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 serves as a melting pot and focal point for disciplines that share a common approach to problems and a common mathematical base with an energized environment that fosters the research needed for the future and instills in our students the disciplinary depth and interdisciplinary understanding required of research and engineering leaders of today and tomorrow.

At present, the faculty members 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 (ORC); 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 (CEE); 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 Sloan School of Management; the Media Lab; and the Harvard-MIT Health Sciences and Technology Program. Most significantly, during this past year LIDS has played a central role in the process of defining the new entity that is to combine a major presence in information and decision systems and a growing presence in 21st-century statistics with a concerted effort to address major societal issues that fall under the broad category of sociotechnical systems. This now places LIDS in a very exciting position in which all of our strengths can be leveraged considerably and where LIDS has the opportunity to play a major leadership role as this new entity unfolds.

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—most importantly the new entity mentioned in the preceding paragraph, but also the Future Urban Mobility project within the Singapore-MIT Alliance for Research and Technology (SMART) and MIT's Transportation Initiative.

LIDS researchers continue to have great success in obtaining funding for a broad and deep research agenda, and continue to develop relationships with industrial organizations and national laboratories such as Draper Laboratory, Lincoln Laboratory, the Los Alamos National Laboratory, Siemens, the Shell Oil Company, Honeywell, the Ford Motor Company, SNCF, 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 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 of these dimensions. The convergence of issues that arise in the challenges of the present and future has led us 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, and 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 large part, LIDS involvement in the creation and leadership of the new entity mentioned previously is a recognition of both the importance of the analytical information and decision sciences and the many interdisciplinary challenges and opportunities that are central to the mission of the new entity and that will provide exciting new opportunities and collaborations across all of MIT's schools for researchers in LIDS.

Driving forces behind research in LIDS (and in this new entity) include the availability of increasingly capable sensing, communication, and computation systems enabling the collection and transfer of large amounts of data pertaining to complex and heterogeneous interconnected systems, including the so-called sociotechnical systems that involve combinations of engineering systems, activities of networks of agents and individuals, and institutional structures. 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 contributed 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 of the emerging challenges that we see before us. As a result, many of the research themes on which LIDS researchers now work are aimed precisely at these challenges. These 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/multigranularity modeling, including methods for describing complex phenomena at multiple granularities, learning of such models from complex and heterogeneous data, and reduction/simplification of 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.

It is our expectation that LIDS involvement and leadership in this new entity will spur further work in these and other foundational areas that are crucial to the role information and decision sciences must play to meet the challenges articulated in the mission statement of the new entity.

Faculty Activities

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

Dimitri Bertsekas

Professor Bertsekas, McAfee professor of electrical engineering, is interested in deterministic optimization problems and the role of convexity in solving them, possibly through the use of duality. He has recently 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 is completing a companion textbook on convex optimization algorithms that includes some of his recent 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, logistics, and 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. Recent efforts have focused on developing policy iteration methods for problems in which measurability is a concern and resolving the related long-standing theoretical issues. In addition, the study of semicontractive dynamic programming models was undertaken, as a follow-up to a research monograph on the subject published in 2013. In particular, two broadly applicable shortest path models, involving robust path planning and infinite space shortest path problems, were studied using the semicontractive framework.

Robert C. Berwick

Professor Berwick's research during the past year made significant advances regarding the biological and computational origins of human language and cognition. Working in conjunction with professor Shigeru Miyagawa, he developed additional experimental support for the integration hypothesis model of the initial evolutionary appearance of human language. By examining many extant languages, it was found that they all obey a particular syntactic structure in terms of alternating expressive and lexical (word-based) "layers," as expected by the integration model. He presented these findings at the American Museum of Natural History in New York. Professor Berwick also commenced a new experiment-driven research collaboration to investigate the connection between learned bird song and human language. The first step has been to investigate how songbirds perceive rhythmic structure—is it in the same form as in human language? Initial results indicate that bird song has the same easy-to-learn, linear structure as human language, enabling it to be acquired from simple examples provided by adults.

In addition, Professor Berwick developed new methods for analyzing the way in which languages change over time and are related to each other. He developed a new model for phylogenetic analysis that is based on the details of how each individual acquires his or her language, rather than lumping all individuals who speak a particular language into one group. This has permitted a more detailed analysis of the spread of particular language groups as far back as 8,000 to 10,000 years ago. The findings have been compared against empirical data on the spread of these language groups. (Professor Berwick was on sabbatical leave from July 2013 to July 2014.)

Munther Dahleh

Efficiency-Risk Tradeoffs in Electricity Markets with Dynamic Demand Response

Professor Dahleh, acting director of the Engineering Systems Division and acting director of LIDS, along with graduate student Qingqing Huang and principal research scientist Mardavij Roozbehani, is studying the impact of dynamic demand responses in future smart grids. In this context we are examining, in an abstract framework, how a tradeoff between efficiency and risk arises under different market architectures. We first assess system performance under noncooperative and cooperative market architectures, both under marginal production cost pricing. The statistics of the stationary aggregate demand processes show that, although the noncooperative load scheduling scheme leads to efficiency loss, the stationary distribution of the corresponding aggregate demand process has a smaller tail, resulting in less frequent aggregate demand spikes. We are also investigating how real-time electricity pricing can be used as a tool by system operators to achieve a target tradeoff between efficiency and risk in a noncooperative set-up. Furthermore, we are providing a convex characterization of the Pareto front of system performance measures, which serves as a benchmark of the tradeoffs for system operators to evaluate pricing rules.

Robust Network Routing Under Cascading Failures

Professor Dahleh continues to study resilience of networks with his collaborators Giacomo Como and Ketan Savla (past postdoctoral fellows). We propose a dynamical model for cascading failures in single-commodity network flows. In the proposed model, the network state consists of flows and activation status of the links. Network dynamics are determined by a (possibly state-dependent and adversarial) disturbance process that reduces flow capacity on the links and routing policies at the nodes that have access to the network state but are oblivious to the presence of disturbance. Under the proposed dynamics, a link becomes irreversibly inactive as a result of overload conditions. The coupling between link activation and flow dynamics implies that links that become inactive successively are not necessarily adjacent to each other, and hence the pattern of cascading failure under our model is qualitatively different than standard cascade models. The magnitude of a disturbance process is defined as the sum of cumulative capacity reductions across time and links of the network, and the margin of resilience of the network is defined as the infimum over the magnitude of all disturbance processes under which the links at the origin node become inactive. We propose an algorithm to compute an upper bound on the margin of resilience for the setting where the routing policy has access only to information about the local state of the network. For the limiting case in which the routing policies update their action as fast as network dynamics, we identify sufficient conditions on network parameters under which the upper bound is tight under an appropriate routing policy. Our analysis relies on making connections between network parameters and monotonicity in network state evolution under the proposed dynamics.

How Networked Agents Make Decisions: Coordination with Local Information

Professor Dahleh, along with his student Spyros Zoumpoulis and professor John Tsitsiklis, studies decision making by networked entities by examining coordination among strategic agents that share local information. In particular, we study the role of local information channels in enabling coordination among strategic agents. Building on the standard nite-player global games framework, we show that the set of equilibria of a coordination game is highly sensitive to how information is locally shared among agents. Specifically, we show that the coordination game has multiple equilibria if there exists a collection of agents (1) that do not share a common signal with any agent outside of that collection and (2) whose information sets form an increasing sequence of nested sets, referred to as a filtration. Our characterization thus extends the results on the uniqueness and multiplicity of equilibria in global games beyond the well-known case in which agents have access to purely private or public signals. We also provide a characterization of how the extent of equilibrium multiplicity is determined by the extent to which subsets of agents have access to common information: we show that the size of the set of equilibrium strategies increases with the extent of variability in the size of the subsets of agents observing the same signal. We study the set of equilibria in large coordination games, showing that as the number of agents grows, the game exhibits multiple equilibria if and only if a nonvanishing fraction of the agents have access to the same signal.

Emilio Frazzoli

Professor 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 work on asymptotically optimal sampling-based algorithms such as RRT*, Frazzoli and his group continued developing the theory of sampling-based algorithms and their application in planning, control, and estimation problems.

Two new algorithms were developed. The first, called RRTX, is the first asymptotically optimal sampling-based motion planning algorithm for real-time navigation in dynamic environments (containing obstacles that unpredictably appear, disappear, and move). Whenever obstacle changes are observed (e.g., by onboard sensors), a graph rewiring cascade quickly updates the search graph and repairs its shortest-path-to-goal subtree. Both graph and tree are built directly in the robot's state space, respect the kinematics of the robot, and continue to improve during navigation.

