuations, the efficient allocation can be computed via linear programming, and as a consequence their auctions implement the efficient outcome using item prices, i.e., offering an anonymous price for each item. However, they establish that this simple pricing structure is not sufficient when the underlying value graph has cycles.

Professor Parrilo and colleagues from the University of Washington are currently investigating the basic geometric questions of when a given convex set is the image under a linear map of an affine slice of a given closed convex cone. Such representations, or "lifts," of convex sets are especially useful if the cone admits an efficient algorithm for linear optimization over its affine slices, and are fundamental in the application of conic optimization to a large class of problems. They have shown that the existence of a lift of a convex set to a cone is equivalent to the existence of a factorization of an operator associated to the set and its polar via elements in the cone and its dual. This generalizes a well-known theorem of professor Mihalis Yannakakis (Columbia University) that established a connection between polyhedral lifts of a polytope and nonnegative factorizations of its slack matrix. In recent work, they have extended their work from exact factorization to approximation factorizations, showing that it is possible to construct approximate second-order cone lifts of convex bodies from a nearby conic factorization of the slack operator.

Yury Polyanskiy

Professor Yury Polyanskiy conducts research in the areas of mathematics of information (information theory), coding theory, and theory of random processes. His current work focuses on non-asymptotic characterization of the performance limits of communication systems, non-Shannon information measures, redundant circuits, and probabilistic methods in combinatorics.

Empirical Distribution of Good Channel Codes

This line of work studies several properties of channel codes that approach the fundamental limits of a given memoryless channel with a non-vanishing probability of error. The output distribution induced by a capacity-achieving code is shown to be close, in a strong sense, to the capacity-achieving output distribution for discrete memoryless channels (DMC) and additive white Gaussian noise (AWGN). Relying on the concentration of measure (isoperimetry) property enjoyed by the latter, it is shown that regular (Lipschitz) functions of channel outputs can be estimated precisely and turn out to be essentially non-random and independent of the used code. It is also shown that the binary hypothesis testing between the output distribution of a good code and the capacity-achieving one cannot be performed with exponential reliability. Using related methods, it is shown that quadratic forms and sums of q-th powers when evaluated at codewords of good AWGN codes approach the values obtained from a randomly generated Gaussian codeword. The random process produced at the output of the channel is shown to satisfy the asymptotic equipartition property (for DMC and AWGN). Results were reported in Professor Polyanskiy's papers "Lp-norms of Codewords from Capacity and Dispersion-achieving Gaussian Codes," presented at the 50th Allerton Conference on Communications and Control, University of Illinois, October 2012; and "Empirical Distribution of Good Channel Codes with Non-vanishing Error Probability," with Sergio Verdu (Princeton University), submitted for publication to *IEEE Transactions on Information Theory*.

Fundamental Limits of Communication on Multi-antenna Wireless Channels

This line of work investigates the maximal achievable rate for a given blocklength and error probability over block- and quasi-static multiple-input multiple-output fading channels. Under mild conditions on the channel gains, it is shown that the channel dispersion is zero regardless of whether the fading realizations are available at the transmitter and/or the receiver. The result follows from computationally and analytically tractable converse and achievability bounds. Through numerical evaluation, it is verified that zero dispersion indeed entails fast convergence to outage capacity as the blocklength increases. In the example of a particular 1x2 single-input multiple-output Rician channel, the blocklength required to achieve 90% of capacity is about an order of magnitude smaller compared with the blocklength required for an AWGN channel with the same capacity. Results were reported in Professor Polyanskiy's paper "Diversity Versus Channel Knowledge at Finite Blocklength," with Wei Yang (Chalmers University of Technology), Guiseppe Durisi (Chalmers University of Technology), and Tobias Koch (Universidad Carlos III de Madrid), presented at the Information Theory Workshop 2012, Lausanne, Switzerland.

Combinatorial (Adversarial) Joint Source-channel Codes

Algorithms for lossy/lossless compression and error-correcting codes have been at the core of the digital revolution. This line of work focuses on the particular set of applications in which both the lossy compression and the noise resilience are required. Examples include storage of the high resolution imagery on non-perfect flash memory and realtime video surveillance over jammed or noisy channel.

