Sensor Webs of SmartDust

Data Fusion/Inferencing In Large Arrays of Microsensors

S. Sastry , K. Pister, J. Chandra(co-PIs)

University of California, Berkeley, George Washington University

, Faculty Investigators: V. Anantharam, L. El Ghaoui, M. Jordan, J. Malik, and S. Russell

University of California, Berkeley

1 Executive Summary

Dramatic advances in MEMS, computing, and communication technology are revolutionizing our ability to build massively distributed, easily deployed, self-calibrating, disposable sensor networks. Such networks gather multi-modal data about the environment, use both local and distributed inference algorithms to determine reliable interpretations at multiple levels of granularity, and communicate those interpretations in response to events or queries. We propose research in the four areas required in order to realize such systems: 1) sensor design; 2) algorithms for reliable communication and localization; 3) algorithms for reliable distributed information fusion and interpretation with limited communication and computation resources; 4) information-theoretic foundations for understanding the design, performance, optimization, and fundamental limits of geographically distributed sensory systems.

In a DARPA-sponsored project called SmartDust at Berkeley, we (K. Pister and co-workers) have demonstrated the possibility of constructing autonomous sensor nodes of the order of several cubic millimeters volume containing sensors (acoustic, vibration, acceleration, pressure, temperature, humidity, magnetic, biochemical), electronics, power supply, and communication hardware. To date this project has demonstrated laser communication from the SmartDust Motes over a distance of about 21 km., RF based sensor networking, and deployment via silicon "dandelion seed" and "maple seed" designs. In the next several months, we expect to construct several hundred one-cubic-inch SmartDust Motes at a cost of $50 each. We will work to redcue the cost and size of these devices by at least an order of magnitude in the next two years and to increase the onboard computational power while reducing power consumption. Our conclusion is that *Using COTS technology it is already possible to build Sensorwebs that present challenges far beyond the capabilities of existing algorithms and theory.*

Our team of experts in MEMS (Pister), Estimation and Signal Processing (Malik, Sastry), Statistical and AI Methods for Distributed Inference, Learning, Data Association and Data Fusion (Jordan, Russell), Robust Optimization and Estimation (El Ghaoui), Networking and Communication theory (Anantharam), and simulation and modeling techniques for future Army command and control (Chandra) will provide a unified approach to the development of theoretical tools, algorithms, and software for Sensorwebs. Research topics include algorithms for distributed sensor localization using robust optimization and multidimensional scaling; adaptation of computer vision techniques, such as optical flow analysis, to the sensorweb setting of scattered sensor "cells"; variational and Markov Chain Monte Carlo (MCMC) techniques for probabilistic data interpretation, including efficient global data association; unification of sensor-based low-level representations and object-based high-level representations through expressive probabilistic languages; a distributed Shannon theory deriving ultimate sensorweb performance limitations from constraints on energy consumption and computation. Our approach will be to connect rigorous theoretical advances with demonstrated capabilities on a variety of application scenarios which we describe below.

2 Motivation: Smart Dust Project

The revolution in sensor networks is being driven by the convergence of exponential improvements in three areas of electronics: digital circuits, communication circuits, and MEMS. Following Moore's Law analogs in the early 90s led several groups to the conclusion that wireless sensor networks were about to become ubiquitous and infinitesimal. The Smart Dust project, one of several DARPA-sponsored projects in the sensor networks area, has perhaps the smallest claims to date: completely autonomous sensor nodes (motes) will be smaller than a cubic millimeter by 2001. Despite their dimunitive size, these motes will by no means have meager performance - the MEMS sensors on-board will be have a wide dynamic range, and the communication systems will be capable of megabit per second links at more than a kilometer of distance. With even the first generation, cubic-inch scale "macro motes" that are functional today, the hardware has already outstripped the capability of the software and algorithms. We can easily create networks of thousands of sensors, and distribute them around areas of interest with any number of delivery systems in a matter of seconds or minutes, and yet we have no existing theory which will tell us what to do with the data being sensed - data which contains all of the information we desire, if we only had the wit to extract it.

Robust communication networks, affordable sensors and information technologies form the backbone (critical enabling layer) in the battlefield decision making architecture. Together they provide the "eyes", "ears", and "noses" for all tactical and strategic decision making. Integrated sensor systems are critical and fundamental to achieve situation awareness on the battlefield to win the information war. Sensors are vital to the survivability of soldiers and the weapon platforms on the battlefield. Intelligent, multi-domain, internetted sensors can provide robust performance at a reasonable cost (for, e.g., low power consumption). For example, low level sensor fusion coupled with integration of sensors such as acoustic, seismic, chemical, biological, and electromagnetic can provide significant capabilities to the warfighter. In a specific situation, for instance, detection of an enemy tank may be possible from an integrated suite of acoustic (sounds like a tank), olfactory (smells like diesel exhaust), magnetic (strong magnetic signature), seismic (strong seismic signature), and infra-red (something is hot) sensors.

The Smart Dust project aims to scale sensing, communication platforms down to cubic millimeter volumes. Combining CMOS technology for logic, MEMS corner-cubes for passive communication and thick-film batteries for power, Smart Dust exploits current mass production technologies. With the ability to economically produce thousands of sensor nodes (known as Smart Dust "motes"), local information can be gathered from a wide geographical area. As a showcase for the capabilities of the goal product, a macroscopic breed of Smart Dust has been developed on printed circuit boards with commercial off-the-shelf products for sensing, transmission and power. With active communication capabilities (RF, visible laser or IRDA), these "macromotes" allow for a more diverse set of network functions. We will now describe our work to date. Different breeds of macromotes are currently in use. The RF-communicating mote has a transmission range of 20 meters in open space (cf. Figure 1) at 4800bps. Visible laser motes have been tested at distances of 20km. A scanning apparatus allows the beam to scan over (this fraction) of the hemisphere and research is underway to allow two scanning motes to locate each other. All current macromotes used Atmel microprocessors that are programmed via a PC serial port.

The next generation of macromotes is currently undergoing fabrication. Major improvements include RF reprogrammability and increased communication bandwidth. The former feature permits one-to-many programming and opens up a domain of passing programs as messages through the network. One can imagine programs selectively permeating a sensor network evolving hop-by-hop based on local information and requirements. It is tempting to seek the purity of a totally distributed control scheme and reject the notion of an observer. To shun external computational possibilities entirely is, however, academic and quixotic. Using a PC interface to the Smart Dust network to send and receive messages not only provides a means of observing network behavior, but also can help maintain stability by controlling the creation of new messages. Additionally, each mote in the network is given a unique ID at the time of programming. This simplification has great logical importance.

2.1 Achieved Networking Goals

The first Smart Dust "network" application was a simple serial port connection of a macromote to a PC, monitoring sensor readings on the computer screen. Using the same hardware with different programming, wireless motes were placed throughout a room using RF communication to send messages to the PC connected mote. By sending a mote ID, the PC mote can demand sensor information from any one of the distributed sensors. Building on this direct communication, a multi-hop message system was tested. Several motes were spaced down a hallway, with the final mote gathering sensor data. This data was relayed linearly back to the PC connected mote. In contrast with the previous experiment in which only one mote collects data, the "Star Trek Locator" network uses all motes as relays and sensors (see Figure 3). The same linearly

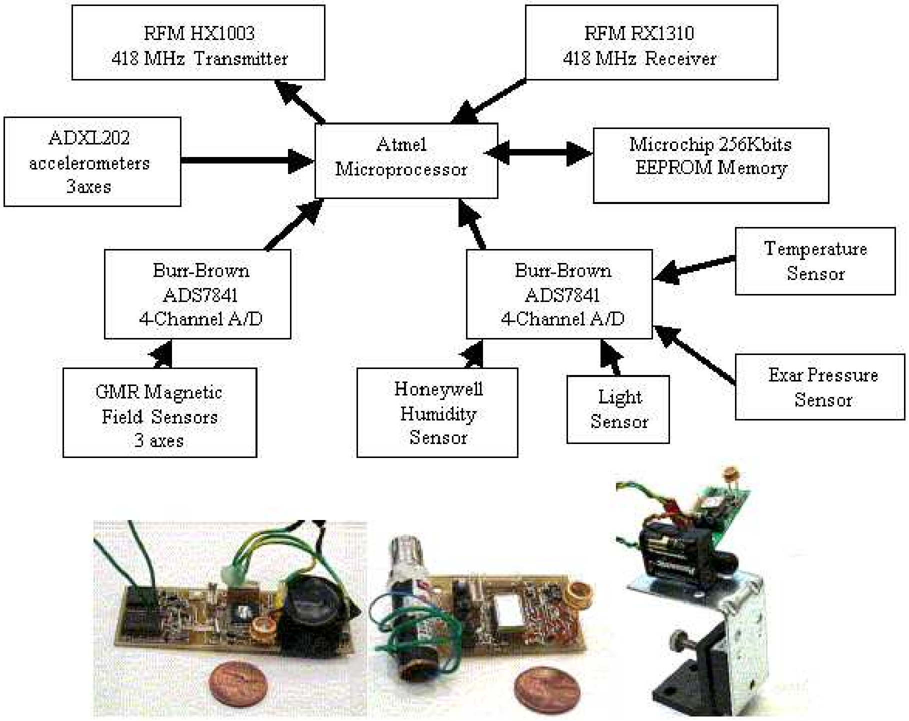

Figure 1: From left to right, RF, laser and mounted laser motes. Above the photos is a block diagram of the components of the RF mote.

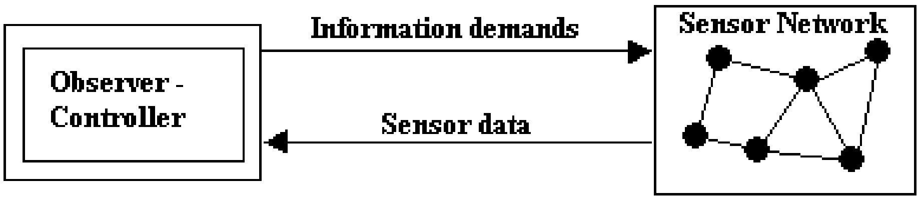

Figure 2: Flow of information. All network messages originate and terminate at the observer-controller. Each sensor mote has a unique ID but has no local or global topological knowledge.