The second algorithm is a deterministic asymptotically optimal feedback motion planning algorithm. The algorithm is based on two well-established numerical practices: (1) the fast marching method, which is a numerical Hamilton-Jacobi-Bellman solver, and (2) the adaptive mesh refinement algorithm, designed to improve the resolution of a simplicial mesh and, consequently, to reduce the error of a numerical solution. Using uniform mesh refinement, we show that a sequence of numerical solutions converges to the optimal one. To reduce the computational cost of the algorithm, we exploit the dynamic programming principle to focus the refinement in the vicinity of the optimal trajectory. Numerical experiments confirm that our algorithm outperforms previous asymptotically optimal planning algorithms, such as PRM* and RRT*, in that it uses fewer discretization points to achieve similar-quality approximate optimal paths.

Autonomous Vehicle Development

Self-driving cars are no longer restricted to science fiction but are becoming a reality. In fact, several companies and universities have demonstrated vehicles that can drive autonomously in traffic, in the process building social awareness and pushing the boundaries of current regulations and risk management practices. Professor Frazzoli, primarily through his involvement in the Singapore-MIT Alliance for Research and Technology and in collaboration with professor Daniela Rus, led the development of novel, low-cost experimental test beds for autonomous (self-driving) car technologies. The current fleet of autonomous vehicles includes several golf carts and a Mitsubishi iMiEV electric car called SCOT ("Shared Computer-Operated Transport").

Current obstacles to autonomous vehicles are no longer found in technological feasibility but, rather, affordability and safety certification. In order to address safety certification, we have developed techniques that allow our cars to generate only trajectories that satisfy all of the "hard" rules while satisfying as many of the "soft" rules as possible. These techniques are based on the synthesis of a number of concepts from robotics and theoretical computer science, as follows.

- *Incremental concretization:* We incrementally construct a graph containing a (large but finite) number of trajectories that the vehicle can execute; this can be done in a way that "explores" efficiently all possible behaviors of the vehicle. In practice, we generate thousands of new trajectories per second.
- *Incremental model checking:* Before adding trajectories to the graph, we check which trajectories are "safe" with respect to a number of prioritized temporal/logical constraints expressed in languages such as linear temporal logic or mu-calculus. This can be done incrementally, in the sense that the computation remains efficient even when the graph of trajectories becomes very large.
- *Anytime control:* At any given time we need to issue a command to the car, the "safest" trajectory (i.e., a trajectory satisfying all of the hard constraints and as many high-priority soft constraints as possible) is extracted, and the relevant commands are sent to the car. Furthermore, the car's behavior converges to optimal as the computational resources increase. Note that the car maintains what we call "safety invariance" by never abandoning a safe state unless it is certain to reach another safe state.

We were able to demonstrate the safety and correctness of our approach on our autonomous vehicles, which, to the best of our knowledge, represents the first instance of a full-size robotic vehicle with real-time control software that is provably correct with respect to a set of temporal-logic formulas expressing rules of the road.

Regarding affordability, one of the objectives of our work was to reduce dramatically the cost of driverless cars while keeping most of the capabilities demonstrated by others. The sensors and computers on our car cost about 30,000 SGD and, while not providing the full functionality of other self-driving cars (such as Google's), they are able to perform in a safe and reliable way a set of tasks that we claim are sufficient to provide a useful service to urban passengers.

Our solution is based on simple and cheap 2D LIDARs (laser range-finders), replacing the 3D LIDARs used by others, which are employed not only for obstacle detection but also for localization. Mounting the 2D LIDARs in a "pushbroom" configuration (i.e., slanted downward), we are able to detect 3D features such as curbs, roadside fittings, and buildings and use these elements to localize the vehicle within a map. The map is constructed automatically as a human operator drives the vehicle along the intended paths, which has allowed us to provide demos at unprepared sites.

Experiments show that our ability to localize the car is excellent and reliable, consistently of the order of a few centimeters in location. (The accuracy of the position in the longitudinal direction may be low on long, featureless straight segments, but it reverts back to the centimeter level as soon as the car approaches a turn, a driveway, or an intersection.)

Mobility on Demand Systems

A natural question to ask when working on self-driving car technology is the following: what is the point of autonomous cars? Some of the often-cited benefits include increased safety, decreased congestion, and reduced environmental impact. While these are undeniable and desirable benefits, we believe that the main benefit of autonomous vehicles is the enabling of widespread car sharing. Car-sharing services are growing worldwide but typically do not offer one-way, door-to-door rental options. If they do, they often suffer from limited car availability. If shared cars could drive to a customer's pickup location and return to a parking station by themselves upon drop-off, they could offer an unprecedented level of convenience. Financially, car sharing distributes the cost (in terms of dollars and time) of purchasing, maintaining, and insuring autonomous vehicles across a large user base. In addition, autonomous shared cars will eliminate the need to look and pay for a parking spot and will reduce the time needed to walk to and from the parked vehicle.

We investigated these claims in a case study in Singapore based on actual census and mobility data. An estimate of the average benefit of autonomous vehicle sharing, assuming a sharing factor of four (i.e., each autonomous car replaces four private cars), is about $8,700 a year, which at the aggregate level amounts to $1.8 trillion each year in the United States. Our analysis shows that a fleet of 300,000 autonomous shared cars could serve the mobility needs of the entire population of Singapore with average waiting times of no more than 20 minutes at rush hour. This is a substantial reduction from the 800,000 passenger vehicles in operation today. While this analysis is the result of a thought experiment (replacing all transportation modes with car sharing) and is clearly not a realistic scenario, it sends a compelling message.

Within the limits of the approximations and assumptions made, the above estimates suggest that autonomous driving is indeed a transformative technology, with a potential financial benefit to the United States on the order of more than $3 trillion per year. It is interesting that the benefits due to increased safety and reduced congestion pale in comparison with those due to sharing and increased productivity. In particular, the synergy between autonomy and car sharing is readily apparent.

Stability and Control of Transportation Networks

Professor Frazzoli and his group explore modern technologies for analyzing, in real time, massive amounts of traffic data as well as drivers' behavior (e.g., through traditional means such as traffic lights or more innovative concepts such as congestion charging and dynamic pricing). In particular, they have 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 the basis of an agreement with Singapore's ST Electronics and Land Transport Authority for a possible field trial. New refinements to the algorithm include provisions for operational constraints and for finite-capacity queues.

Bacterial Motility

After a conversation with professor Roman Stocker (a faculty member in CEE), an experimental microbiologist with strong interests in bacterial mobility and chemotaxis, Professor Frazzoli learned that he had observed some phenomena that could not be explained with the "standard" models in the literature. Professor Frazzoli noticed similarities with current problems in control theory and robotics and began looking at Professor Stocker's experimental results to understand the reasons for his observations. What he found is that by relatively simple differences in their motility mechanisms (i.e., one flagellum versus many), different species of bacteria can exhibit widely different behaviors. For example, he and Professor Stocker were able to show mathematically that while the steady-state density of many species of bacteria is proportional to the density of nutrients (a known phenomenon), the steady-state density of other species is proportional to the square of the nutrient density. The latter species are thus able to concentrate better at peaks in the nutrients, a behavior that is advantageous when nutrient sources are sparse and ephemeral, such as in marine environments (as opposed to, e.g., the human gut).

Indeed, this is exactly the phenomenon that Professor Stocker noticed in his experiments, and it is essentially due to the fact that a bacterium with a single flagellum is able to reverse its direction of motion (i.e., maintain a "memory" of the previous direction), while bacteria with multiple flagella effectively "forget" their direction of motion very quickly.

This work points out previously unknown connections between control systems theory (stochastic hybrid systems, nonholonomy) and microbiological systems. These new insights can not only be used to better understand bacterial mobility; they also can provide new ideas to build swarms of small, minimalistic robots that, instead of rejecting the randomness present in the environment, can embrace it and exploit it as bacteria do, to exhibit highly sophisticated behaviors with minimal individual complexity.

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 automated landing for uninhabited helicopters, learning in games, and simultaneous state and input estimations.

Jonathan How

Professor 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 principal investigator for the Aerospace Controls Laboratory (ACL). Recent research includes the following.