The state-of-the art solution is "separation": pair an off-the-shelf compressor with an off-the-shelf code. However, as was shown by the PI, for worst-case guarantees the separated solution is far from being (even asymptotically) optimal. This realization is the main motivation for the omnidirectional investigation of the combinatorial, geometric, algebraic, and information theoretic aspects of the joint source-channel coding problem. Some of the initial results are reported in Professor Polyanskiy's papers "The Adversarial Joint Source-channel Problem," with Yuval Kochman (Hebrew University of Jerusalem) and Arya Mazumdar (University of Minnesota), presented at the Information Theory Proceedings (ISIT) 2012, Cambridge, MA; and "Results on the Combinatorial Joint Source-channel Coding," presented at 2012 Information Theory Workshop (Lausanne, Switzerland.

Hypercontractivity and Analysis of Boolean Functions

Consider a linear space of functions on the binary hypercube and a linear operator T_δ acting by averaging a function over a Hamming sphere of radius δn. It is shown that such an operator has a dimension independent bound on the norm $L_p \to L_2$ with $p = 1+(1-2\delta)^2$. This result evidently parallels a classical estimate of Bonami and Gross for $L_p \to L_q$ norms for the operator of convolution with a Bernoulli noise. There are significant difficulties in proving the estimate for T_δ, since the latter is neither a part of a semigroup nor a tensor power. The result is shown by a detailed study of the eigenvalues of T and $L_p \to L_2$ norms of the Fourier multiplier operators Π_a with symbol equal to a characteristic function of the Hamming sphere of radius a.

An application of the result to additive combinatorics is given: any set $A \subset \boldsymbol{F}_2^n$ with the property that $A+A$ contains a large portion of some Hamming sphere (counted with multiplicity normalized by cardinality of A) must have cardinality a constant multiple of 2^n. The interest of the PI in these questions arose from the earlier results on hypothesis testing with a restricted kind of tester. Results were reported in Professor Polyanskiy's papers "Hypothesis Testing Via a Comparator," presented at ISIT 2012, in Cambridge, MA; "Hypothesis Testing via a Comparator and Hypercontractivity," submitted in February 2013 to *Problemy Peredachi Informatsii*; and "Hypercontractivity of Spherical Averages in Hamming Space," preprint, June 2013.

Mardavij Roozbehani

Principal Research Scientist Mardavij Roozbehani led several research efforts related to understanding the sources of robustness, fragility and systemic risk in power systems, as well as robustness and safety properties of software systems. These interdisciplinary research efforts are at the interface of control theory, computer science, power systems, and optimization and operations research. Funding for his research comes from NSF, Siemens AG, and Draper Laboratory.

Systemic Risk in Power Systems

The statistics of outages in the US electric power grid show that both frequency and size of the power outages in the US have been steadily on the rise. Careful studies of the anatomy of these outages show that large failures in power systems are almost never caused by a single large disturbance that brings down an entire system. Usually, such large failures are the result of an increased level of risk or reduced level of robustness, which makes the system more vulnerable to a sequence of contingencies, or certain patterns of disturbances. The contingencies can be the result of a variety of events, such as volatile weather events, local component failures, intermittencies in renewable generation, etc. But the cascades that follow such events are often reinforced by mechanisms that are put in place to improve efficiency under normal conditions, or increase robustness to withstand other types of disturbance. For example, the tripping of protective relays isolates faults but also creates a buildup of load on the rest of the system, sometimes in a way that can lead to a cascade of other faults. Therefore, while such systems may perform well under normal operating conditions or be very robust to most common disturbance scenarios, they can exhibit fragility in response to certain large or small disturbances. Such fragility in power systems and more broadly in interconnected systems is characterized by systemic risk.

At an abstract level, fragility of a feedback control system can be explained by some of the results in the classical control literature, which establish trade-offs between performance and robustness. In view of these tradeoffs, fragility can be interpreted as follows: increasing robustness with respect to certain types of disturbances inherently increases sensitivity with respect to uncertainties and disturbances that are not included in the model. Such concepts are, however, less developed and less formalized for interconnected systems. Dr. Roozbehani's research is motivated by both the theory gap for characterization of such trade-offs in networked systems, and by the need to address a pressing and increasingly more important problem: systemic risk and cascaded failures in power systems.