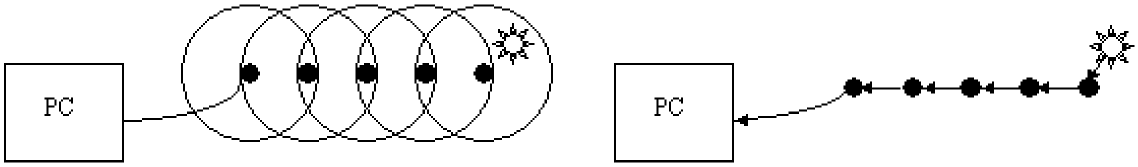

Figure 3: Star Trek Locator network. Solid dots represent motes with a range of communication given by the radius of the surrounding circle. The star represents the moving transmitter, seen by the rightmost mote. Positive identification of the target travels left to the PC as shown on the right.

located chain of motes is used, but now a single mote is programmed to periodically emit locating signals from a moving target. When it enters the range of one of the sensing motes, the information is relayed down the communication chain to the PC. Finally, we have demonstrated sensor information being displayed and continually updated on the Internet. While not a networking achievement in itself, it hints at the potential of incorporating existing telecommunications into sensor networking.

Existing multi-hop networks have used preprogrammed communication chains. While demonstrating the feasibility of networking the sensor motes, it does little to exploit the dynamic aspects that a distributed and adaptive sensor network can provide. The PC plays an important role in the system: while the sensor motes completely mediate communication among themselves, the additional computational power of the PC allows for simple observation of sensor data and keeps track of all messages in the network.

2.2 Current Plans

The innate beauty of having thousands of Smart Dust motes is that a system can continue to operate in the case of many and inevitable mote failures (or movement). To this end, communication paths must be dynamically updated in the network. The next goal is to construct a partial connection list of motes that can communicate with each other from an unknown topology, partial in the sense that a list of hops from the PC to any reachable mote in the network is determined. To construct this list, the PC selects a mote (the red circle in Figure 4) of unknown logical whereabouts and broadcasts a message to every mote in range asking for the destination mote. The message floods through the network accumulating hop information within the path field of the message as it propagates. Repeat message broadcast is avoided by a unique message ID assigned by the PC. Once the message reaches the desired destination n, it follows the hops back to the PC. A logical map from the PC to the destination and any intermediate motes can be added to the routing list. The number of such topological queries required to build the entire list is hence less than the number of motes in the network and the time required to find any single mote should grow only as the network diameter.

Having built the connection table for all destinations, the PC can now demand sensor information from any mote as though the channel was preprogrammed. Clearly, this simplistic model works perfectly in the dream world where messages are never lost or collide and nodes and connections are reliable and constant. If a node or connection fails, this will be revealed to the PC through nonresponse to an information request. If this occurs, the connectivity-establishment algorithm can be rerun. Alternation between connectivity and information demand allows the PC to monitor sensors at various positions in a dynamic network.

Another approach that can be taken is to provide some degree of autonomy to individual nodes in establishing connections with neighbors. In Figure 4, for example, mote 2 could learn that it is within communication range of mote 0, 3, and 4. Failure of one connection in the "routing table" could lead the mote to try another path. This issue and many related questions are studied extensively in the Internet and mobile networking literatures. Subsequent sections of this proposal also deal with the distributed establishment of connectivity and geographical location.

With the proof of concept of the connectivity/demand strategy, we will seek methods to improve the robustness of communication in the network. Issues to explore include studying the relevance of various well-known MAC standards in highly connected, low bit-rate networks. Studies are under way to improve the physical transmission and reception of RF messages with more effective antennae designs. Additionally,

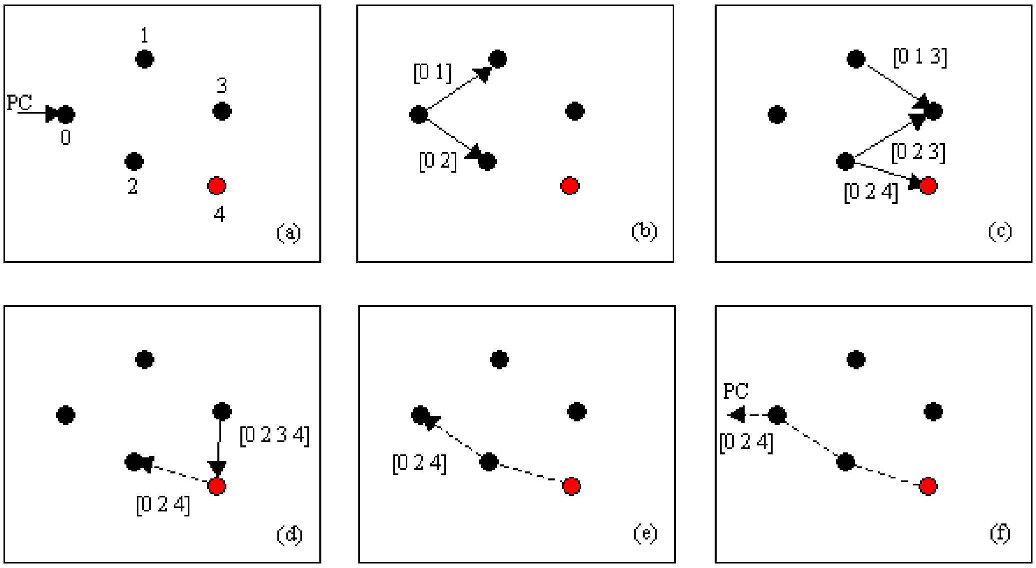

Figure 4: Example of the PC determining a path to mote 4 as defined in (a). Images (a)–(c) show forward flooding of the demand through the network. Bracketed lists show the path field of the propagating messages. Each mote responds only to the first message reaching it. (d)–(f) show how the path information is used to reverse-link the communication path along the dashed lines.

both current macroscopic and future miniature Smart Dust have a limited amount of stored energy, and power consumption is a prime concern. Ideally, the motes could be run at low duty cycles to conserve energy (current RF motes can survive for 2 years at a 1% duty cycle). Synchronization of the waking periods of the motes throughout the network is another interesting topic.

Without the PC in the exploration above, topological exploration is complicated. A network consisting only of identical motes confronts such issues as leadership, local storage of routing tables, and instability due to the unbridled creation of messages. The lucidity of the PC observer solution comes at the price of communication flexibility—in this regime the motes do not communicate with one another other than to relay a message to or from the observer. There is, however, no lack of application or research potential for our case. The popularization of Smart Dust has fomented sundry requests for sensor networks including a desk-environment monitoring system for improving office conditions and a box tracking network for use in an assembly plant.

The questions discussed thus far have been directed at networks of RF motes, but since efforts are made to preserve mote compatibility, there is no restriction against multimedia networks. For example, consider three RF mote networks isolated around the Bay Area. Now place a laser mote in communication range of each network, and in the line of sight of the other two laser motes. A PC connected to one mote can now monitor sensor data from each mote in each network; once the science is developed, arbitrarily complex multimedia connections are engineering problems. With the addition of the Internet to the mix, yet another level of eclectic networks is conceivable.

As pointed out earlier, the new generation of macromotes will be programmable by RF. The distinction between message and program is no longer necessary. With this new faculty, motes will not only be prompted for sensor information, but can also have their behavior altered by a distant PC. With this paradigm, the network can be viewed as a parallel computational entity as well. However, before such complex objectives are pursued, it is important to thoroughly explore and mathematically understand the tools for determining and using communication links.

3 Objectives: Application Scenarios

In this proposal, our primary focus is on the fundamental scientific issues and engineering processes that undergird several classes of distributed sensing problems. We aim at the development of robust analytical framework comprising of concepts, models, analyses, algorithms, architectures, metrics, and software that will enable solution of hard technical problems that are exemplified by the application scenarios given below.

1. **Decentralized Estimation of Single Moving Targets (Tank Tracker).** Given a random distribution of SmartDust motes in an area of operations, (initially chosen to be a square, and later with more complex geometry), a model of sensor sensitivity (say accelerometers or acoustic or magnetic) give scalable algorithms for the optimal trajectory estimation for a single moving object. Sensitivity of the sensors is initially specified in terms of a spatial (2 D) probability distributions of detection and false alarm. Efficiency of the algorithms will be evaluated both in terms of guarantees of overall detection accuracy subject to a given acceptable false probability threshold. Analyze the detection efficiency of the algorithms as a function of acceptable false alarm rate, and compute lower bounds on the required sensor density. We will develop partially centralized (groups of the motes communicate to mother nodes) and decentralized algorithms, where some of the motes act as relay motes possibly with greater communication capabilities. In the former case, algorithms for the establishment of the position of the motes relative to each other will need to be developed using ideas from multi-dimensional scaling.

2. **Decentralized Estimation of Multiple Moving Targets (Intelligent Mine Field).** Given a probabilistic distribution of SmartDust motes and models of sensor activity develop algorithms for the recovery of estimated motion and type of multiple moving targets. To be able to distinguish multiple targets the sensors on board need to perform some analysis and data fusion of multiple sensors before broadcasting their results. Efficient multi-sensor distributed data association algorithms will be developed for fusing the trajectory maps of multiple sensors, allowing accurate triangulation of targets. Both structured environments (such as roads and bridges) and unstructured environments will be considered.

3. **Minimum Complexity Sensing ("Star Trek" Locator).** It is of obvious interest to have the ability to communicate with, monitor the health and status of friendly forces and distinguish them from foes. This involves attaching SmartDust to friendly forces and tracking them with the help of a distributed Sensor Web on the battlefield. In this context, we will also study the use of minimum complexity sensing. We will quantify the tradeoffs between using complex high bandwidth, computationally intensive sensors (such as cameras, 2-D sensors, biochemical analysis sensors) with more rudimentary ones (such as magnetic, acoustic) for the purposes of recognition and possibly solider health monitoring. Complex sensors represent a tradeoff not only in the cost of the sensor, but also the power consumption caused by more sophisticated processing and communication requirements.