Real-Time Multiagent Decision Making

Professor How, with student Luke Johnson, developed a distributed task-planning algorithm that provides provably good, conflict-free approximate solutions for heterogeneous multiagent/multitask allocation problems. His past work introduced the Hybrid Information and Planning Consensus (HIPC) algorithm, which incorporates all information available in a network to implicitly coordinate with network neighbors while still retaining performance and convergence guarantees. Recent work has extended HIPC to be usable with corrupted information. The solution detects when information quality is too poor to reliably predict other agents' decisions and then reverts to reactive planning techniques, preserving the convergence and performance guarantees of the final allocations. The result is an algorithm that efficiently uses information in decentralized and communication-contested environments.

Professor How and students Yu F. Chen and Kemal Ure developed a hierarchical decomposed multiagent Markov decision process (MMDP) approach that autonomously identifies different degrees of coupling in reward functions and decomposes the MMDP into a hierarchy of smaller MDPs that can be solved separately and cheaply. The solutions to the smaller MDPs are intelligently combined to form a policy for the original MMDP.

Mission Technologies

In collaboration with John Vian of Boeing and Rajeev Surati of Scalable Display Technologies, Professor How, postdoc Ali Agha, and students Shayegan Omidshafiei, Yu Fan Chen, and Kemal Ure have developed real-time animation software and corresponding hardware tailored toward presenting the belief space of multiagent planning and learning algorithms. A visualization demonstrating the belief space is rendered in real time and then projected onto the experiment area, with latent data and physical systems (e.g., robots) being run synchronously. This "window into belief space" allows spectators to observe hardware while simultaneously gaining an intuitive understanding of underlying decisions made by planning and decision-making algorithms. Experiments involving autonomous teams of ground and air vehicles have been run in conjunction with the projector visualization. The developed system has been tested in several different scenarios, such as multiagent intruder monitoring. In each case, the visualization is used to improve understanding of the vehicles' behaviors through meta-data such as vehicle position, health state, and viability of future actions.

Decentralized Bayesian Learning

Professor How and student Trevor Campbell developed an approximate decentralized Bayesian inference algorithm for unsupervised models, with a focus on addressing the issues presented by mobile robotic networks. Past distributed inference algorithms for unsupervised models required high rates of communication over reliable, high-bandwidth networks. In contrast, the present method is designed to allow a team of mobile robots to collaboratively learn about their environment from streaming data despite a dynamic network structure with long periods of poor connectivity. The key insight in this work is that, for many Bayesian models, approximate inference schemes destroy symmetry and dependencies in the model that are crucial to the correct combination of local posteriors. The proposed method addresses this issue by including an additional optimization step in the combination procedure that accounts for these broken dependencies. This research paves the way for learning Bayesian nonparametric models from unbounded data sets in massively parallel architectures.

Low-Variance Asymptotic Batch-Sequential Inference

Professor How and student Trevor Campbell, along with external collaborators from Duke University and Ohio State University, have developed the dynamic means and spectral dynamic means algorithms. The dynamic means algorithm is 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. The spectral dynamic means algorithm is derived from a spectral relaxation of the cost function underlying dynamic means and extends the modeling capabilities of dynamic means to graph clustering problems. These advancements are the first low-variance asymptotic clustering algorithms suitable for the persistent streaming data commonly encountered in autonomous systems.

Information-Gathering Networks

In many large-scale inferential settings, acquisition or processing of data can be resource intensive. Selecting what data to supply to the inference engine requires computing a measure of data informativeness and, given the combinatorial nature of the problem, bounding the suboptimality of heuristic observation selectors. Professor How and student Dan Levine have proposed a novel algorithm (called embedded tree-based mutual information quantification, or ET-MIQ) that uses embedded substructures to quantify observation informativeness in the presence of nuisance variables for large-scale Gaussian Markov random fields with cycles (i.e., Gaussian distributions with nontrivial correlation structures). Whereas nuisances prevent application of submodularity-based performance bounds, How and Levine have introduced the concept of submodular relaxations, which can be used to generate online-computable suboptimality bounds in general Markov random fields (i.e., arbitrary distributions) with nuisances. These contributions significantly extend the theory of sparsity exploitation in Bayesian experimental design.

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 postdoc Girish Chowdhary, and Professor How extended the incremental feature dependency discovery (iFDD) algorithm to the decentralized iFDD (Dec-iFDD) algorithm, enabling scalable distributed learning for large-scale planning missions with incomplete models. Dec-iFDD allows agents to assess and exchange relevant features to improve model quality.

Health-Aware Planning in Partially Observable Environments

In persistent missions, taking system health and capability degradation into account is an essential factor to predict and avoid failures. The state space in health-aware planning problems is often a mixture of continuous vehicle-level and discrete mission-level states. This in particular poses a challenge when the mission domain is partially observable and restricts the use of computationally expensive forward search methods. Postdoc Ali Agha, student Kemal Ure, and Professor How (in collaboration with John Vian of Boeing) developed a method that exploits a structure existing in many health-aware planning problems and uses a two-layer planning scheme. The lower layer acts on vehicle-level states, while the higher layer focuses on mission-level variables. This planning scheme limits the expensive online forward search to the mission-level states and thus predicts system behavior over longer horizons in the future. The performance of the method has been demonstrated on a long-duration package delivery mission using quadrotors in a partially observable domain in the presence of constraints and health/capability degradation.

Cooperative Planning in Partially Observable Environments

Cooperative planning under uncertainty aims at exploiting capabilities of different (possibly heterogeneous) agents in a team to increase the chances of successfully performing a given task. Student Shayegan Omidshafiei, postdocs Ali Agha and Christopher Amato, and Professor How proposed a POMDP (partially observable Markov decision process) formulation that takes communication cost and the limited interaction of the agents into account to produce a policy that utilizes communication when necessary and allows the execution to be independent when possible. Then, using funnel-based techniques in belief space, the intractable multiagent POMDP is reduced to a tractable problem over a representative graph. While minimizing the communication between robots, the resulting cooperative behaviors reveal the strengths of the proposed method in exploiting different capabilities of individuals in a group to increase mission success probability and achieve planning results that are beyond the sum of the capabilities of the individual agents.

Bayesian Nonparameteric Models for Planning and Control

Classical approaches for learning models to improve mission planning typically assume that 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 leverages powerful nonparametric learning techniques to build models for control without requiring such background knowledge.

Professor How and students Sarah Ferguson, Brandon Luders, and Robert Grande developed a framework for safe trajectory planning for autonomous vehicles. In urban environments, autonomous vehicles must be able to quickly assess the future intentions of dynamic agents. Pedestrians are particularly challenging to model, as their motion patterns are often uncertain and/or unknown a priori. This project involves a novel changepoint detection and clustering algorithm that, when coupled with offline unsupervised learning of a Gaussian process mixture model, enables quick detection of changes in intent and online learning of motion patterns not seen in prior training data. The resulting long-term movement predictions demonstrate improved accuracy relative to offline learning alone, in terms of both intent and trajectory prediction. When these predictions are embedded within a chance-constrained motion planner, trajectories that are probabilistically safe with respect to pedestrian motions are identified in real time. Hardware experiments demonstrate that this approach can accurately predict motion patterns from onboard sensor/perception data and facilitate robust navigation within a dynamic environment.

Students Robert Klein and Shayegan Omidshafiei and Professor How have collaborated with Hongchuan Wei, Wenjie Lu, and professor Silvia Ferrari of Duke University on an approach for learning unknown nonlinear target dynamics. This work utilizes a nonparametric Dirichlet process/Gaussian process mixture model, which provides the flexibility required for clustering time-invariant spatial phenomena. The method is supplemented with a particle filter search in order to allow efficient tracking of a multiagent system. The group's work has been demonstrated through numerical simulations and physical experiments in the Real-time indoor Autonomous Vehicle test Environment (RAVEN) testbed at MIT.

Patrick Jaillet

The research of Professor Jaillet (Dugald C. Jackson professor and ORC codirector) has primarily focused on formulating and analyzing online and dynamic versions of classical optimization problems such as the shortest path problem, the traveling salesman problem, the assignment/matching problem, and the matroid 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 uncertainty. Methodological tools include those from online optimization (competitive analysis), stochastic optimization (robust analysis), online learning (min-max regret analysis, Bayesian updates), game theoretic concepts (price of anarchy), and their integrations.