Part of his research in this area has focused on analysis of the feedback loops and the interactions between the market layer (where pricing and economic decisions take place) and the physical layer (where power flows from generators to loads) of power systems.

In earlier research, his research team showed that volatility increases in the system when price-taking consumers actively respond to wholesale electricity market prices due to the uncertainty in their response to price signals, and uneven sensitivity to prices induced by threshold policies.

Recent work has revealed more subtle sources of fragility in power systems. Researchers developed abstract models that show how fragility and endogenous risk can be inherent to the architecture of the system, and arise from the dynamics of the system even under the most ideal assumptions of fully rational agents with perfect information. They developed a model of a dynamic oligopolistic energy market in which a set of distributed agents with market power dynamically update their output (consumption or production) decisions. In this model, the agents have complete knowledge of how their decisions affect the market price, and are fully rational in strategizing their decisions to minimize their expected cost. By characterizing the statistics of the stationary aggregate output process across a spectrum of networks from fully cooperative to fully non-cooperative, it was shown that a trade-off exists between efficiency (aggregate system cost) and risk (tail probability of aggregate output). While the non-cooperative network leads to an efficiency loss—widely known as the "price of anarchy"—the stationary distribution of the corresponding aggregate output process has a smaller tail, whereas the cooperative network achieves higher efficiency at the cost of a higher probability of output spikes. Furthermore, the cooperative network has a smaller output variance, which can be interpreted as higher robustness; but it also has a higher probability of large output spikes, which can be interpreted as higher fragility. This particular fragility emerges when a large accumulated backlog in the system coincides with lack of flexibility to either absorb the backlog or schedule it for the future. Intuitively, the cooperative scheme allows for shifting flexible loads more aggressively, thereby increasing the probability of a large backlog in the system, which eventually leads to a large spike. The non-cooperative scheme, however, is more conservative in shifting flexible loads, which, as a result, is less likely to lead to a large backlog.

This work has significant implications on the market operations and pricing in power systems of the future. In particular, it shows that the current practice of marginal cost pricing is inadequate for engaging consumers and active demand response in real-time operation of power grids and balancing supply and demand. In order to implement real-time pricing, the system operators need to consider the "risk" of aggregate demand spikes and design dynamic pricing mechanisms that incorporate the backlogged demand as part of the state information. The design of such mechanisms and their implication on efficiency, robustness, and risk, and the type of information required from consumers in order to implement these mechanisms is an important part of the ongoing research.

Mathematical Models of Consumer Agents in Smart Grids

In collaboration with Draper Laboratory, Dr. Roozbehani's research team investigated and characterized the behavior of individual agents in real-time energy markets. These characterizations are extremely detailed, taking into account various practical constraints in order to obtain a realistic description of the behavior of a single agent consumer. Using these characterizations of individual consumers, researchers obtained powerful synthetic models for characterization of the behavior of a representative consumer consisting of an aggregate of heterogeneous consumers. These models map aggregate quantities, such as price and aggregate demand data, to aggregate future consumption as a function of price. These models were derived in an "open loop" setting with exogenous prices—in other words, by assuming that the price signal is an external input of the system. This assumption can be considered a valid approximation as long as the number of consumers responding to the price signal is a relatively small fraction of all consumers' population. However, once the number of consumers responding to dynamic prices reaches a critical mass, they cannot be considered price-takers anymore, since their behavior will affect the clearing price. In future work, researchers will move forward to developing such models and identification of the parameters of the aggregate consumer model directly from data. This identification procedure is fundamental for the actual implementation of any price mechanism. The ultimate goal of this project is the analysis of how the dynamics of the price and the behavior of consumers influence each other in a real-time market. Researchers are working on the analysis of fundamental properties of the system, such as robustness, efficiency, and stability, and how trade-offs between these characteristics can be obtained, and the design of proper pricing and control mechanisms for striking the trade-offs optimally.