4. **Biochemical Hazard Monitoring (ChemBio Plume Detectors).** Diffusing gradients of hazardous biochemical contaminants represent a significant threat to forces. Tracking the origin and the wave front of propagation of these contaminants requires different decentralized algorithms than the kinds developed for target or people tracking, and involve algorithms for the decentralized solutions of inverse problems. Other applications include using temperature data to determine storm front propagation.

4 Approach: Technical Research Concentration Areas

The distributed sensing problems described above, while appealing in their simplicity and their obvious usefulness are beyond the reach of current analytical methods: they belong to the domain of an emergent field which may be called real-time decentralized information theory. This is a marriage of communication complexity in computer science with the information theory of noisy channels. The issue is how distributed agents, each privy to some information, can coordinate over *noisy* channels to achieve the overall computational goal as efficiently as possible. The distinguishing feature is the need for *adaptive* response of each sensor in *real-time* to the information it gets from other sensors. One cannot fully exploit the law of large numbers as one does in information theory, and the challenge is to see how well one can do in the face of *real-time*, and *energy constraints*.

1. **Self Organization of Distributed Sensor Networks** (Anantharam, Chandra, Jordan, Pister, and Sastry).

   When SmartDust is scattered throughout a natural or urban environment, the spatial distribution and communication connectivity of the sensors will be initially unknown and will depend on both the scattering process and the physical structure of the environment. The first task is for each sensor to establish local (and hence global, via routing algorithms) communication connectivity with a set of neighboring sensors. Research will address methods that are communication-efficient and robust with respect to transient and permanent failures of nodes and links.

   A second task is geographical localization of sensors. GPS-equipped nodes can do this easily unless GPS is unavailable for reasons of stealth, foliage, etc. In such cases, or if sensor design precludes GPS, localization must be done in a cooperative, distributed fashion. We will develop localization algorithms that use qualitative nearest-neighbor information, using ideas from multidimensional scaling, and will extend these to incorporate noisy metric data from nearest-neighbor distance estimates and orientation estimates from magnetic sensors. Trade-offs involving the use of some GPS-equipped *mother nodes* along with non-GPS nodes will be studied.

2. **Application of Vision to Sensorweb** (Malik, Pister, Sastry). One can conceptualize distributed sensor networks as *random* retinas, i.e., ones in which the location of the sensels (sensory pixels) where flow information is received are random. Many powerful techniques from computer vision can be brought to bear on this problem, with interesting challenges in guaranteeing performance in the presence of the traditional signal noise as well as new sources of noise in position and information flow. Because many vision algorithms use local computation and communication, they are ideally suited for adaptation to decentralization using Sensorwebs. Edge detection, feature tracking, flow and motion recovery, and image segmentation are simple examples of problems which have direct analogues between vision and Sensorweb.

3. **Probabilistic Inference for Multi-Modal Sensing** (Jordan, Russell). We propose to develop a computational architecture that explicitly addresses the multi-scale, multiply-reliable nature of information sources, grounding in probability theory the interpretation of sensor data at levels ranging from individual target trajectories and attributes to global hypotheses about large-scale maneuvers or ChemBio weapon deployment. Building on recent results on efficient approximate inference in graphical models and for multisensor, multitarget data association and fusion, we will address problems of distributed computation and communication, flexible topologies, and continuous online learning and redistribution of sensor and target models. The algorithmic approaches will combine ideas from randomized methods (MCMC), variational techniques, and robust estimation. Within the overall architecture, connections will be made from situation assessment to decision making, using recently developed hierarchical reinforcement learning methods. The research will unify disparate approaches from AI, statistics, and control theory. We will provide algorithms, convergence results, error analyses, robustness results and prototype tools.

4. **Distributed Shannon Theory for Sensor Fusion with Energy Constraints** (Anantharam, El Ghaoui, Jordan, Sastry). Sensor fusion is deceptively subtle; classical results state that even fusing information from identical sensors depends in a delicate way on what and when they communicate to a centralized unit. Thus, a centralized unit can adaptively apprise the field agents of how they should send information, and strike tradeoffs between the accuracy of this information sent and the time it takes to send it. This tradeoff is achieved between the convergence rate of the algorithms running at the field agents (how many iterations should we run) and an data compression issues at these agents (more coarsely compressed information can be transmitted faster). Further the question of what kind of protection to provide for missing data is a coding theory problem; the more resilient the codes that are used are, the longer it takes to transmit the information. Since channels between the field agents

and a mother node are vulnerable to being compromised by an adversary, issues of game theory become important here. In addition, the communication will typically be wireless so the special features of multiaccess communication in a wireless environment, such as a-priori architecting to facilitate data fusion become particularly relevant in this context.

5. **Robustness and Graceful Degradation of Inferencing in SensorWebs** (Elghaoui, Malik, Pister, Russell, Sastry) Uncertainty issues are expected to come in at various levels in sensor webs. For example, the precise location of the motes will be largely unknown. However, it is reasonable to assume that the motes will be known to lie in some given area (say, a more or less circular area around a nominal point). Another example arises in wireless communications systems, in which the direction-of-arrival matrix is not perfectly known. The robustness of a decision-making process relates to its ability to withstand uncertainty, modeled as deterministic and unknown-but-bounded. An extensive body of work on robust optimization is highlighted by a general framework and efficient robust decision algorithms, mostly based on convex optimization relaxations developed by El Ghaoui and co-workers. We plan to develop the research effort in this area along two lines. One is to increase the ability of the distributed sensing and communications algorithms to withstand uncertainty at several levels, with graceful degradation of performance. A second objective is to assess the impact of approximations on the actual worst-case values, depending on problem size.

6. **Performance Metrics and Assessment of Algorithms** (Anantharam, Chandra, Elghaoui, Jordan, Malik, Pister, Russell, Sastry) Possibly the most fundamental problem to be addressed in the theory underlying Sensorweb is a grammar or language for describing the tradeoffs between computation, communication, and energy usage. Optimistic numbers for energy consumption in RF communication over short distances are in the ten nano-Joule/bit range. For modest quality CMOS, this corresponds to thousands of multiply/accumulate operations on 8 bit data. We will form an abstraction of the energy constraints of a given electronics process and combine that with an abstraction of the tradeoffs in computation versus transmission for a particular class of algorithms. To provide a specific "ground truth" for the grammars, algorithms, and software developed, we will implement each of the four application scenarios on a real wireless sensor network with hundreds of sensor motes distributed in urban or natural terrain, according to the requirements of the scenario. Concrete metrics can then be measured: Was the tank detected? What was the average error in the reconstructed shape of the anthrax plume? How many minutes or years of battery life do we have with this particular algorithm?

7. **Emergent Behavior in Decentralized Sensor Networks (from local to global)** (Chandra, Pister, Russell, Sastry). The Sensor Web will have sufficiently many (order of thousands) components to bring it within the regime of asymptotic analyses. Research in distributed data fusion has shown the value of adopting a macroscopic approach to systems consisting of a large number of sensors. For instance, it is known that the sometimes only a small fixed number of different data representation rules need to be used by the sensors and the problem of optimal data fusion can be transformed into that of deciding what fraction of the individual sensors should use each rule. From a computational point of view, this converts the problem from one with size proportional to the number of sensors into one of a fixed size, independent of the number of sensors.

   There is a continuum of design choices for system decomposition, ranging from strict hierarchical control to a fully distributed multi-agent system. Furthermore, different choices may be appropriate at different levels of abstraction, ranging from the (typically continuous-domain) low-level control systems concerned with safety and smooth execution to the (typically symbolic/discrete) strategic levels concerned with optimization and planning for high-level goals. Centralized computation will not work. Computation must be distributed, and in fact the computation performed will depend on the spatial and temporal context of the query and location of the node. We will explore a spectrum of hierarchies: from the comparative anarchy of a cellular automata approach to the flexible but regimented adaptive multi-scale promotion approaches used by insects.

5 Detailed Research Objectives

5.1 Self Organization of Distributed Sensor Networks

A number of important issues arise when we consider—in the absence of on-board GPS—algorithms for establishing and maintaining relationships between (network) topology, (environmental) topography, and geometric representations. Different applications will have different needs in this regard, some requiring only local scalar signals, others requiring notions of differentiability (gradients), others requiring locally Euclidean neighborhoods, and others requiring an embedding space. It is useful to try to solve these problems as an integrated package, providing an upgrade path for increasingly sophisticated applications. It is also important to recognize at the outset of the design of such a package the implications of the stochastic, low-bandwidth environment in which the network nodes will operate. There are a number of classical techniques that provide initial directions for research. Broadly speaking, these fall into two classes of algorithms: (1) "nonmetric multidimensional scaling" algorithms and (2) self-organizing topographic map algorithms. Let us briefly discuss some of the issues that arise in trying to adapt these methods to our problems.

Nonmetric multidimensional scaling (NMDS, [1]) solves the following problem. Let us suppose that we have N nodes that we wish to embed as points in a d-dimensional space. We are given a *similarity matrix*; that is, a real value for each pair of nodes characterizing the similarity (equivalently, inverse dissimilarity) between the nodes. In our SensorWeb domain there are a variety of such signals that could provide an appropriate similarity measure, such as the strength of received RF signals at nodes within the network as transmitted by other nodes, or the relative strength of extra-network signals. We imagine that there is a monotone function linking the similarity measure to a metric (a distance) for the space in which we wish to embed the points. We do not need to know this monotone function a priori; it is an output of the algorithm. Given the matrix of similarities as input, NMDS produces coordinates for each of the nodes in a Euclidean space. This can be done with or without knowledge of the dimension d. In our case d can be assumed known and fixed.

There are a variety of algorithms for NMDS; the most popular pose the problem as a nonlinear optimization problem which is solved by conjugate gradient methods. In the setting in which all of the nodes communicate back to a mother PC, offline conjugate gradient methods may provide a workable solution to our basic embedding problem. Using node IDs, the pairwise similarities can be uplinked and coordinates can be returned. Note that NMDS solutions are routinely found for problems involving dozens to hundreds of nodes. For larger-scale problems, one can envision multiresolution versions of NMDS.