Motivating applications include routing problems that arise from transportation and logistics networks, data communication and sensor networks, and autonomous multiagent systems, as well as dynamic resource allocation problems in various applications arising from the Internet economy (search engines and online auctions), health care (kidney exchange programs), and social interactions.

Professor Jaillet's research group at MIT this past academic year included two postdocs (Vahideh Manshadi from EECS/ORC and Thibaut Vidal from EECS), six graduate students from ORC (Arthur Flajolet, Chong Yang Goh, Virgile Galle, Swati Gupta, Nikita Korolko, and Maokai Lin), two graduate students from EECS (Dawsen Hwang and Andrew Mastin), one graduate student from CEE/ORC (Setareh Borjan), and one graduate student from Leaders for Global Operations/EECS (Augusta Niles). The research group in Singapore included three postdocs from SMART (Yossiri Adulyasak, Jie Chen, and Ali Oran) and several graduate students (from the National University of Singapore, Nanyang Technological University, and Singapore Management University).

Current funded research programs originate from the National Science Foundation (NSF) ("Online Optimization for Dynamic Resource Allocation Problems"), 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"), and SMART ("Future Mobility," a large project involving nine other MIT principal investigators).

Sertac Karaman

Charles Stark Draper assistant professor Karaman carries out research in the areas of control theory, optimization, formal methods, stochastic processes, and applied probability, with applications to robotics, mobile sensor networks, cyber-physical systems, and data-driven application systems. His current work focuses on joint sensing and control for agile robotic vehicles, high-performance robotic networks, and information-theoretic methods in robot motion.

Information-Based Planning Algorithms for Mapping and Persistent Monitoring

This area investigates planning algorithms that can govern mobile sensors to collect the most valuable information. For example, the algorithms work to obtain the lowest-entropy estimates or maximal mutual information in minimum time. The applications include, but are not limited to, environmental monitoring, robotic inspection, surveillance, and military reconnaissance. Our recent contributions include the development of efficient algorithms for the computation of mutual information in mapping applications. We also show that the widely used heuristics in planning for mapping applications are optimal under certain technical assumptions. In addition, we consider the problem of robotic monitoring of discrete events occurring in a large number of stations, where the precise times of arrivals are not known but their statistics are assumed to be available. We show that the resulting complex inference and control problem can be reduced to a quasi-convex optimization problem in one variable that can be solved efficiently. Our research in this area leads to a better understanding of some widely studied problems. We also propose the first computationally efficient algorithms that enable efficient real-time mapping using on-board computational hardware. (Joint work in this area involves professor Daniela Rus, postdoctoral associate Jingjin Yu, and graduate student Brian Julian.)

Agile Robot Motion in Stochastic Environments

In this area of study, we consider mobile robotic vehicles that travel in a cluttered environment or attempt to fulfill tasks that require visiting several locations. We answer foundational questions of the following form: What is the maximum speed that such vehicles can achieve and maintain for long periods? How does this speed depend on the agility, perception, actuation, or computation capabilities of the vehicle? To answer these questions, we formulate various control and planning problems in stochastic obstacle fields or stochastic reward fields; subsequently, we establish novel connections between these problems and suitable fundamental problems of statistical mechanics. In particular, we point out critical phenomena, phase transitions, and universality classes in certain control and planning problems. With the help of these results, we propose efficient algorithms with provable performance guarantees. Our results have applications in the design and development of agile robots that can navigate through densely cluttered environments such as forests or urban canyons. We also envision novel distributed sensing and monitoring systems that allow spatially and temporally extended situational awareness, with the new algorithms able to provide secure, rapid data collection. (Joint work involves graduate students Fangchang Ma and Quanwen Zhu.)

High-Performance, Large-Scale Agile Robotic Networks

Driverless cars are considered to be one of the most important technological advances that will occur in the next few decades. In fact, looking forward, many researchers and policymakers envision autonomous transportation networks consisting of driverless cars, trains, and airlines. We consider the following question: What is the capacity of transportation networks that employ autonomous vehicles that can negotiate shared resources such as intersections, lanes, and air space? In our recent work, we have developed algorithms that allow autonomous vehicles to negotiate intersections. Assuming stochastic arrival times for the vehicles, we develop algorithms that help schedule intersection traffic. We prove rigorous performance bounds on delays through such traffic intersections. We also consider a large transportation network consisting of several intersections and analyze the delay propagation, especially that due to rare events with significant effects. Our technical approach is inspired by foundational problems in data communication networks and nonequilibrium statistical mechanics. (Joint work involves graduate student David Miculescu.)

Sanjoy Mitter

Professor Mitter's research has spanned the broad area of systems, communication, and control. He has made fundamental contributions to control of infinite dimensional systems, nonlinear filtering, and hybrid systems. Over the last 20 years, his research has focused on the unification of communications and control. Current technological advances in control over networks have necessitated the development of a framework that encompasses information theory, statistical inference, and stochastic control. Modern ideas of statistical mechanics, in particular nonequilibrium statistical mechanics, are also relevant.

In seminal work (jointly with Nigel Newton of the University of Essex), Professor Mitter has shown how Bayesian inference is equivalent to the concept of free energy minimization in statistical mechanics. Moreover, Shannon's noisy channel coding theorem can be derived by considering it as a problem in Bayesian inference. In related work (jointly with Newton, Jean-Charles Delvenne of the Université catholique de Louvain, and Henrik Sandberg of the KTH Royal Institute of Technology), Professor Mitter has studied the thermodynamics of noisy electric circuits and shown how work can be extracted from the circuits through the intervention of a Maxwell demon. This work has deep connections to Landauer's work on fundamental limits to implementation of computation.

Professor Mitter's work on control with communication constraints (jointly with Anant Sahai of the University of California, Berkeley; Sekhar Tatikonda of Yale; Vivek Borkar of IIT Bombay; and Nicola Elia of Iowa State University) examines necessary and sufficient conditions for stabilization of control systems when sensors and controllers are linked via noisy communication channels. This work has led numerous other researchers to continue efforts in this area. In related work (with Tatikonda), he has shown how computation of the channel capacity of Markov channels can be viewed and solved using the methodology of stochastic dynamic programming.

Eytan Modiano

Professor Modiano (associate head of the Department of Aeronautics and Astronautics) 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 and scalable and meet emerging needs for high-data-rate and reliable communications. In order to address critical needs in the areas of 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 an ONR-funded project titled "Optimal Control of Wireless Networks: From Theory to Practice." This project is a collaboration between CNRG and researchers at the Naval Research Laboratory (NRL). The goal of the project 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., throughput and utility) but so far have been largely limited to theoretical frameworks. 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 NRL will facilitate the transition of this promising technology for use in military communication systems. Over the past year, the group developed new network control mechanisms that account for practical considerations in the way networks are designed and deployed. For example, the group developed a new optimal scheduling and routing algorithm that separates the two functions. This allows routing to be done at the network layer and scheduling to be done entirely at the link layer (as in most network implementations), with throughput optimality being maintained. This is in contrast with previous approaches wherein throughput optimality is achieved through a joint routing-scheduling scheme that is difficult to implement within existing network architectures. Dr. Hulya Seferoglu, who was a postdoc under this project, has joined the faculty at the University of Illinois-Chicago as an assistant professor.

In recent years, the group has turned its attention to the design of robust network infrastructures, including wireless networks, fiber-optic backbone networks, and power networks. In particular, the group has been working on an NSF-funded project focusing on the design of robust wireless networks. The project develops mechanisms for reliable communications over wireless mesh networks, which have recently emerged as a promising solution for providing last-mile Internet access. Such networks operate at the edge of the Internet and consist of wireless routers that provide a backbone over which users can access the Internet. While wired networks are highly dependable, wireless networks are notoriously unreliable due to interference, obstructions, jamming, and mobility issues. The goal of this project is to develop mechanisms for reliable communications in this challenging environment. In particular, the project adapts some of the techniques that have been successfully used to provide reliability in wired networks to the wireless environment.

The group continues to work on a project involving the design of highly robust telecommunication networks that can survive massive disruptions resulting from natural disasters or intentional attacks. The project, funded by the Department of Defense (DOD), 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 provide insights into the design of network architectures that are robust to natural disasters or physical attacks.