Layered Architecture and Distributed Load Scheduling

The design and operation of power systems consists of solving various optimization problems at different time scales and by various agents at different levels of a hierarchy, with different objectives and with access to different type of information. A layered distributed approach consists in assigning the decision variables to different decision makers, embedded in a communication architecture that allows them to share a limited amount of information. The decomposition of the original problem into sub-problems assigned to different distributed decision makers with partial information yields suboptimal operation of the power system. Proper decomposition choices exploit time scale separation, weak coupling of some phenomena, and redundancy of the devices, in order to achieve a certain level of performance with a limited exchange of information between the decision makers. An intermediate layered architecture is expected to provide a trade-off between the performance of the exact solution and the robustness and scalability of the fully decoupled solution. Researchers are working on layered architectures that allow distributed scheduling of shiftable loads in power systems, while respecting global constraints such as voltage and line limit constraints that are dictated by global constraints including power flow equations. The idea is to have a scheduling layer at the level of individual loads make distributed decisions about deferring consumption or storing energy while a congestion control layer is in charge of satisfying the network constraints. In order to achieve an acceptable level of performance and robustness, the two layers must communicate and exchange the relevant information and implement properly designed protocols based on the information. Design of the proper layering, information architecture, and the associated algorithms and protocols is an important part of the ongoing work, while some preliminary results based on max-weight optimization have been obtained.

Compiler Design for Automated Verification of Robustness and Safety of Software Systems

This is a project under the EECS Super Undergraduate Research Opportunities Program initiative, and its goal is to automate the process of software verification (i.e., proving that a given computer program does not produce run-time errors) for a limited class of software. The theoretical framework for finding certificates of safety for a high-level model of a program (e.g., a graph model) based on convex optimization of Lyapunov Invariants was developed in the past several years. Lyapunov Invariants are "energy-like" functions that decay along the execution trace of the program. The decay of the "energy" then provides a certificate for "good behavior" of the program. The verification technique relies on finding Lyapunov Invariants (i.e., certificates of safety) using convex optimization techniques. This project focused on building a compiler that takes a structured class of software (e.g., a limited subset of c codes) as input, and produces a high-level model (e.g., a graph model) that can be analyzed in an optimization-based framework. A compiler with very limited parsing capabilities was developed under this project, and in the next year researchers will be working on adding additional features to both the compiler and the analysis techniques to expand the class of programs and properties that can be verified.

Devavrat Shah

Professor Devavrat Shah and his research group are involved in developing understanding of complex networks so as to be able to better engineer them. This includes communication networks such as the internet and networks of statistical dependencies observed in large datasets captured through graphical models and emerging social networks like Twitter. A salient feature of Professor Shah's work across this variety of networks is the use of distributed, iterative, or so-called message-passing procedures as operational building blocks in communication networks (e.g., medium access) for efficient information processing in statistical networks (e.g., belief propagation), and as behavioral models in a social network (e.g., gossip algorithms). These collective research activities span computer science, electrical engineering, operations research, and management sciences, and utilize tools from algorithms, graphs, information theory, Markov chains, optimization, stochastic processes, and queuing theory.

Processing Social Data: Recommendations Using Comparisons

A recommendations system primarily determines what a user, community, or entire population may like/dislike given its prior history. A recommendation system has become essential for any "consumer" facing system—it is the primary force behind Amazon becoming a successful e-retailer or Netflix becoming a dominant media portal.

An important challenge in designing a good recommendation system lies in its ability to quickly identify an individual's preference from heterogeneous sources of partial preference data. Professor Shah and his group (Ammar Ammar, PhD candidate; Sewoong Oh, former postdoctoral researcher, now assistant professor at the University of Illinois at Urbana-Champaign; and Sahand Negahban, former postdoctoral researcher, now assistant professor at Yale University) have developed a novel approach by using comparison as the basic unit of preference data and thus capture signal about preferences from all sorts of preference data simultaneously.

This work has managed to successfully impact the industry (Netflix is interested in trying out aspects of the technology developed) and has received popular press coverage, including in the *New York Times*, in addition to academic credentials (publications at the major venues).

This work has been supported through projects funded by NSF.

Processing Social Data: Predicting Trends on Twitter

Twitter captures the "pulse of society" through tweets made by people. In particular, the "trends" on Twitter reflect what is of interest to the dominant segment of society. Predicting trends, that is determining what "topic" might become "trending" before it truly does so, is of great interest—to social advertisers for being able to capture the trending moment to reach massive target audiences, for social resource planners to be able to plan better, or for governments to react quickly to impending events. Intellectually, the question of trend prediction on Twitter has become a major challenge for the entire social media and information processing community.