Such solutions, however, are less satisfying in dynamic domains, in which node movement may rapidly render fixed coordinates useless, and also in static domains in which a distributed solution is desired. In these cases, one can consider local algorithms that iteratively adjust coordinates based on local "stress" of the triangle inequality. Such algorithms are studied in the literature on "self-organizing topographic maps" (SOTM; [2]). In this setting, each node is associated with a Gaussian model for a local region of space, in which the means (and covariance matrices) are unknown a priori. Moreover, the nodes are assumed to be embedded in a d-dimensional space, where once again the coordinates are unknown. The overall goal is to produce a density model of sources in the domain, and to embed the density model in a Euclidean space.

In this setting we can imagine that each node serves as a "data point" for the other nodes, with its current coordinates (the mean of its Gaussian) serving as the location of the data point. Each node broadcasts its coordinates, with both the strength of the received signal as well as the coordinates themselves being used by the algorithm at the receivers. Essentially, the receivers move their own coordinates toward the coordinates of the transmitter, with the step size weighted by a function of signal strength. Moreover, coordinates are pushed apart when this serves to relieve "stress" in fitting a local monotone function to the distance-strength function (i.e., each node keeps a cache of its neighbors current coordinates and solves a local NMDS problem).

While this algorithm provides only a set of consistency equations for node coordinates, we can also provide boundary conditions (e.g., fixed coordinates for nodes on the boundary) to calibrate the algorithm. Similarly we can rely on an off-line NMDS algorithm associated with a mother PC to provide an initial solution which can be updated locally.

If the density of nodes is low the consistency equations may not provide sufficient constraint on a solution. In such cases a complementary approach would involve using allowing naturally occurring signals as "data points." That is, a naturally occuring signal in the environment provides an implicit constraint that each node

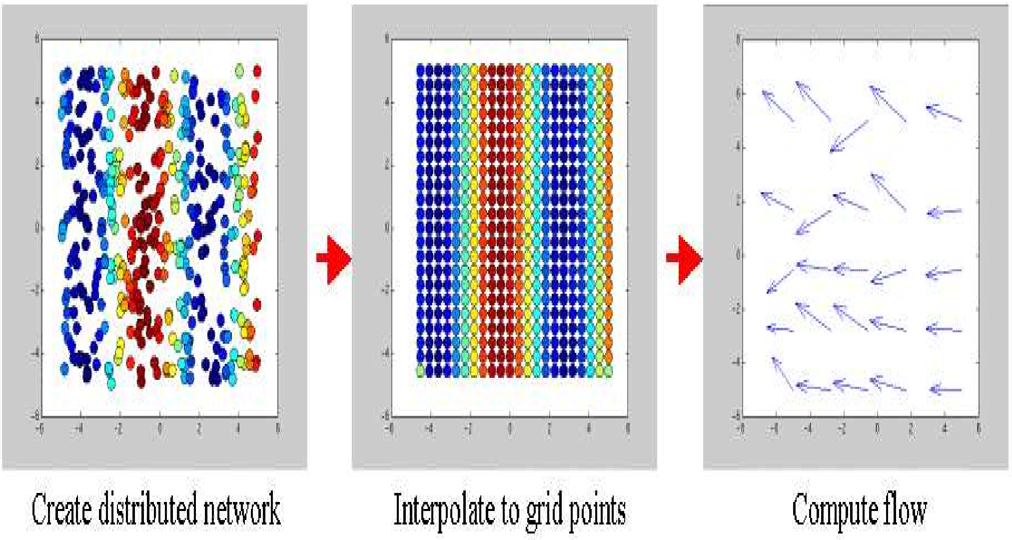

Figure 5: This figure illustrates the steps in flow recovery for a right-to-left horizontally moving flow field

registering the signal should have similar Euclidean coordinates. By broadcasting coordinates and averaging (weighted this time by the strength of the naturally occurring signal) this constraint is embedded into the algorithm. Both the NMDS and SOTM algorithms have natural robustness properties; the former because of the weak, non-metric assumptions on which it is based, and the latter because of it stochastic model. Coupled with the relatively favorable computational complexities of these algorithms and the possibilities for parallel, distributed implementation, we feel that these algorithms provide a reasonable starting point for our research effort.

5.2 Application of Vision to Sensorweb

Data collected from Smart Dust can be used to determine atmospheric occurrences or other large-scale motion patterns. For example, periodically sampling temperature data from distributed dust motes may be useful in determining the direction in which a storm front is moving. With this example in mind, it is clear that recovery of the temperature flow from a distributed network of dust motes could be useful in an area that may not lend itself to conventional means of data collection. In the computer vision community, optical flow recovery algorithms have been well characterized, as discussed by Simoncelli [3]. The general formulation of gradient-based optical flow algorithms relies on spatial and temporal differentiation of a regular arrangement of pixel brightness values, otherwise known as an image. Implementing optical flow algorithms for randomly distributed networks can be achieved by interpolating mote data into a fully populated, regular grid. In this scenario, the temperature values in the grid correspond directly to pixel brightness in an image. The primary factor considered in determining the accuracy of a flow recovery is number of nodes in a given area, or node density. Also to be considered are interpolation techniques and derivative calculation techniques.

Our development of flow recovery for a distributed network (Figure 5) begins by simulating a flow field over a set of two-dimensional, randomly distributed nodes. The flow field is generated by a continuous function of time that varies two-dimensionally in space. A horizontal flow may be represented by the function $\cos(x+t)$, where x is a position along the horizontal axis and t represents time. Each node in a randomly distributed network is then assigned a value based on its position and the current value of t. Next, a fully populated grid is generated by interpolating the values in the distributed network. The final step in recovering flow is implementing the Lucas-Kanade [4] gradient-based optical flow algorithm.

The fineness of the flow estimation (i.e., number of flow vectors recovered for the 10x10 area) is determined by the number of interpolated grid points and the window size used in the linear least squares estimate of the flow field. These parameters must be chosen carefully to avoid temporal aliasing [3]. Here,

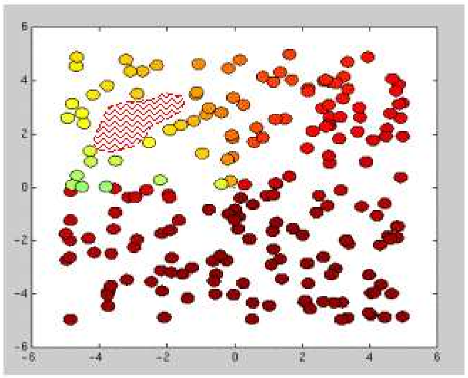

Figure 6: Distributed network of 200 nodes withn rotating flow field

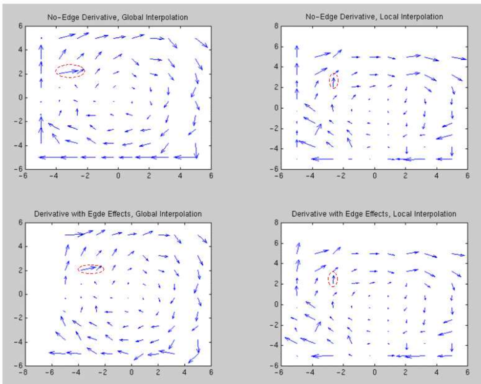

Figure 7: Flows recovered using 4 variants of a flow recovery algorithm.

they have been determined by trial and error for each flow field function.

Within a 10x10-unit area, a specified number of nodes were distributed randomly, and the flow was computed. Figure 6 shows an example of a distributed network of 200 nodes, with colors representing increasing temperatures from blue to red, determined by a rotating flow function. Figure 7 contains the corresponding flow determined by four variations of the flow recovery algorithm. By comparing the flow field with the distributed network, it is clear that errors in the flow field occur on or near areas of sparse population in the network, as indicated by highlights in the figures. Another observation is that the global interpolation method generates a smoother flow field with fewer drastic changes between adjacent flow vectors than in the local interpolation case. The addition of edge effects to the derivative filter does not appear to significantly enhance algorithm performance. Further analysis of flow recovery was performed for a purely horizontal flow function. Although motion in only one direction is a poorly conditioned case for the optical flow algorithm, the inherent inaccuracies of interpolating data from a distributed network allow the recovery of near horizontal motion. Since the desired flow vectors have no y component and a constant x component, the error function, V_y/V_x, is used to evaluate the flow recovery. A plot of the number of nodes present in the 10x10 area versus V_y/V_x indicates that increasing the number of nodes from 10 to 100 by incrementally adding nodes to the original distribution (Figure 8, part a) results in a decrease in the error function. Although this plot also seems to suggest that local interpolation with edge effects results in the smallest error, visual observations reveal that while V_y/V_x is small, many of the vectors face the wrong direction. In general this case does not perform better than the others do. Increasing the number of nodes incrementally from 100 to 1000 (Figure 8, part b) reveals that the global interpolation with no edge effects generally recovers flow most effectively; however, this is largely dependent on the specific distribution of nodes. Above approximately 100 nodes, flow estimations cease to improve, suggesting that a finite number of nodes is sufficient to obtain optimal flow vectors for the given function. Additionally, the grid size (24×24) is small

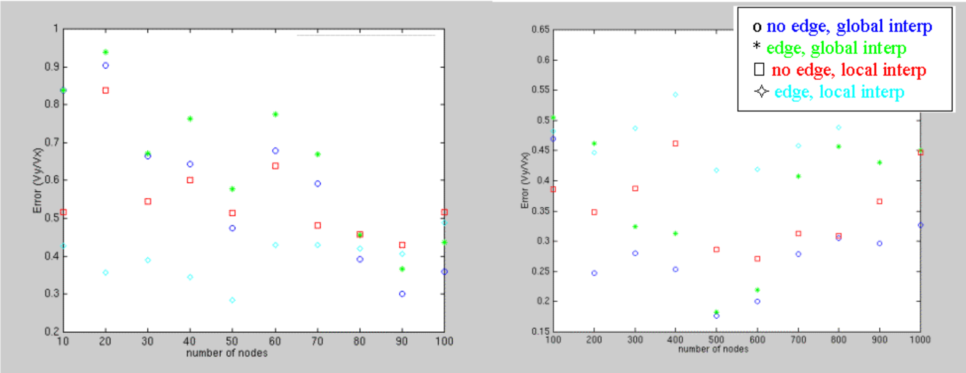

Figure 8: This figure illustrates the mean flow error versus number of nodes: a) 10-100 nodes and b) 100-1000 nodes.

enough that edge effects impact the error function. For a much larger grid, edge effects can be neglected. Other general observations include an improvement in flow estimates when averaging flow vectors over several distributions of nodes. Performance was not improved, however, by averaging flow vectors over multiple sets of time sequences for the same distribution. These two results confirm what was stated previously: sparsely populated areas in the network induce errors. However, when nodes are redistributed and flow results are averaged, these errors are minimized. Another important result is that the flow recovery algorithm breaks down for the case of a small moving object in the flow field. This is similar to the case of a white square moving on a black background in computer vision. For this type of motion recovery, an EM-based algorithm is discussed next. Improvements for future implementations of flow recovery include the use of multiscale algorithms to avoid temporal aliasing, such as that described by Simoncelli [5]. Also, a weighting function could be implemented which minimizes the weight given to flow estimates near sparse areas in the network.