Over the past year, the group started work on a new DOD-funded project addressing 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 lows 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 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 that mitigate the effects of 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 Ozdaglar's research group focuses on modeling, analysis, and optimization of large-scale, dynamic multiagent networked systems. Their research draws on advances in game theory, optimization theory, dynamical systems, and stochastic network analysis. It focuses on both investigating substantive problems 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 sociotechnological networks, which include communication, decision, and power networks. The operations 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 interplay of the human element and engineering aspects in the context of networks.

One strand of her work focuses on developing models and control mechanisms for dynamic processes (e.g., opinion dynamics and epidemics) over networks. In recent work, conducted jointly with graduate student Yongwhan Lim and postdoctoral associate Alex Teytelboym, she developed a stochastic linear threshold model of diffusion of innovation in networks. Previous efforts focused on deterministic versions of this model in which agent thresholds are fixed, making the analysis intractable. We provide explicit analytical characterizations for expected numbers of adopters when thresholds are random in terms of a novel network centrality measure. This characterization allows us to identify networks that maximize overall diffusion when the seed set is random or targeted. In two recent papers, she and her colleagues (Daron Acemoglu, Giacomo Como, and Fabio Fagnani) developed 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). This work provides conditions for network structures and the extent of influence of prominent agents under which we obtain homogeneity or polarization in opinions. In another recent effort (with Professor Tsitsiklis and graduate student Kimon Drakopoulos), she studied the design of dynamic intervention mechanisms to influence epidemics over networks. 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 states. These efforts are supported by an Army Research Office joint multidisciplinary university research initiative (with the University of Pennsylvania, Stanford, and Cornell) and Draper Laboratory's directed research and development program.

Professor Ozdaglar's recent work focuses on market-based mechanisms for efficient allocation of power in decentralized electricity markets. Traditional operation of electricity markets involves the market operator (i.e., the independent system operator) forecasting a demand profile over a time horizon (for example, over 24 hours in the day-ahead market) and choosing the prices such that the supplier (generator) cost is minimized while ensuring that supply is matched to demand at each point in time. The generators participate in this market by bidding their supply curves (price and quantity produced). This operation ignores the heterogeneity of user demand on the consumer side and does not involve consumers actively participating in the market. In a recent paper (with graduate student Ermin Wei and professor Azarakhsh Malekian from the University of Toronto), Professor Ozdaglar provides a competitive model of a market structure in electricity networks that incorporate user heterogeneity together with adjustment costs due to ramp-up and ramp-down constraints of the generators. We characterize the competitive equilibrium and provide an optimization framework for its decentralized computation. Also, we use the optimization formulation to introduce a systematic framework for analyzing and controlling price volatility. In another project (with postdoctoral associate Ali Kakhbod), Professor Ozdaglar studies the market power of generators in deregulated wholesale markets. This project develops an oligopoly model of competition with conventional and wind energy and investigates the resulting market prices as a function of wind penetration into the market and diversification of portfolios of conventional thermal energy companies. This work is supported by a joint NSF Cyber-Physical Systems grant with Berkeley, Michigan, Vanderbilt, and a joint Masdar Institute/MIT project.

Another strand of her 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 Acemoglu from the MIT Economics Department and professor Tahbaz-Salehi from 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. It provides a general framework for analyzing 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 thus lead to a more fragile financial system.

Her recent work investigates the implications of strategic, selfish agents for system-wide performance on a variety of problems. In a project with Professor Acemoglu and Professor Malekian from the University of Toronto, she has 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 project involves a systematic analysis of security investments in general random networks subject to an attack and provides a characterization of equilibrium and optimal security investments in terms of novel network centrality measures. In other projects, she has developed game-theoretic models for resource allocation problems in communication networks, with a focus on pricing and investment incentives of providers and implications of competition for network performance.

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 (with professor Pablo Parrilo and professor Ozan Candogan from Duke University) addresses efficient (i.e., social welfare maximizing) 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 project focuses on valuation functions that admit a compact graphical representation (a natural assumption that holds in practical settings) and provides efficient iterative auctions with simple pricing structures. In another project (with Professor Acemoglu and professor Mohamed Mostagir from the University of Michigan), 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.

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 that 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 (sub)gradient-based methods, which suffer from a slow rate of convergence. Her recent work introduced new distributed optimization algorithms with much faster rate of convergence. In a series of papers with graduate students Ermin Wei and Ali Makhdoumi, she developed new algorithms for distributed multiagent optimization problems based on the alternating direction method of multipliers, which is a classical method for sequentially decomposing optimization problems with coupled constraints. This work shows that these algorithms converge at a faster rate and provide bounds on performance as a function of network structure. This research effort is supported by a joint Air Force Research Office multidisciplinary university research initiative with the Georgia Institute of Technology and the University of Maryland and by a joint ONR project with the University of Pennsylvania.

Pablo A. Parrilo

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

Professor Parrilo coedited (with Grigoriy Blekherman and Rekha Thomas) the book Semidefinite Optimization and Convex Algebraic Geometry (Society for Industrial and Applied Mathematics [SIAM], 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 it can be used for classes or seminars.

In collaboration with PhD student Hamza Fawzi, Professor Parrilo has developed a framework to prove lower bounds on the nonnegative rank. The nonnegative rank of a matrix A is the smallest integer r such that A can be written as the sum of r rank-one nonnegative matrices. The nonnegative rank has received a substantial amount of attention recently due to its applications in optimization, probability, and communication complexity. We study a class of atomic rank functions defined on a convex cone that generalize several notions of "positive" ranks such as nonnegative or cp (completely positive) ranks. This work's main contribution is a new method to obtain lower bounds for such ranks that improve on previously known bounds. Additionally, the bounds we propose can be computed by semidefinite programming. The idea of the lower bound relies on an atomic norm approach wherein the atoms are self-scaled according to the vector (or matrix, in the case of nonnegative ranks) of interest. This results in a lower bound that is invariant under scaling and is at least as precise as other existing norm-based bounds. We have mainly focused our attention on the two important cases of nonnegative and cp ranks in which our bounds satisfy interesting properties. For the nonnegative rank, we show that our lower bound can be interpreted as a noncombinatorial version of the fractional rectangle cover number, while the sum-of-squares relaxation is closely related to the Lovasz theta number of the rectangle graph of the matrix. In addition, we prove that the lower bound inherits many of the structural properties satisfied by the nonnegative rank (e.g., invariance under diagonal scaling, subadditivity). We also apply our method to obtain lower bounds on the cp rank for completely positive matrices. In this case we prove that our lower bound is always greater than or equal to the plain rank lower bound, and we show that it has interesting connections with combinatorial lower bounds based on edge-clique cover numbers.

Professor Parrilo and PhD student James Saunderson have studied the convex hull of SO(n), thought of as the set of n × n orthogonal matrices with unit determinant, from the point of view of semidefinite programming. We have shown that the convex hull of SO(n) is doubly spectrahedral; that is, both it and its polar have a description as the intersection of a cone of positive semidefinite matrices with an affine subspace. Our spectrahedral representations are explicit and of minimum size, in the sense that there are no smaller spectrahedral representations of these convex bodies.

In collaboration with PhD students James Saunderson and Hamza Fawzi, Professor Parrilo has studied the important question of the semidefinite representability of convex sets. A central issue in optimization is to maximize (or minimize) a linear function over a given polytope P. To solve such a problem in practice, one needs a concise description of the polytope. We are interested in representations of P using the positive semidefinite cone: a positive semidefinite lift (psd lift) of a polytope P is a representation of P as the projection of an affine slice of the positive semidefinite cone Sd+. Such a representation allows linear optimization problems over P to be written as semidefinite programs of size d. Such representations can be beneficial in practice when d is much smaller than the number of facets of P. We are concerned with so-called equivariant psd lifts (also known as symmetric psd lifts) that respect the symmetries of P. We present a representation-theoretic framework to study equivariant psd lifts of a certain class of symmetric polytopes known as regular orbitopes. Our main result is a structure theorem wherein we show that any equivariant psd lift of size d of a regular orbitope is of the sum-of-squares type, with the functions in the sum-of-squares decomposition coming from an invariant dimension subspace smaller than d^2. We use this framework to study two well-known families of polytopes, namely the parity polytope and the cut polytope, and we prove exponential lower bounds for any equivariant psd lifts of these polytopes.