Professor Shah's group (Stanislav Nikolov, ME alumnus) and George Chen, PhD candidate, EECS) have successfully solved this challenge by developing a novel method for time-series classification. The prediction algorithm manages to predict correctly with 95% true positive rate and 4% false positive rate, and on average one hour and 46 minutes in advance. The work was featured on *Forbes News*, *reddit*, *Gigaom*, and other popular media.

This work has been supported through projects funded by NSF, ARO MURI, and AFOSR.

Efficient Reliable Transmission Over Un-reliable Medium: Rateless Spinal Codes

Professor Shah and his collaborators (PhD student Jonathan Perry and professor Hari Balakrishnan, EECS) have developed a novel rateless code called the Spinal Code. Methodically, this work provides an implementable approach to make the conceptual contribution of Shannon practical. The initial hardware prototype developed for this code shows promising performance—more than 10% improvement over the state-of-art approach across the channel parameters.

This work has been supported through projects funded by NSF and AFOSR.

John Tsitsiklis

Professor John Tsitsiklis (associate director of LIDS) and his research group—Systems, Networks, and Decisions Group—work on system modeling, analysis, optimization, and control in possibly stochastic and dynamic environments and in the possible presence of multiple agents with conflicting interests. Research activities have focused on developing methodologies, mathematical tools of broad applicability, and computational methods. Motivating applications for recent work have come from domains such as computer networks, social networks, and the life sciences.

Computer Networks and the Value of Flexibility

This research concerns the analysis and optimization of resource allocation methods (routing and scheduling) in queuing networks, with communication networks being the prime motivation. Recent work has focused on the subject of flexibility and resource pooling. While the advantages of having flexible servers (each server able to process all job types) are well known, the group's research has established that providing even a small amount of flexibility (each server is able to process a small number of different job types), when suitably architected and exploited, can recover the benefits of full flexibility.

In other research, the group has continued the development of a performance analysis methodology for network scheduling strategies in the presence of highly variable ("heavy-tailed") traffic.

Social Networks

The group is investigating a number of models related to inference and decision making in social networks. Examples include a game theoretic analysis of "global games" (a model of phenomena such as regime change, monetary attacks, etc.), methods for optimal deployment of resources to control a spreading epidemic, and inference algorithms for detecting who influences whom on the basis of observed actions.

Life Sciences

Besides ongoing research on adaptive radiation therapy, the group has been involved in the development of algorithms for the Bayesian inference of the structure of biochemical networks based on limited data, and in the problem of constructing a parsimonious model of sleep apnea with strong explanatory power.

Alan Willsky

Professor Alan Willsky (director of LIDS) leads the Stochastic Systems Group (SSG), whose research focuses on developing statistically based algorithms and methodologies for complex problems of information extraction and analysis from signals, images, and other sources of data. The work extends from basic mathematical theory to specific areas of application. Funding for this research comes from sources including AFOSR, ARO, ONR, and the Royal Dutch Shell Group.

A major thrust of research in SSG continues to be the extension and exploitation of its methodology for statistical inference, information fusion, and estimation for problems involving complex graphical models, such as those that arise in military command and control, mapping from remote-sensing data, and monitoring complex systems. SSG's work involves discovering and exploiting structure in complex graphical representations that lead to new processing algorithms. This research continues to yield significant advances, including new classes of signal and image-processing algorithms that have provable performance properties, that can be applied to very large problems in a scalable manner, and that outperform previous methods. During this past year, the group has had significant success in developing new and scalable algorithms for sampling from very large random field models, including detailed analysis of the performance of highly parallel algorithms. In addition, it has developed a fully distributed approach to message-passing algorithms for inference in graphical models that builds on the methods developed in the preceding year for highly accurate (but centralized) inference, exploiting a combination of both graph and statistical structure. These methods are being applied to problems involving geophysical data analysis. This part of SSG's research portfolio has received considerable international attention, as evidenced by a string of best paper awards, recent plenary and endowed lectures by Professor Willsky, as well as extensive citations and influence on the work of others in fields ranging from systems and control to chemical engineering to groundwater hydrology.