5.3 Probabilistic Inference for Multi-Modal Sensing

The overall problem of interpreting and combinng the sensor data from a Sensorweb can be divided into three levels:

- *Low-level sensor interpretation*, in which raw sensor data at each sensor is processed to extract meaningful features. For example, data from an acoustic sensor might be processed to extract several engine signals, each of which is described in terms of base frequency, loudness, and direction.
- *Tracking and data association*, in which individual objects are tracked by combining and individuating sensor observations at single and multiple sensors.
- *High-level situation assessment*, in which descriptions are generated involving complex behaviors of multiple objects—for example, a group of soldiers digging in with artillery, or civilians evacuating a village.

The first of these three items involves the application of a wide variety of sensor-specific signal processing methods, which we shall not address here except to say that our group has very substantial expertise in this area [?, 6, 7, 8]. We do not propose to conduct research on specific algorithms for local sensor interpretation under this proposal. [[is this correct?]] The second and third items aare dealt with in the following subsections.

5.3.1 Tracking and data association as probabilistic inference

One of the primary tasks of the Sensorweb will be to monitor objects residing in moving through the environment—tanks, truck convoys, individual soldiers, civilians, etc. Traditionally, this problem has been divided as follows (see, for example, [9]):

- *Tracking*: Given a sequence of observations of an object, estimate and predict its trajectory and other properties.
- *Data association*: Given a sequence of observations of multiple objects, associate observations with distinct objects.

Clearly, these two tasks cannot be handled separately, because tracking requires identification of observations associated with a particular object and data association requires track predictions in order to decide which observation belong with which track. Our approach will be to model and solve these problems as a single probabilistic inference problem for which we can derive approximately optimal solutions and for which we expect to be able to design completely distributed and robust algorithms.

Our research will distinguish two settings: the *sparse* sensor setting, where the detection radius of each sensor is much less than the nearest-neighbor distance; and the *dense* setting, in which the detection radius is less than or comparable to the nearest-neighbor distance, and hence the detection zones overlap. In the former case, which might occur with laser communication bewteen sensors, the problem is greatly simplified because each object is in at most one detection zone at one time and usually at most one object is in any one detection zone. In the latter case, each sensor may simultaneously detect several objects, some of which may also be detected by other sensors. This gives rise to substantial data association difficulties, but allows for methods such as triangulation that can give very accurate tracks. We propose novel and efficient approximation algorithms for this case, including the possiblility of full decentralization.

Sparse detector setting The sparse detector setting is effective when long-distance communication among nodes is feasible and wide-area detection coverage is needed. We have conducted a pilot study of a simplified version of the problem: in the first instance, to estimate the position and heading of vehicle with straight-line motion. Our aim is to understand the sensitivity of network performance to changes in three independent parameters: the distribution of nodes, the number of nodes in network, and the sensing radius of each node. We derive qualitative insights that lead to more extensive and realistic investigations and to directions for theoretical development.

Our simulations assume that a vehicle "activates" any sensor if it passes within a fixed detection radius. Trajectories are inferred from the set of activated sensors (Figure 9). To determine the line trajectories, the Expectation Maximization (EM) algorithm, as outlined by Weiss [10], is implemented with a modification to use the perpendicular distance from the line rather than the vertical distance for the residue function. This gives a simple fit of the observations to the most likely trajectory.

Our experiments used a network of 200 nodes with sensing radii of 2.0 units. We found that the detection reliability and track accuracy was approximately independent of the spatial distribution of sensors—we tried distributions that were purely random, random within fixed bucket regions, restricted to a perimeter, and in a fixed grid. This suggests that a random distribution is quite reasonable as a deployment mechanism, and our subsequent experiments were limited to this case.

The next set of experiments measured the detection reliability and track accuracy as a function of the detection radius of the sensors and the number of sensors in a fixed-size region. Figure 10(left) shows that detection reliability increases with the number of sensors and with detection radius; the right-hand graph shows that the track heading estimate was reliable in all cases, but that above a certain optimal radius the error increases because the larger-radius sensors prodcue more uncertainty as to the actual location.

The approach above can be extended to more than one track by first attempting to fit the data to one straight line, then two, and so on. The loop terminates when the addition of another line to the data does not significantly reduce the overall residues. The algorithm can successfully resolve most cases with two trajectories, particularly when the trajectories are far from parallel.

The EM approach is unlikely to scale up without the additional use of timing data to resolve multiple trajectories, and will also be confused when multiple trajectories intersect a single sensor. These problems are addressed in the next section. The approach may be able to handle curved trajectories quite easily, at least mathematically, since the EM algorithm can simply be applied to a larger parameterization of the trajectories.

An obvious direction for theoretical research is to analyze system performance to determine the optimal sensor density and sensor radius subject to fixed total cost. This would appear to be a fairly simple problem

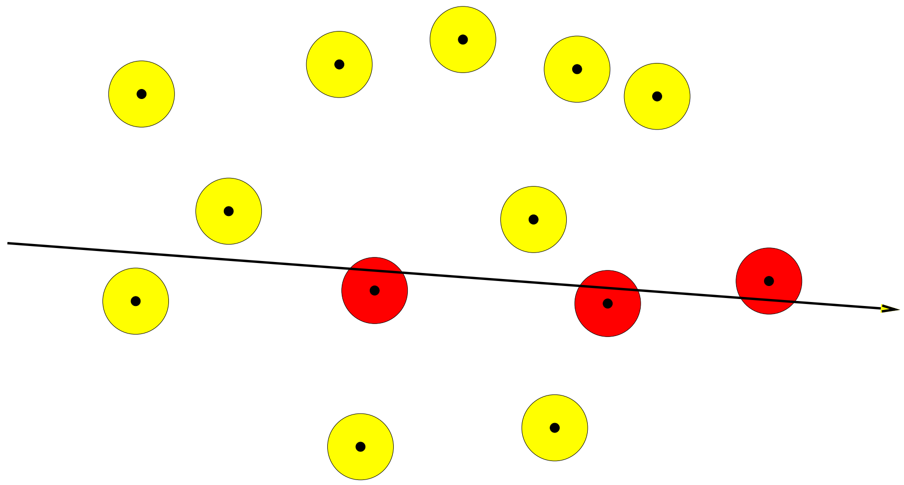

Figure 9: Detection of a trajectory from the set of activated sensors.

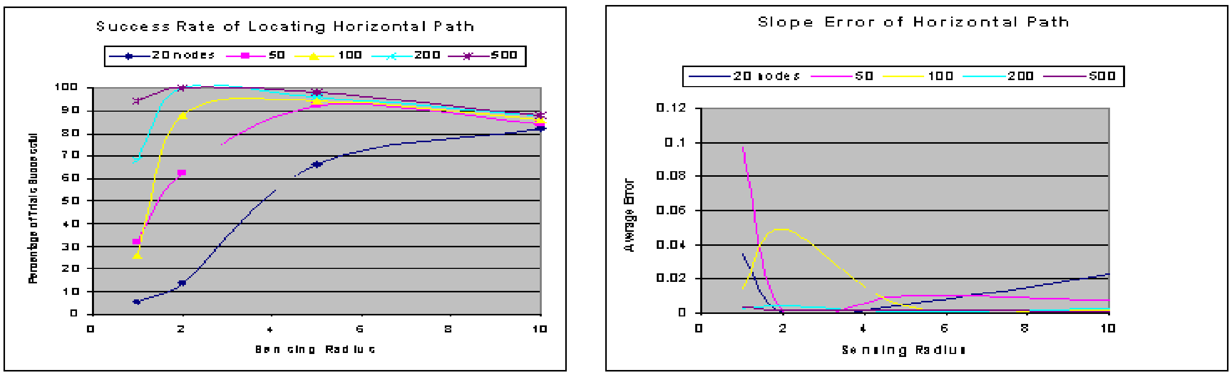

Figure 10: Detection reliability (left) and heading error (right) as a function of detection radius, for various numbers of sensors.

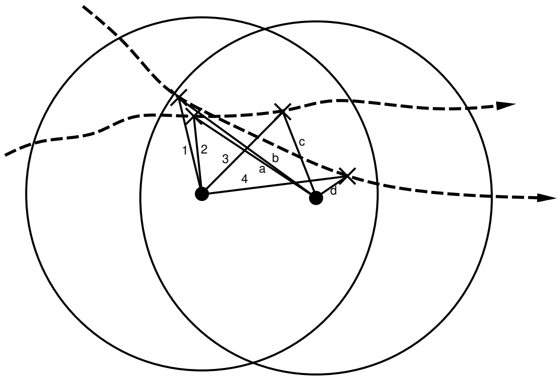

Figure 11: Multiple tracks and multiple sensors: vehicles are detected by the first sensor at headings 1, 2, 3, 4 and by the second sensor at headings a, b, c, d. There are 24 possible assignments; the correct assignment is $\{(1,b),(2,a),(3,c),(4,d)\}$. Given the correct assignments, triangulation identifies the location for each reading. Additional assignment of locations in sequence gives the tracks.

in computaitonal geometry. It becomes more complex when we deal with more realistic models of detectors, where detection probability falls off with distance.

Dense detector setting Our current plan for Smart Dust motes is to use RF communication with a nearest-neighbor range of about 20m. Sensors such as acoustic and vibration sensors, which are most useful for detecting vehicle and personnel movements, may have significantly higher range. Therefore, the dense detector setting is likely to be pervasive. As mentioned above, the dense setting requires solving the data association problem.