Professor Parrilo and PhD student Frank Permenter have developed a practical semidefinite programming (SDP) facial reduction procedure that utilizes computationally efficient approximations of the positive semidefinite cone. The proposed method simplifies SDPs with no strictly feasible solution by solving a sequence of easier optimization problems and could be a useful preprocessing step for SDP solvers. Through choice of approximation, the method allows one to naturally trade off preprocessing effort with problem simplifications. We have demonstrated the effectiveness of the method on SDPs arising in practice and provided a publicly available software implementation. We also have created a postprocessing procedure for dual solution recovery that generally applies to techniques based on facial reduction.

Professor Parrilo and former PhD student Ozan Candogan (now an assistant professor at Duke University) designed simple and efficient iterative auctions for selling multiple items in settings where bidders' valuations belong to a special class (tree valuations) that can exhibit both value complementarity and substitutability. We provided a compact linear programming formulation of the efficient allocation problem and used it to establish the existence of a Walrasian equilibrium for tree valuations that satisfy an additional technical condition. This result reveals a new class of valuations for which a Walrasian equilibrium exists in the presence of value complementarities. We also provided an iterative algorithm that can be used for the solution of this linear programming formulation. Complementing the algorithm with an appropriate payment rule, we obtained an iterative auction that implements the efficient outcome (at an ex-post perfect equilibrium). This auction relies on a simple pricing rule and compact demand reports and uses a novel (interleaved) price update structure to assign final payments to bidders, guaranteeing truthful bidding.

In collaboration with colleagues from the University of Washington, Professor Parrilo and his student Hamza Fawzi have investigated a new concept called the positive semidefinite rank of a matrix. The positive semidefinite rank has many appealing geometric interpretations, including semidefinite representations of polyhedra and information-theoretic applications. The research group has extensively studied the main mathematical properties of the positive semidefinite rank, including its geometry, relationships with other rank notions, and computational and algorithmic aspects.

Yury Polyanskiy

Robert J. Shillman assistant professor Polyanskiy conducts research in the areas of mathematics of information (information theory), coding theory, and probability. His current work focuses on nonasymptotic characterization of the performance limits of communication systems, non-Shannon information measures, and fault-tolerant circuits.

Hypothesis Testing via Comparator and Hypercontractivity

Professor Polyanskiy investigates best achievable performance through a hypothesis test satisfying a structural constraint: two functions are computed at two different terminals, and the detector consists of a simple comparator checking whether the functions agree. Such tests arise as part of the study of fundamental limits of channel coding and hypothesis testing with communication (rate) constraints. A simple expression for the Stein exponent is found. Connections to Gács-Körner common information and to hypercontractivity properties of the conditional expectation operator are identified. In the case of the zero Stein exponent, a nonvanishing lower bound on probability of error is established by pairing specific estimates.

Quasi-Static MIMO Fading Channels at Finite Blocklengths

This area explores the maximal achievable rate for a given blocklength and error probability over quasi-static multiple-input multiple-output (MIMO) fading channels, with and without channel state information at the transmitter and/or the receiver. The principal finding is that outage capacity, despite being an asymptotic quantity, is a sharp proxy for the finite-blocklength fundamental limits of slow-fading channels. Specifically, channel dispersion is shown to be zero regardless of whether the fading realizations are available at both transmitter and receiver, at only one of them, or at neither of them. These results follow from analytically tractable converse and achievability bounds. Numerical evaluation of these bounds verifies that zero dispersion may indeed imply fast convergence to the outage capacity as the blocklength increases. In the example of a particular 1×2 single-input multiple-output Rician fading channel, the blocklength required to achieve 90% of capacity is about an order of magnitude smaller than the blocklength required for an additive white Gaussian noise channel with the same capacity. For this specific scenario, the coding/decoding schemes adopted in the LTE (Long-Term Evolution) Advanced standard are benchmarked against finite-blocklength achievability and converse bounds. (Wei Yang, Giuseppe Durisi, and Tobias Koch took part in this work.)

Dispersion of Compound Discrete Memoryless Channels

Code for a compound discrete memoryless channel (DMC) is required to have a small probability of error regardless of which channel in the collection perturbs the code words. The capacity of the compound DMC has been derived classically: it equals the maximum (over input distributions) of the minimal (over channels in the collection) mutual information. In this work, the expression for the channel dispersion of the compound DMC is derived under certain regularity assumptions on the channel. Interestingly, dispersion is found to depend on a subtle interaction between the channels encoded in the geometric arrangement of the gradients of their mutual information. It is also shown that the third-order term need not be logarithmic (unlike single-state DMCs). By natural equivalence with the compound DMC, all results (dispersion and bounds) carry over verbatim to a common message broadcast channel.

Bit Error Rate of Repeated Error-Correcting Codes

Error-correcting codes have typically been studied with respect to performance metrics such as minimum distance (combinatorial) or probability of bit/block error over a given stochastic channel. In this work, a different metric is considered. It is assumed that the block code is used to repeatedly encode user data. The resulting stream is subject to adversarial noise of a given power, and the decoder is required to reproduce the data with a minimal bit error rate. This set-up may be viewed as combinatorial joint source-channel coding. Two basic results are shown for the achievable noise-distortion tradeoff: optimal performance for decoders informed of the noise power and global bounds for decoders operating in complete oblivion (with respect to noise level). General results are applied to the Hamming [7,4,3] code, for which it is demonstrated (among other findings) that no oblivious decoders exist that attain optimality for all noise levels simultaneously. (Weihao Gao took part in this work.)

Scalar Quantization with Noisy Partitions and Its Application to Flash ADC Design

Motivated by recent circuit designs for flash ADCs with imperfect comparators, Professor Polyanskiy investigates the problem of scalar quantization with noisy partition points, where partition point locations are perturbed from the designated values by noise during the placement process. For this problem setting, a high-resolution approximation for mean square error is derived, and the optimal partition point density is analyzed accordingly. The results indicate that it is necessary to take the effect of noise into account in the design process. In particular, the optimal partition point density when the input distribution is Gaussian or uniform is derived, and, when noise variance exceeds a certain threshold, it is shown that a peculiar phase transition occurs and the optimal point density degenerates into a delta function at the origin. These theoretical results allow optimization of the design of flash ADCs. (Da Wang and Gregory Wornell took part in this work.)

Mardavij Roozbehani

Principal research scientist Roozbehani led several research efforts related to mathematical modeling, optimization, and control for power systems, with a particular emphasis on robustness analysis and architecture design for decentralized scheduling. Dr. Roozbehani also participated in other research projects on analog to digital converter design, compiler design for software verification, and robustness analysis of switched linear systems. The long-term goal of his research is to understand the sources of robustness, fragility, and systemic risk in power systems and, more generally, networked systems. Funding for his research comes from NSF, Draper Laboratory, the MIT-Masdar Institute Cooperative Program, and the MIT Skoltech Initiative.

Robustness of Networked Systems

Questions related to the robustness and fragility of networked systems have recently renewed the attention and interest of researchers across multiple disciplines, including control theory and optimization. This is driven by both the theory gap in understanding the sources of fragility and risk in networked systems and the need to address pressing and increasingly important practical problems (e.g., systemic risk and cascaded failures in power systems, transportation networks, or financial networks).

Last year, we reported results on how fragility and endogenous risk can be inherent to the architecture of the system. In particular, we reported results that showed how dynamic interactions among fully rational agents might aggregate small exogenous disturbances in the system and create endogenous risk. In our continued research effort in this area, we have identified a new metric for assessing the sensitivity of networked systems to exogenous and stochastic disturbances. This metric is called the incremental mean squared gain. It is a measure that approximates the ratio of the variance of deviations in the nominal output to the variance of stochastic perturbations in the input. We developed new criteria based on a notion of stochastic differential contraction for assessing and computing the sensitivity of networked systems as measured by the incremental mean squared gain. These results have been derived for systems described by implicit nonlinear equations that define the underlying dynamics. This is an attractive feature because many engineering systems including networks with optimization in the loop and systems with algebraic constraints (which arise in networked settings) are naturally described in this form. In our future work, we intend to pursue practical application of these results in analyses of the robustness of power systems and energy markets to disturbances.

Layered Architecture and Decentralized Scheduling for Integration of Flexible Loads

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 types of information. A layered approach consists of 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 subproblems assigned to different distributed decision makers with partial information leads to suboptimal power system operations. Proper decomposition choices exploit time scale separation, weak coupling of certain phenomena, and device redundancy to achieve a specific level of performance with a limited exchange of information between the decision makers. An intermediate layered architecture is expected to provide a tradeoff between the performance of the exact solution and the robustness and scalability of the fully decoupled solution.