The other central component of research in SSG is in the area of machine learning, in particular the extraction or discovery of statistical models, usually in graphical form, of complex phenomena. Together with work on scalable inference algorithms, this work is of particular current relevance as the challenges of "Big Data" represent some of the most important problems in information technology. This research arcs across the EECS boundary, bringing together new perspectives on these challenges. Recent work in this area has had three components. The first is the continuation of SSG's work on bringing advanced methods of optimization—and in particular semi-definite relaxations—to the learning of models with sparse and exploitable structure. A second involves currently emerging methods for learning models with graphical structure that is ideally adapted to the methods for inference mentioned in the previous paragraph and that also yield models with structure found in many naturally occurring phenomena. These methods can also discover "hidden" variables explaining the statistical structure of the observed variables, a problem of central importance in fields such as social network analysis. In addition, SSG continues work on so-called Bayesian Nonparametric Methods for discovering complex behaviors in temporal or spatio-temporal data. Its work in this area continues to attract attention from researchers in different areas, including the analysis of power utilization in homes; new approaches to long-standing and well-studied problems in speech analysis, such as extracting characterizations of phonemes in speech; and most recently to collaborative work with researchers at Harvard Medical School who are studying motion behavior of mice.

Moe Win

The Wireless Communication and Network Sciences Laboratory led by Professor Moe Win is involved in multidisciplinary research that encompasses developing fundamental theories, designing algorithms, and conducting experiments for a broad range of real-world problems. Current research topics include network localization and navigation, network interference exploitation, intrinsic wireless secrecy, adaptive diversity techniques, and ultra-wide bandwidth systems. Details of a few specific projects are given below.

The group members have continued their research on network localization and navigation. In particular, they developed efficient techniques for sensory data fusion in harsh propagation environments and resource allocation for network localization.

First, they presented a new methodology for nonlinear/non-Gaussian filtering called belief condensation filtering (BCF). By fusing global positioning system (GPS) and inertial data, the simulation results for the navigation problem by fusing GPS and inertial data have shown that BCF can provide accuracies approaching the theoretical bounds and significantly outperform existing techniques in terms of the accuracy versus complexity trade-offs. Moreover, they introduced the concept of ranging likelihood (RL), which captures the positional information inherent in the range measurements for localization, and developed effective techniques for obtaining real-time RL based on unsupervised machine learning. The results using data from an extensive indoor measurement campaign have shown that the proposed approach significantly improves the performance of wireless localization in harsh environments. Second, in network localization, resource allocation not only affects the network lifetime and throughput but also determines the localization accuracy. To obtain the optimal strategy for power allocation, they developed both optimization and geometric methods that exploit the special properties of the performance metric in localization. They also determined the sparsity property of optimal power allocation, showing that at most three anchors need to be activated for the optimal localization accuracy. Currently, the group is working on the design and operation of localization networks with proposed resource allocation strategies.

Professor Win's students are also investigating various topics in intrinsic network secrecy. First, they studied the network performance in terms of network secrecy throughput accounting for intrinsic properties of the network physical layer, such as spatial distribution of the nodes, radio propagation conditions, and aggregate network interference. The results have shown that these intrinsic properties can be used to enhance the network secrecy performance. Based on these results, the students have developed strategies to engineer the interference to further improve the wireless network secrecy. In particular, they have established a mathematical framework to characterize the feasible region of interference engineering strategies, designed algorithms that converge to global optimal, and quantified the contribution of these algorithms in network secrecy enhancement. Second, they considered secret key generation using common unknown deterministic sources and developed a mathematical framework from a non-Bayesian perspective. They proposed a new information measure called intrinsic information to characterize the achievable length of the secret key that can be generated from an unknown deterministic source. As a case study, they investigated a wideband propagation medium in mobile wireless networks as a source and derived its intrinsic information as a function of various network parameters. The main contribution includes a new methodology for information-theoretic secrecy problems as well as insights into the potential of wireless medium for secret key generation. Finally, Professor Win's group recently started a new research project aimed to develop a new spectrum-sharing paradigm for networks with secrecy. In this context, network coexistence can be seen as a mutualistic interaction where the primary network allows secondary activities in non-idle channels to enhance its secrecy. The goal of this work is to establish a new framework using interdisciplinary theories of communication, information, mathematics, and ecology that aims to the design and analysis of spectrum sharing strategies under secrecy constraints. This project is carried on in collaboration with the group of theoretical ecology, led by professor Carlos Castillo-Chavez (Arizona State University), which will help explore network subpopulation dynamics and their effects on spectrum sharing and network secrecy.