Formally, the data association problem can be viewed as the problem of computing posterior probabilities for specific propositions ϕ given observations $\emptyset$. For example, suppose that a number of vehicles are observed using acoustic sensors that have reasonable directional resolution but almost no absolute range capability. If we can decide which objects are which, then the directions obtained by each sensor for a given vehicle can be combined using triangulation to obtain the position of the vehicle, and hence a reliable track (Figure 11).

An *assignment* is, in general, a grouping of observations into equivalence classes, each of which is associated with a single object. Given an assignment, the problem reduces to n separate tracking problems for n objects; these will typically be solved by Kalman filtering and variants thereof [9] for spatial motion tracking, or by dynamic Bayesian networks [11] and improvements thereon for the more general case (see next section).

The difficulty, of course, is that assignments cannot be observed. The probability of each assignment ω given the observations, $P(\omega|\emptyset)$, can usually be computed efficiently. Moreover, given an assignment (and hence a set of tracks), most questions of interest can be answered easily. The answer to any particular question ϕ—e.g., how many tanks are heading west through this sector—is given by *summing over* assignments:

$$P(\phi|\emptyset) = \alpha \sum_{\omega} P(\phi|\omega) P(\omega|\emptyset)$$

But with n objects and m overlapping sensors, the number of assignments grows exponentially in both m and n. It can be shown that exact solution of the data association problem is computationally infeasible (#P-hard).

In recent work [12], we have developed a provably polynomial-time approximation algorithm for this problem, based on the use of Markov chain Monte Carlo inference (MCMC). An MCMC algorithm generates samples from the set of assignments; it can be shown [13] that our particular MCMC formulation converges rapidly to give a guaranteed good approximation to the true solution. We have demonstrated the effectivenes of the approach in practice for freeway traffic monitoring with dozens of sensors and thousands of vehicles—far larger than any previous tracking system could handle. The computational load grows roughly linearly with network size, and the sampling computation is extremely simple and flexible.

The MCMC approach also allows for a very straightforward application of online EM, combining ideas of [14, 15]. Using this technique, the system can adaptively identify models of sensor behavior and object behavior without the need for calibration and is robust with respect to changing conditions (weather, lighting, object population, etc.).

In the Sensorweb setting, the variety and characteristics of the available sensors and the richness of the possible object behaviors motivate a rich research agenda of domain-specific problems—for example, fusion of multimodal data (e.g., directional acoustic and scalar vibration). We will also investigate algorithmic improvements inthe MCMC scheme, such as *adaptive directed sampling* to improve the rate of convergence by biasing the sampling process towards *prima facie* likely assignments.

The most interesting problem will be to arrange for the MCMC computations to be distributed over the network in such a way as to minimize communication and local computational loads. The process of generating each sample in the MCMC process involves exchanging pairs of observations between equivalence classes in an assignment; if this can be done locally, where each sensor node is responsible for that portion of the complete assignment involving its own readings, then the entire computation can be distributed over the Sensorweb.

5.3.2 High-level situation assessment: representation, inference, and learning with uncertain knowledge

Recent developments in the area of *graphical models* are having significant impact in many areas of computer science, engineering and applied mathematics. Graphical models, referred to in various guises as "Bayesian networks," "Markov random fields," "influence diagrams," "decision networks," or "structured stochastic systems," are an elegant marriage of graph theory, probability theory and decision theory. They provide a system-theoretic, computational framework in which to do statistics and decision theory, justifying and unifying many classical techniques for inference, learning, and decision-making, and providing a firm foundation on which to design new systems.

In AI, Bayesian networks (directed graphical models) have become the method of choice for designing systems that cope with uncertainty. Bayesian networks have long provided the core technology for inductive probabilistic inference. More recently, Bayesian networks and influence diagrams (Bayesian networks augmented with decision variables and utilities) are being used in research on planning and learning. In particular, Markov decision processes (MDP's) provide the banner under which much recent work on planning and reinforcement learning is being carried out. Many unsupervised learning systems are graphical models and much recent work in machine learning on supervised learning, model selection and model averaging is being pursued within the graphical framework

While acknowledging that graphical models represent a significant success story, it is important to acknowledge as well that much remains to be done before graphical models can be viewed as a full-fledged engineering discipline, capable of providing robust, systematic solutions to large-scale problems in inference, learning and decision-making. In our view there are three problems that stand out:

- Graphical models need more expressive symbolic representational capabilities. This includes the ability to represent*entities* and *relations* within an overall probabilistic framework.
- Graphical models require a suite of *approximate inference methods* with (1) good scaling properties, (2) guaranteed convergence properties, and (3) error analyses.
- Graphical models need *robust analysis and design procedures* that provide performance guarantees.

We plan to make significant progress on all three of these topics.

Many applications of graphical model technology involve dynamical systems, where the probabilistic inference problem becomes the classical "filtering" problem of calculating optimal state estimates based on past and current data, as well as the related "smoothing" problem in which past estimates are updated based on new data.

If we utilize a Bayesian network to model a single probabilistic state transition—a single "time slice"—and repeat this basic structure for subsequent time slices, we obtain a graphical model known as a "dynamic Bayesian network (DBN)" (Xrefs). DBN's are generalizations of classical state-space methods such as hidden Markov models (HMM's) and Kalman filters. While these latter methods assume an unstructured state transition matrix, DBN's allow structured dynamics to be represented and exploited during probabilistic inference.

The class of structured dynamical systems represented by DBN's include the class of "coupled chains" or "multiple models" that are often studied in sensor fusion problems ([?, ?]). Indeed DBN's provide a general setting for dealing with the fusion of multiple sensory channels over time. The ability to manipulate structured dynamical systems in crucial in scaling up to large filtering and smoothing problems. In recent work, we have been able to perform probabilistic inference in DBN's involving thousands of states and millions of parameters (Xrefs) and to do so in near real-time.

5.4 Distributed Shannon Theory for Sensor Fusion with Energy Constraints

The dust motes need to coordinate as well as communicate to achieve their surveillance or detection objective. The architecture chosen for coordination will influence the efficiency of the communication (power efficiency versus accuracy) and the communication protocols will inform the choice of the coordination architecture. This is a new kind of problem in information theory on several grounds :

1. Traditional information theory treats the coordination architecture as fixed and discusses only communication efficiency. Here the two aspects of the problem have to be viewed together: the results of the communication may require that one adapts the architecture, and the choice of architecture may require the communication scheme to adapt to it.

2. Traditional information theory appeals to law of large numbers, allowing for substantial delays to improve communication efficiencies. In coordination scenarios, the response of a sensor will depend on what it has heard from other sensors and allowing for too much delay in this will substantially degrade performance, in addition to being sometimes infeasible.

We propose to address this new class of problems drawing on several new technical developments that have appeared in the recent literature. To flesh this out, we briefly discuss some of these results. Suppose the input to a function f is split between two processors connected by noiseless binary channels. The *communication complexity* of f measures the number of bits they must exchange to compute f. Since its introduction by Yao ([16]), this interactive model of computation has been widely studied, with the objective of characterizing the inherent cost of communication in distributed computation. For a survey of this work, see Orlitsky and El Gamal ([17]). Since all practical channels are noisy, it is of interest to study the effect of channel noise on the complexity and reliability of this communication. In this direction, L.J. Schulman ([18]) recently proved that: any noiseless-channel deterministic protocol of complexity L that computes f correctly can be deterministically simulated over noisy channels with $O(L)$ transmissions, while incurring error probability $2^{-\Omega(L)}$ in computing f.

The proof of Schulman develops a class of codes, called *tree codes* which have the remarkable global error correcting properties : essentially, as the communication (in order to compute) progresses even errors in the distant past can be corrected based on information just gained. Codes such as these are likely to play a significant role in extracting communication effiency from smart dust networks. We plan to investigate and develop this mathematical tool in the context of the Berkeley testbed.

Smart dust networks are more complicated than just two nodes interacting. In recent work we have developed a new kind of "type theory" for handling interactive communication via networks of processors. This technique is developed in Venkatesan and Anantharam, [19], [20]. We believe that this technique will be of significant importance in bringing information theoretic ideas to bear on coordination in multisensor smart dust networks. In the cited papers we prove remarkable results on the common randomness capacity of networks of processors using this technique : in particular, we identify this capacity precisely as the solution to a simple finite dimensional optimization problem where randomness "flows" across a network with capacity constraints given by the Shannon capacities of the noisy links connecting the individual processors.

Other remarkable recent progress that we propose to exploit is based on work by Borkar, Mitter, and Tatikonda [21] and by Borkar and Mitter [22]. In this work the avergae cost control problem for linear stochastic systems with Gaussian noise and quadratic cost is studied when the control actions have to be based on quantized observations, and a form of the separation principle is derived. The ideas used in this work are likely to lead to significant insights into the design of smart dust networks, where the communication must necessarily be simple (using few symbols), and estimating the intended messages based on the received symbols is an essential part of the overall problem.

5.5 Robustness and Graceful Degradation of Inferencing in SensorWebs

In a sensor web, the data is largely uncertain. In some situations, it is possible to assume a known distribution of the data. Quite often, though, this distribution is simply not known, or very poorly. For example, the location of the motes, if they are distributed by plane, depends highly on local meteorological conditions and the resulting distribution is highly uncertain. Uncertainty can be inherent to the system, but can also be used as a way to address complex relationships, using a technique known as embedding. For example, one might address quantization issues by embedding quantization errors in a larger class of unknown-but-bounded errors. The uncertainty problem impacts on the choice of algorithm for distributed sensing. An algorithm that is very efficient in the case that the location of the motes is preciserly known should exhibit a high sensitivity to sensor location, for the very reason that it is highly tuned to all available information, including the motes' location.

Robust optimization addresses problems where the data is subject to nonlinear, unknown-but-bounded errors. A robust solution is one which remains acceptable (*i.e.*, achieves the desired specifications) despite the errors, and minimizes the worst-case value of the cost function. Recent progress has been made in the field, with work from Ben-Tal Nemirovskii, El Ghaoui and co-authors [23, 24, 25], resulting in methods to solve these problems via convex optimization techniques. These techniques are particularly fast and efficient if the "nominal" problem (*i.e.* with no uncertainty) is convex.