An important aspect of architecture design for future power systems pertains to integration of flexible loads, that is, loads with fulfillments that can be deferred up to a deadline. Because of their dynamic nature, these loads cannot be described with the tools of elastic consumption. If properly managed, such loads have the potential to provide virtual storage at much lower cost, to enable the integration of larger shares of renewable energy generation, and to relieve peak hour congestion of the grid. Meanwhile, the presence of operational constraints of the grid (supply-demand balance, power flow constraints) requires a level of coordination among these loads that cannot be guaranteed with sufficient reliability by real-time responses to prices alone.

The objective here is to design modular and scalable protocols and computational tools that enable the participation of flexible loads in the energy market to increase grid efficiency while guaranteeing satisfaction of individual hard deadlines and safe and reliable grid operations. We researched the problem of clearing the real-time market subject to scheduling problem constraints, which consist in both the bound on total consumption and the dynamic feasibility of the scheduling problem after one-step clearing. We obtained new results on the structure of the set of all feasible decisions, which lie in a partially ordered set, and showed how they can be used for reducing the complexity of the decision process and devising decentralized scheduling schemes.

In order to extend the research toward more complex scenarios with more interesting constraints on the aggregate power consumption of the loads, a more tractable model for the interaction between the energy market and physical constraints needs to be developed. In fact, while the computational problem is so far tractable, the current approach seems difficult to extend to the complex combination of voltage, power flow, and demand/supply balance constraints. Such constraints can be effectively approximated with a set of linear constraints with positive coefficients. We therefore plan to approach the problem of coordinated scheduling of the loads via a layered approach. In this approach, a higher coordination layer produces a set of individual constraints on the maximum power rate of each bus in the power distribution grid. Once these auxiliary constraints are produced, the scheduling problem at the individual bus level becomes a problem similar to the one considered earlier and therefore becomes tractable.

Devavrat Shah

Associate professor Shah and his research group are involved in developing an understanding of complex networks so as to be able to better engineer them. This includes communication networks such as the Internet, networks of statistical dependencies observed in large data sets, and emerging social networks such as Twitter. A salient feature of 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), as efficient information processing tools in statistical networks (e.g., belief propagation), and as behavioral models in social networks (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. Specific areas of progress made in the past year are described next.

Processing Social Data: Recommendations Using Comparisons

A recommendations system primarily determines what users, communities, or entire populations may like or dislike given their prior history. Recommendation systems have become essential for all consumer-facing systems; they are the primary force behind Amazon becoming a successful e-retailer and 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 preferences from heterogeneous sources of partial preference data. Professor Shah and his group (PhD candidate Ammar Ammar, graduate student Luis Voloch, and former postdoc Sewoong Oh, now a faculty member at the University of Illinois) developed a novel approach by using comparison as the basic unit of preference data, thus capturing signals about preferences from all sorts of preference data simultaneously. Work on this ongoing project has progressed in terms of developing novel theoretical foundations as well as scalable algorithms. Earlier results had a successful impact on the industry (e.g., Netflix) and were covered in both the popular press (e.g., the New York Times) and academic publications. This work has been supported through funding from NSF.

Processing Social Data: Local Algorithms for Scalable Computation

Computing the stationary distribution of a large finite or countably infinite state space Markov chain has become central to many problems such as statistical inference and network analysis. Standard methods involve large matrix multiplications, as in power iteration, or simulations of long random walks as in Markov chain Monte Carlo methods. Such methods do not all scale well with the size of the data. We have developed a novel local algorithm that scales well with very-large-scale data (e.g., web-scale data).

John Tsitsiklis

Professor Tsitsiklis (Clarence J. LeBel professor of electrical engineering) and his 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. Their 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 analysis and optimization of resource allocation methods (routing and scheduling) in queueing networks. While the advantages of having flexible servers (each server being able to process all job types) are well known, our past research had 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.

During the last year, we focused on comparisons of various architectures: in one architecture each server is assigned at random a certain number of job types, in another architecture the overall network is organized modularly into distinct subnetworks, and so forth. Analytical results have been used to compare the qualitative performance properties of the different architectures.

Social Networks

We have continued our research on the subject of inferring the web of influences in a social network and developed analytical results on the amount of data that are needed toward this purpose. In addition, we have developed a particular formulation of the problem of controlling the spread of an epidemic in a network and designed a control policy with favorable performance characteristics.

Life Sciences

We have initiated research on the subject of central sleep apnea. A relatively simple but faithful model has been developed. The ultimate objective is to be able to calibrate such a model based on patient data and use it to make predictions and inform therapeutic interventions.

Education

While on sabbatical during AY2014, Professor Tsitsiklis focused on the development of the online 6.041x Introduction to Probability undergraduate class, offered through EdX. The course was successfully offered in spring 2014 and was well received.

Alan Willsky

Edwin Sibley Webster professor (retired) Willsky (director of LIDS through June 2014) led the Stochastic Systems Group (SSG), whose research focuses on developing statistical 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 various sources, including the Air Force Research Office, the Army Research Office, ONR, and Shell Research and Development.

One of the major directions of research in SSG continues to be the extension of our 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 structures 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, we have had significant success in developing and analyzing a class of highly parallel algorithms for sampling from very large random field models. These algorithms achieve this high level of parallelism with some loss of accuracy, and, prior to our work, the promise of these methods had been demonstrated through experiments. However, detailed analyses of these methods and theoretically justified quantifications of their accuracy (or loss thereof) were largely lacking. Our recent research provides a very thorough analysis of so-called Gaussian random fields, making clear when these methods work and what the tradeoffs are in accuracy as a function of key parameters in the algorithms. In addition, we provide for some important cases a method to recover any lost accuracy through a simple correction. We have also completed our development of a fully distributed approach to message-passing algorithms for inference in graphical models that builds on the methods we had developed in preceding years for highly accurate (but centralized) inference, exploiting a combination of both graph and statistical structures. In particular, we have developed a distributed messaging protocol that generalizes so-called belief propagation and allows each node in a network to receive messages and use the protocol information both to perform local processing and to route the results of those computations to each of its network neighbors. 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, and a major IEEE (Institute of Electrical and Electronics Engineers) award that Professor Willsky received this May in recognition of his body of work in this area.

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 our 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 electrical engineering/computer science boundary at MIT, bringing together new perspectives on these challenges. Our work in this area continues to have several components. The first involves bringing advanced methods of optimization (in particular, semidefinite programming methods) to the learning of models with sparse and exploitable structures. Our most recent contributions in this area involve the learning of hierarchical or multiresolution models that can capture correlations among variables at many different scales. A second area has focused on learning models with graphical structures that are ideally adapted to the methods for inference mentioned above and that also yield models with structures found in many naturally occurring phenomena. As with our optimization-based methods, these new algorithms 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, we have continued our work on so-called Bayesian nonparametric methods for discovering complex behaviors in temporal or spatiotemporal data. The major advances during the past year have been motivated by two different applications: motion behavior analysis of mice and speech analysis. For the former, a major issue is that of developing computationally efficient algorithms for learning with massive data sets. To address this issue, we have developed some very powerful algorithms for stochastic variational inference, algorithms that are already being employed by other researchers focusing on quite different problems. Our research motivated by problems in speech analysis involves learning more complex layered models of the dynamics of speech—more complex and accurate than, for example, the methods based on hidden Markov models that are widely known and used. Our methods involve hierarchical models that capture dynamics at the relevant time scales for speech analysis. Our work has been successfully demonstrated in this context through collaboration with speech researchers at MIT; even more important, however, these algorithms are far more widely applicable, and software for them is being made available.

Moe Win

The Wireless Communication and Network Sciences Laboratory, led by Professor Win, is involved in multidisciplinary research that encompasses developing fundamental theories, designing algorithms, and conducting experiments for a broad range of real-world problems.

To advocate 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 (UROP) and the MIT Summer Research Program. The group recently hosted an MIT UROP student and two undergraduate students as summer research interns. Professor Win maintains dynamic collaborations and partnerships with academia and industry, including the University of Southern California; the University of California, Santa Barbara; 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 Singapore University of Technology and Design and Nanyang Technological University in Singapore; Tsinghua University in China; Kyung Hee University in Korea; Draper Laboratory; the Jet Propulsion Laboratory; and Mitsubishi Electric Research Laboratories.

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 provided below.