To improve outreach and diversity, the group is committed to attracting undergraduates and underrepresented minorities and to giving them exposure to theoretical and experimental research at all levels. The group has a strong track record for hosting students from both the Undergraduate Research Opportunities Program and the MIT Summer Research Program. In spring 2013, the group hosted two minority students from Arizona State University. Professor Win maintains dynamic collaborations and partnerships with academia and industry, including the University of Southern California; the University of California, Santa Barbra; Arizona State University; the University of Bologna and the University of Ferrara, in Italy; the University of Lund, in Sweden; the University of Oulu, in Finland; the National University of Singapore and Nanyang Technological University, in Singapore; Draper Laboratory; Jet Propulsion Laboratory; and Mitsubishi Electric Research Laboratories.

Highlights, Awards, and Events

AY2013 was another very successful year for LIDS. In addition to a continuing record of research accomplishments and intellectual excitement, the laboratory has grown in every meaningful dimension. Research volume continues to grow, exceeding our ambitious strategic plan and more than doubling over the four-year period that has just ended. LIDS welcomed additional faculty as members of the community, made major strides in engaging other units across MIT, took leadership roles in several major Institute-wide initiatives and centers, continued the series of events and activities that have added so much to the environment, and continued to lead and host major research symposia and workshops, thus continuing and enhancing its position of leadership in the international research community.

Successful activities within LIDS include the colloquium series and the 17th annual LIDS student conference. A seventh edition of the LIDS community-oriented magazine *LIDS-ALL* was produced, and the eighth edition is being readied for the start of AY2014. LIDS faculty are playing leadership roles in several major Institute-wide initiatives, including the Future Urban Mobility program within SMART, and the Virtual Center on Connection Science and Engineering. All these activities bring together researchers from all Schools at MIT, bridging disciplines from engineering to economics to the social sciences.

LIDS continues to host major workshops and symposia in emerging areas, including network science, analysis of social networks, and game theory. Finally, LIDS students and faculty members continue to receive awards and significant recognition for their accomplishments.

Awards

- LIDS alumnus Amir Ali Ahmadi received the Simons Travel Award from the American Mathematical Society.
- Former LIDS postdoctoral associate Anima Anandkumar was awarded a 2013 Microsoft Research Faculty Fellowship.
- Professor Ozdaglar was the inaugural recipient of the Steven and Renee Finn Innovation Fellowship, given in recognition of her accomplishments and research push into new areas of great potential.
- Professor Parrilo was a recipient of the EECS Faculty Research and Innovation Fellowship, given to senior faculty members in recognition of outstanding research contributions and international leadership.
- Professor Polyanskiy was selected for an NSF CAREER Award for his proposal "Information Theory Beyond Capacity." In addition, he was named the Robert J Shillman Career Development Professor of EECS.
- Alborz Geramifard, a postdoctoral associate supervised by Professor How, received a scholarship from the Conference on Uncertainty in Artificial Intelligence 2013.
- Kuang Xu, a student supervised by Professor Tsitsiklis, was awarded a 2013 Claude E. Shannon Research Assistantship.

Honors

- LIDS alumnus and former EECS and LIDS faculty member Nils Sandell, Jr. was named director of DARPA's Strategic Technology Office.
- Professor Shah was named to the advisory board of Compass Labs in March 2013.
- Professor Tsitsiklis was elected chair of the Council of the Haropio University in Athens, Greece.
- Ermin Wei was selected as one of MIT's Graduate Women of Excellence. She is advised by Professor Ozdaglar.
- Alan Willsky delivered the William Gould Dow Distinguished Lecture at the University of Michigan, in January 2013; this lectureship is the highest honor offered by the Department of Electrical Engineering and Computer Science at the University of Michigan. Professor Willsky also gave the Dean Lytle Electrical Engineering Endowed Lecture at the Department of Electrical Engineering and Computer Science at the University of Washington in May 2013; the lecture series is the department's premiere annual event, featuring internationally renowned researchers in the field of communications and signal processing.
- Professor Mitter gave a lecture series titled Connections between Nonlinear Filtering, Information Theory, and Statistical Mechanics at Carnegie Mellon University in April 2013; at Tata Institute of Fundamental Research, Bangalore, in January/February 2013; and at Princeton University in November 2012.