Our goal is to achieve a graceful trade-off between robustness and performance in various aspects of distributed sensing, with robust optimization algorithms that require low computational effort. In addition, we expect to come up with performance metrics and *a priori* assessments of these algorithms (see section 5.7). In the direction of evaluating the conservativeness of the bounds used, as a function of problem size, a preliminary set of results has been obtained by Ben-Tal and Nemirovskii [25]. The contribution of robust optimization to the global research effort should be embedded within the other lines of research. The following examples illustrate this interaction.

Worst-case simulation and inferencing for complex uncertain systems. The robust optimization approach can be used in the context of worst-case simulation and estimation of complex, uncertain dynamical systems [26]. The basic premise is to encapsulate the entire knowledge about the state of the system in a "confidence ellipsoid" that is guaranteed to contain the state vector. The problem of *worst-case simulation* is to compute an approximation of the state of a system, given some (uncertain) initial condition, and an (uncertain) dynamical model that describes its behavior. The problem of *estimation* is formulated as one of computing ellipsoids of confidence for the state of an uncertain dynamical system, given uncertain measurements. This allows to generalize Kalman filtering to cases when the noise enters corrupts the data (the model) in a nonlinear, complex fashion. The tank tracker problem falls in this class of problems.

Robust optimization and approximate inference. The similarity between variational inference approaches (based on Fenchel relaxations and convex optimization) and robust decision techniques (based on similar convex analysis tools) calls for a close study of their mutual relationship. A first perspective pertains to the *robustness of approximate inference.* The techniques devised for variational inference lead to convex optimization algorithms, for which the robustification procedures could be applied. This approach would lead to upper bounds on quantities of interest (likelihoods, posterior probabilities, etc) that are valid *despite* uncertainty (for example, interval covariance matrices in a Gaussian channel) and yield results on time/accuracy trade-offs in the context of robust estimation. The problem of approximate inference is currently addressed via robust decision techniques, mostly in the form of ellipsoidal (set-membership) techniques for identification and estimation of models with error structure.

Robust power allocation in wireless communication. This problem is concerned with adjusting power of user units within a cellular network, where the users (the motes) have to communicate with one base station (the case of multiple base stations might be handled as well). The base station is assumed to be a phased-array antenna. We seek to maximize the lowest capacity among each individual channels, in order to ensure proper communication with all the motes. Our design variable is the power allocated to each unit. It turns out that this problem can be formulated as a convex optimization problem. This opens the interesting possibility of handling robustness in the direction of arrival matrix, and obtaining power allocations that are robust with respect to the uncertain locations of the motes.

Robust least-squares and estimation. One very classical method of inferencing states from measurements is least-squares. Flow recovery for example, necessitates the solution of a least-squares problem to estimate the flow field (see section 5.3). This problem involves uncertain data due for example to the uncertainty in the sensing (of temperature) and/or the location of the mote. To obtain an estimate that takes into account these uncertainties, one may apply the robust least-squares technique developed by El Ghaoui and Lebret [27]. When the uncertainty is modeled as "unstructured" (all elements in the data matrices are affected similarly), the resulting estimation algorithm in of the same complexity as ordinary least-squares (via singular value decomposition). This raises the issue of a faster, recursive implementation of robust least-squares.

Mote location optimization. The distribution of the motes will contain a random element, but it is reasonable to assume that the distribution process will not be *completely* random. We may assume that we are able to assign *regions* to groups of motes, and make sure that each group will end up in the assigned region. This raises the following issue: how can we optimize this assignment, given some objectives of accurate measurement, and given the uncertainty on the precise location of the motes within the regions? Let us first consider the case when the motes' location can be controlled with high precision. Each mote results in a measurement y_i, which we assume is linear in the variable to be estimated x, and contains a (Gaussian) random variable. That is, $y = Ax + v$, where A is a matrix, and v is a (Gaussian) measurement noise. Inevitably, some of the motes will send extremely inaccurate measurements, due to failure, loss of coverage due to foliage, etc. Therefore, we are to select the set of measurements (the columns of the matrix A) that are maximally informative. This problem is known as "experiment design" in statistics [28], and can be addressed using convex optimization [29]. Thus, we may apply the robust optimization approach and handle the case when the motes' location is uncertain, that is, when the matrix A in the experiment design problem is uncertain.

Antenna array design. In a variation of the above problem, we assume that each mote contains an antenna. The problem is to place these antennas, *and* optimize their gains, so as to optimally adjust the overall gain pattern. The problem has been addressed in the context of *fixed* positions of the antennas along a line in [30].

5.6 Performance Metrics and Assessment of Algorithms

One of the key strengths of this proposal is the existence of a hardware testbed. This hardware is imperative not to verify the theortical algorithm performance claims, but more importantly to provide a source of inspiration and intuition into the design of the algorithms themselves. In the initial stages of the project we will have a testbed of roughly 100 motes capable of being programmed and monitored over the RF network. This will allow for rapid experimentation in the early brainstorming stages of the project.

By the middle years of the project we should have several sensor networks with hundreds of sensors each available for real-time interaction and programming over the net. These facilities will be available to researchers outside of the UCB Sensorweb team as well. These networks will provide a direct physical evaluation of algorithm performance. The first metric is simple functionality, an essentially binary metric: can the algorithms be implemented in the computational resources provided? After basic functionality we will look for scaleability - are the decentralized algorithms truly decentralized, or are there "hidden master" requirements? What are the actually coefficients of the scaling to larger networks? Even an $O(N)$ algorithm may be unacceptable given the real constraints on memory, processor speed, and data rate in real network.

Metrics for energy constrained optimization of computation and communication do not currently exist. Do you reduce connectivity to hundreds of motes in order to save energy in a gateway mote with a failing battery? How valuable is the life of a single mote in relation to the health of the whole network? The answers to these questions will be application dependent, and implementation dependent as well. With real algorithms running on real hardware we will get the data we need to understand the underlying mathematics.

5.7 Emergent Behavior in Decentralized Sensor Networks (from local to global)

The sensorweb we implement will have of the order of a thousand individual dust motes. This is a scale at which macroscopic descriptions of the control and communication strategies begin to become relevant. For a concrete instance of what we mean by macroscopic description consider the problem of data fusion on the basis of observations of the environment by a large number of smart dust motes. Each mote will execute a mapping of its observation into a small number of possible symbols and these symbols will be exchanged between the motes to arrive at an estimate of the quantity being measured. What rules should the motes use to quantize their observations ? When there is a large number of motes it is advantageous to pose the problem in terms of the *number of motes that use a specific rule*. Indeed, the number of possible rules is fixed (independent of the size of the network of motes), so a description in terms of these fractions has a fixed complexity that does not grow with the number of motes. Further, one can discover important general rules based on such a description. For instance, Tsitsiklis [31, 32] proves that when the problem is to decide between one of M states of the world, and the motes communicate with a leader who makes the final decision, it suffices, for a large number of motes, to consider only $\binom{M}{2}$ distinct quantization rules and still achieve asymptotically optimal performance.

Another important aspect that needs to be explored in this project is whether the communication protocol between the motes can be arranged so that optimal behavior emerges in an adaptive way. To illustrate this point, let us first indicate the nontrivial nature of optimal strategies in some simple distributed smart dust tasks. Consider the problem of determining the parity of a set of observations made by the motes (each mote observes a bit, and the problem is to efficiently – with the minimum number of steps of communication – determine the parity of these bits). Gallager [33] proves the remarkable result that, over a noisy broadcast network, an exponential improvement in the communication cost of solving this problem is achievable if the N dust motes execute a hierarchical protocol : they should form roughly $\frac{N}{\log N}$ clusters, each consisting of roughly $\log N$ motes. The motes in a cluster should first communicate among each other to estimate their own parity, and then the clusters should communicate to estimate the overall parity. The key question of interest to us is : is it possible for the nodes to *discover this strategy on their own* ? Clearly, the optimal protocol will be problem dependent, and it is unlikely that one can foresee a priori all the scenarios that will be faced in the field. We are interested in studying the design of adaptive protocols in which there is emergent optimal behavior as a result of proper incentivization and learning from many attempts at the same problem.

6 Impact Assessment

The principal goal of this effort is the creation of a basic and rigorous framework for real-time distributed/ decentralized information processing. This will include a re-examination of a classical view of information theory. We propose to pull together all the various methods and algorithms to be developed under this effort towards the solution of difficult technical problems exemplified in our application scenarios. The impact of our work is likely to be broad in the form of conceptual design principles and synthesis procedures, metrics for evaluation, analyses tools, and engineering practices. The immediate realization of this will be best seen through our experimental test-beds. Specifically, we plan to leverage our hardware test bed as we describe below. Another important measure of impact of our proposal is the training of students and post-doctoral fellows. We discuss this in Section 10. Over the life of this grant, we intend to train at least 30 students in the broad areas of intelligent sensor networks and decentralized information fusion under various constraints.These students will not only be trained in formal scientific methods as indicated above, but will be involved in “hands-on” experience with dealing and interacting with some practical situations. We plan to develop education material and syllabus for real-time , decentralized estimation and information fusion. One of the salient features of this proposal is the integration of various reaearch results via a proof of concept test-bed.The students and post-doctoral fellows will play a key role in interacting and collaborating with potential user community in identifying not only additional specifics of application scenarios mentioned earlier but also in identifying other appropriate application domains, and accomplish integration of reaearch results as proof of concept around such test-beds.

6.1 Validation and Experimental Testbeds

We will leverage the hardware developed on the DARPA funded project (PI Pister) to validate the algorithms developed on the current project. We expect to have about a hundred macro-motes (about an inch cubed) with laser and RF communication, working solar cells, lithium batteries, accelerometers, pressure, temperature, humidity, magnetic field sensing and some processing power ready by December 1999, and miniaturized versions of these at the sub-centimeter scale by July 2000. Methods for aerial delivery of these sensors using a molded polysilicon winged structure for Silicon dandelion or maple seed like structures are also being developed. These silicon wings currently autorotate with a descent rate of 1 m/sec, and soon will provide several milliWatts of electrical power from integrated solar cells. Six-inch-scale Micro Air Vehicles (MAVs) with a SmartDust magazine and ejector for delivering the sensors have also been demonstrated, along with solid rocket micro-motors designed to give the motes a *one-time* motion capacity. Each one of these technological advances will provide a new challenge for the Sensorweb theory—mobility is certainly of great importance—how do we use it efficiently? Optimally? Algorithms, techniques, and software tools developed in the Sensorweb program will be tested first in software and then on different versions of the SmartDust hardware test bed.