Network Localization and Navigation

The group has continued its research on network localization and navigation, in particular developing techniques for efficient fusion of sensory data and environmental information for network localization and navigation in harsh wireless environments. Based on the proposed belief condensation (BC) engine, the group has successfully developed BC-based filtering techniques for representing the positional belief in cooperative navigation networks. The results proved that the proposed techniques can handle the nonlinear dynamics and non-Gaussian beliefs that occur in real scenarios with an improved accuracy-complexity tradeoff relative to conventional techniques. Moreover, the group has designed new algorithms that exploit the map information in real-time navigation networks. These algorithms significantly improve localization accuracy in medium to low signal-to-noise ratio scenarios. The group has also developed a framework for context-aided inertial navigation and developed efficient BC-based algorithms for its implementation. The results have provided a new methodology to fuse information from inertial measurements and situational contexts through Bayesian inference over an augmented hidden Markov model. The proposed techniques enable efficient and accurate context modeling and data fusion through BC. The theoretical framework and algorithmic techniques can enable seamless integration of contextual information in inertial navigation systems and, hence, significantly improve localization accuracy and robustness.

Intrinsic Network Secrecy

The group has also devoted effort to studying how physical-layer security can strengthen secrecy of communications by exploiting intrinsic properties of communication channels and interference. First, the group has developed a framework for the design and analysis of wireless networks with secrecy provided by properties such as node spatial distribution, wireless propagation medium, and aggregate network interference. This framework enables analysis of how network interference affects the network secrecy throughput density of a large-scale wireless network, as well as evaluation of various interference-engineering strategies that mitigate eavesdropping capabilities. Both theoretical analysis and numerical results show that network interference can be exploited to provide better secrecy performance. Second, the group has established a foundation for analyzing the performance of distributed network secrecy (DNS) and provided a new perspective on the relationship between DNS and the network lifetime of energy-efficient, self-organizing multilevel wireless networks (MWNs). The insights obtained from the research demonstrate that different configurations of transmitted energy and communication resources induce different relationships between DNS throughput and network lifetime, thus providing guidelines for the design and analysis of reliable, energy-efficient self-organizing MWNs with DNS.

Finally, the group has developed a foundation for secret-key generation using unknown deterministic sources and provided a new perspective on the potential of wideband multipath channels in mobile networks for wireless intrinsic secrecy. A new information measure called intrinsic information is proposed to characterize the achievable rate of the secret key that can be generated from an unknown deterministic source. Since statistical knowledge of the source parameter is often unavailable, a non-Bayesian scenario is considered wherein the parameters are unknown but deterministic and the intrinsic information has been derived as a function of the parameter value and noise distribution. The group has also determined the intrinsic information of a wideband propagation medium in mobile wireless networks, showing the potential of time-varying wideband channels for secret-key generation.

Highlights, Awards, and Events

Continuing successful activities within LIDS include the colloquium series and the 18th annual LIDS student conference. An eighth edition of the LIDS community-oriented magazine, LIDS-ALL, was produced, and the ninth edition is being readied for the start of AY2015. LIDS faculty are playing leadership roles in several major Institute-wide initiatives, including increased leadership in the Transportation Initiative at MIT. Most significantly, as mentioned above, LIDS has played a central role in the process of defining the new entity that is to combine a major presence in information and decision systems and a growing presence in 21st-century statistics with a concerted effort at addressing major societal issues that fall under the broad category of sociotechnical systems. These activities bring together researchers from all of MIT's schools.

LIDS also 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, as noted below.

Awards

- Dimitri Bertsekas received the 2014 American Automatic Control Council Richard Bellman Heritage Award. This prestigious award recognizes professional achievements by US control systems engineers and scientists.
- Emilio Frazzoli was awarded a Singapore research professor chair for the 2013 calendar year.
- Sertac Karaman received an NSF Faculty Early Career Development (CAREER) Award in March 2014.
- Sanjoy Mitter was elected a fellow of the International Federation of Automatic Control in acknowledgment of his fundamental contributions to nonlinear filtering, stochastic control, optimization, optimal control, hybrid systems theory, and the unification of communication and control. In addition, Professor

Mitter has been named the recipient of the 2015 IEEE Eric E. Sumner Award for outstanding contributions to communications technology.

- Asuman Ozdaglar received the Spira Teaching Award from EECS.
- Pablo Parrilo was presented the 2013 Institute for Operations Research and the Management Sciences (INFORMS) Optimization Society Farkas Prize. This prize is awarded annually at the INFORMS national meeting to a midcareer researcher for outstanding contributions to the field of optimization.
- Alan Willsky won the 2013 Society Award from IEEE's Signal Processing Society for fundamental contributions to probabilistic modeling and for pioneering work in the development and application of multiresolution statistical methods.
- Moe Win received the International Prize for Communications Cristoforo Colombo, established by the city of Genoa in 1954. The prize is awarded annually to an individual who has made significant contributions to advancing communications.

Honors

- Munther A. Dahleh was appointed to the William A. Coolidge professorship.
- Jonathan How's student Alborz Geramifard received a scholarship from the Uncertainty in Artificial Intelligence Conference.
- Patrick Jaillet was inducted as an INFORMS Fellow in 2013.
- Eytan Modiano delivered a keynote talk ("Control of Wireless Networks: From Theory to Practice") at Qualcomm's QTech Forum conference.
- Asuman Ozdaglar was a plenary speaker at the SIAM Conference on Control and Its Applications and the Modeling and Optimization: Theory and Applications conference. She gave invited talks at the Institute for Pure and Applied Mathematics Stochastic Gradient Methods workshop, the Joint Workshop on Pricing and Incentives in Networks and Systems, and the Isaac Newton Institute Systemic Risk: Models and Mechanisms workshop.

Paper Awards

- Emilio Frazzoli (with Shan Jiang, Gaston Fiore, Yingxiang Yang, Joseph Ferreira, and Marta Gonzalez) won the Best Paper Award for "A Review of Urban Computing for Mobile Phone Traces: Current Methods, Challenges and Opportunities" at the Association for Computing Machinery (ACM) Special Interest Group on Knowledge Discovery and Data Mining (SIGKDD) International Workshop on Urban Computing.
- Jonathan How's group (including LIDS student Kemal Ure and former postdoc Girish Chowdhary) won the American Institute of Aeronautics and Astronautics Best Paper Award at the Guidance, Navigation, and Control Conference for the second consecutive year.

- The Applied Probability Society of INFORMS presented its 2013 Best Publication Award to Devavrat Shah and his student Jinwoo Shin. The award recognizes outstanding contributions to the field of applied probability.
- John Tsitsiklis and his student Kuang Xu received the Best Paper Award at the 2013 ACM SIGMETRICS (Special Interest Group for the Computer Systems Performance Evaluation Community) conference; Kuang also received the Kenneth C. Sevcik Outstanding Student Paper Award. This is the first time that both awards were given to the same paper at the SIGMETRICS conference since its inception in 1973. (Note that these awards were presented in June 2013, and were not included in last year's report.)

Future Outlook

In last year's report, we indicated that our efforts over the recent past had placed us in a position from which we expected that LIDS leadership in activities within and outside MIT would continue to grow in 2013–2014. That has certainly taken place, in particular with the emerging entity bringing together LIDS, statistics, and a broad array of both units and individuals to tackle some of the most challenging and important sociotechnical problems that society faces. This now places us in a very exciting position in which all of our strengths can be leveraged considerably, and we have the opportunity to play a major leadership role as the new entity unfolds. This opportunity is accompanied by several challenges as well. One is maintaining all of our strengths, particularly the invigorating environment and strong sense of intellectual community that have made LIDS such an important international center in the information and decision sciences and an exciting home for its members. Doing so while also playing a leadership role in the new entity and leveraging the additional resources this entity will bring represents both an opportunity and a challenge. Finally, LIDS leadership will change as of July 1, 2014, and we must ensure a smooth leadership transition over the upcoming academic year. Given the commitment and quality of the people within LIDS and the vision of both LIDS and this new entity, we are confident that these challenges can be met and the associated opportunities transformed from promise into substance to the benefit of those in LIDS and MIT more broadly.

Alan S. Willsky
Director (through June 30, 2014)
E.S. Webster Professor of Electrical Engineering and Computer Science (retired as of July 1, 2014)

Munther A. Dahleh
Acting Director (effective July 1, 2014)
Professor of Electrical Engineering

Pablo Parrilo
Associate Director
Professor of Electrical Engineering and Computer Science