- Professor Ozdaglar was a plenary speaker at the International Conference on Game Theory at Stony Brook University, NY; a keynote speaker at the 3rd International Federation of Automatic Control Workshop on Distributed Estimation and Control of Networked Systems; a plenary speaker at the 6th International Institute for Computer Sciences, Social Informatics, and Telecommunications Engineering Conference on Performance Evaluation Methodologies and Tools; a keynote speaker at the FuturICT Workshop 2013 at the Media Lab, MIT; and a plenary speaker at the SIAM Conference on Control and its Applications 2013.

Paper Awards

- Amir Ali Ahmadi, together with MIT co-authors Anirudha Majumdar and Russ Tedrake, received the Best Paper Award at the 2013 IEEE International Conference on Robotics and Automation.
- Professor Win, together with student Henghui Lu and postdoctoral fellow Santiago Mazuelas, received the Best Paper Award at the 2013 IEEE International Conference on Communications, for "Ranging Likelihood for Wideband Wireless Localization." He received three 2012 IEEE Communications Society paper awards: the Fred W. Ellersick Prize, for "Network Localization and Navigation Via Cooperation," with students Yuan Shen and Wesley Gifford, and Dr. Mazuelas; the Stephen O. Rice Prize in the Field of Communications Theory, for "Optimized Simple Bounds for Diversity Systems," with student Gifford; and the William R. Bennett Prize, for "Cognitive Network Interference," with postdoctoral fellow Alberto Rabbachin.
- Professor Modiano, alumnus Krishna Jagganathan, and their collaborator at Qualcomm received the 2012 IEEE WiOpt Best Paper Award.
- Professor Tsitsiklis and student Kuang Xu received the Best Paper Award at the 2013 Association for Computing Machinery's Special Interest Group on Measurement and Evaluation (SIGMETRICS), where Xu also received the Kenneth C. Sevcik Outstanding Student Paper Award for the same paper. This is the first time that both awards were given to the same paper at SIGMETRICS since its inception, in 1973.
- Professors Parrilo and Tsitsiklis, together with alumni Amir Ali Ahmadi and Alex Olshevsky, received the 2012 Institute for Operations Research and the Management Sciences (INFORMS) Computer Society Award, recognizing the significant contributions of three papers.
- Professor How was a recipient of the American Institute of Aeronautics and Astronautics (AIAA) Best Paper Award, from the 2011 AIAA Guidance, Navigation, and Control Conference, for "Decentralized Information-rich Planning and Hybrid Sensor Fusion for Uncertainty Reduction in Human-robot Missions," with former student Sameera Ponda, et al., presented August 2012.
- Yuan Zhong (PhD September 2012, co-advised by Professors Shah and Tsitsiklis), received honorable mention for the 2012 INFORMS George Nicholson Student Paper Award.

Future Outlook

During the past year, LIDS has moved well beyond the goals it articulated in late 2008 to rebuild (and more than double) its research base, to re-establish its intellectual leadership within and beyond MIT, and to re-energize its community. The laboratory's efforts over the past few years have also led to its taking leadership roles in several major Institute-wide initiatives, and LIDS expects 2013–2014 to be a year in which its leadership in activities within and outside MIT will continue to grow. It also anticipates continuing its pattern of hosting major symposia and workshops in emerging areas, including social networks, network science, and big data.

While gratified that LIDS has established its leading role in all these activities and has cemented its standing as one of the leading international centers in the information and decision sciences, the laboratory is especially proud of the exciting and fulfilling environment it has developed for all those who call LIDS home. LIDS intends to build on its successes to ensure continued intellectual leadership and fulfillment of its promise. LIDS will experience a change of leadership during the upcoming year and is confident that the new leaders will inherit the responsibility for a laboratory that is strong in all dimensions, from the financial to the intellectual.

Alan S. Willsky
Director
Edwin Sibley Webster Professor of Electrical Engineering and Computer Science

John N. Tsitsiklis
Co-associate Director
Clarence J. Lebel Professor of Electrical Engineering

Pablo A. Parrilo
Co-associate Director
Professor of Electrical Engineering