7 Expected Outcomes: Timeline and Deliverables

The main deliverables on the project are algorithms, techniques and software tools for distributed sensing, and capstone demonstrations of some milestone projects. The primary contributions of the MURI will be the creation of a new field of real time decentralized information fusion over distributed networks. We expect to develop new methods, algorithms, techniques and train graduate students and researchers in fundamentally new and different ways. However, to be able to keep our research focussed and to have some touchstones for our theory we have formulated some specific milestone deliverables for the five year project as indicated here. The deliverables will be in the form of Software, Experiments and Demonstrations.

1. **Year One.** Software and algorithms for the Tank Tracker and Intelligent Mine Field. Primarily centralized algorithms will be demonstrated. Comparison of performance: false alarm, accurate prediction and sensitivity of multiple algorithms. Implementation of Tank Tracker on SmartDust hardware platform.

2. **Year Two.** Decentralized software and distributed algorithms for Tank Tracker and Intelligent Mine Field. Implementation of distributed algorithms on hardware platform. ChemBio Plume Detection software to be demonstrated (centralized). Comparison of centralized and decentralized schemes with respect to failure modes and speed of response. Preliminary algorithms on Star Trek locator.

3. **Year Three.** Decentralized algorithms for Friend/Foe tracker plus status monitor with multiple sensing modalities. Hardware demonstration of Star Trek locator. Software demonstration of of robust estimation and sensitivity capabilities of the Tank Tracker and Intelligent Mine Field software. Fault detection and network reconfiguration after useful life of sensor modules. Decentralized ChemBio Plume Detection Algorithms.

4. **Years Four and Five (Option Years).** Asymptotic behavior of algorithms as number of sensors increases. 3 D sensor gathering and networking using networks of MAVs plus ground sensor networks.

Coordination of terrestrial with aerial reconnaissance. Evaluation of 3 D versions of trackers for multiple ground based plus aerial assets.

8 Relevance to Army Needs

The PIs have had lengthy experience of technology transfer to Army and other DoD laboratories. Chandra was formerly Director of Intelligent Systems at the ARL, and and the Director of Atmospheric Research at the Army Research Laboratory,invovled in the formulation and execution of of programs in intelligent systems, He has considerable experience with distributed survivable communication and information networks, and high resolution battlefield environment models with special emphasis on transport /diffusion, detection, and identification of chemical and biological agents in the battle field. He was also the Diector of Mathematical and Computer sciences Division at the Army Research Office, with experience in developing and executing basic research programs in response to challenges of the digitized battlespace, imperatives of Force XXI and future requirements of Army After Next.He is currently involved in stochastic dynamic analysis of complex interactive networks and systems, as a critical step towards the national goal of a robust information infrastructure.

As a member of the Defense Science Study Group (DSSG) and the DARPA Information Science and Technology (ISAT) research directions team, Pister has toured in depth both active duty and research facilities in every branch of the DoD, and other agencies such as the CIA, DIA, NSA. Sastry, Russell and Malik have many contacts with the Unmanned Ground Sensors program at ARL, and the Intelligent Minefield program of ARDEC, as well as programs at MICOM. Anantharam has been involved in smart decentralized sensing work for DoD agencies. The investigators Jordan, Malik, Russell, and Sastry are also investigators in the Army MURI center *An Integrated Approach to Intelligent Systems.* We have also had extensive contacts with AFRL and NRL. We have personal contacts with Army scientists at these agencies, such as Dr. Norm Coleman (ARDEC), Drs. Phil Emmermann and Som Karamshetty (ARL), Dr. Chester Caroll (MICOM), and several others. These will be our primary contacts for technology transfer to the Army. In the area of UAVs/MAVs for distributed sensing Pister and Sastry have extensive contacts with the Air Force Research Laboratories as well as with the Naval Research Laboratory, for example Dr. Siva Banda (AFRL), Dr. Allen Schultz (NRL), etc. with whom we expect to continue our discussions.

The design of modern smart weapon systems brings with it severe new modeling, simulation, and validation problems beyond those of traditional weapon system design because of the diversity of interacting physical and logical systems which have to be represented and optimized. Weapon systems such as tanks and artillery are characteristically described by evolution of solutions of ordinary and partial differential equations with parameters to allow for adaptation to changing circumstances. Automated command and control systems can be formulated as logical systems in classical logic plus logics developed for AI and common sense reasoning. Evolutionary and adaptive behavior of knowledge and belief is represented in command and control by evolving knowledge bases of rules and facts, which rely on automated deduction to recommend decisions. Command and control, human behavior, and weapon systems interact with each other and the environment as a single feedback system. Many systems being developed have feedback loops which control the systems in which some of the controllers are human agents using decision aids, some are automated decision software, and some are conventional physical controllers. The research proposed on this grant will address the issue of design of hierarchies and heterarchies for the design of an effective management of a digital battlefield without cognitive overload.

Interactions with Army Labs

Each of the PIs as well as a number of our faculty investigators in this effort have considerable experience in interacting with several Army activities. Specifically, we had several technical interactions with scientists and engineers at places such as ARDEC,TACOM, MICOM, ARL, CAA, and the Soldier System Command at Natick. As a long time investigator and currently the Director of the MURI Center (1996-2001) on Integrated Approach to Intelligent Systems, Sastry has considerable interactions in the area of hybrid system design of fire control systems on helicopters, platooning and coordination of armored vehicles, and intelligent systems application to telemedicine and remote surgery.In this effort we intend to have an aggressive and planned program of technology transfer.For instance, it is our intention (Chandra, GWU) to coordinate across all tasks of the proposal towards the development of a coherent body of knowledge and engineering tools that can profittably be applied to specific applications of interest to the Army/DOD. Specifically, we will conduct workshops that promote technology transfer and facilitate substantial collaboration with Army/DOD activities on a sustained basis.

9 Facilities Available

Each of the two campuses has an extensive network of workstations, Suns, HPs, SGIs, and other graphics, simulation and visualization workstations with extensive user interactive facilities for recording interacting and viewing simulations at different levels of granularity. In addition there are three specific laboratories at Berkeley which will be used to support this project:

1. *The Intelligent Systems Laboratory* which specialises in having real time embedded systems prototyping facilities along with testing and measuring equipment for several real time sensing and control experiments. Facilities for CAD layout, design and testing of prototype boards is also available to the pproject, having been purchased on funds from our industrial, government partners and also from DURIP grants.
2. *The Vision Systems Laboratory* This laboratory has extensive hardware for real time video capture, multiple cameras and framegrabbers and TMS real time boards for processing of images.
3. *The Micro Electro Mechanical Systems (Smart Dust) Laboratory.* This laboratory has the facilities for sensor and actuator fabrication at sub-micron to centimeter scale. This laboratory leverages the resources of the Berkeley Sensor and Actuator Center which is a Center for MEMS expertise in fabrication miniature sensors and systems. This Center has use of the world's best "university" clean room microfabrication facility.

Our present budget *does not request any additional new equipment for the proposed research.*

10 Training of Young Scientists

We expect to train a new generation of young researchers on the area of real-time, decentralized estimation, with robustness and total energy constraints. This will be critical, since current curricula in statistics, probability theory, applied analysis, theoretical computer science, AI, signal processing and control theory do not provide students with the set of analytical tools required to address decentralized, real-time, robust sensing and data association with energy and bandwidth constraints. Course development will go hand in hand with the research since in five years. Sensorwebs of millions of devices will ubiquitously turn on and self configure, enabling a revolutionary new kind of information theory to reprogram entire networks with ease in both military and civilian scenarios. Our MURI project will train the next generation of researchers to be able to harness this incredible technological revolution.

There has been a thirty year separation of the teaching of software in computer science departments to the instruction in hardware in engineering departments. Certainly human behavior modelers and human factors specialists claim knowledge of neither computer science nor AI nor engineering. Guaranteeing correct interactions between high level decision software and low level hardware has to be carried out responsibly by those with interdisciplinary training or by interdisciplinary teams with the full gamut of skills in the areas mentioned. Thus software and hardware specifications for smart military systems combining command and control, weapon systems, and human behavior, simply cannot be separated into independent AI-software, hardware, and human factors specifications. Such complex systems have to be designed, tested, simulated as single integrated systems if one hopes to guarantee system performance without perpetual trial and error. The ideal is to express weapon systems, Logic-AI command and control, human behavior, and the environment all at once in a principled single model in which all these elements fit equally at the conceptual, mathematical, and algorithmic level. By providing an environment where students are exposed to hardware, software and also a wide variety of approaches, we will hope to train students to develop meaningful top down system architectures, design procedures, simulations, and training exercises.

11 Planned Collaborations

Each of the listed industrial partners either already supports research proposed on the MURI or has pledged to do so. They all currently support research at one or both of the MURI campuses and have promised to provide more support for the current project. Of the governmental and state laboratories, DARPA, ARO, Caltrans and NASA Ames are major supporters of research at Berkeley and George Washington University on the test bed areas.

Industrial

1. *Honeywell Technology Center.* Sastry and Nerode have been working with the Flight Vehicle Systems and the Air Traffic Management Systems group at Honeywell. The group at Honeywell is large and includes D. Godbole, J. Krause, J. Ward, and D. Cofer. It also includes contacts with Boeing. Honeywell funds focused research in air traffic management systems at Berkeley.

2. *United Technologies Research Center* supports the research of Pister to instrument their rotating turbines with Smart Dust.

3. *Berkeley Sensor and Actuator Center* which Pister is one of the Directors of has a list of 25 member companies, such as Hewlett Packard, Honeywell, Lockheed Martin. They especially support Pister in distributed wireless sensor networks.

4. *Interval Technologies* support the generation of 3D reconstructions from multiple camera views (Malik).

5. *SGI* supports the simulation and visualization environment at Berkeley in the form of equipment discounts. Its data mining group supports work at Berkeley and Stanford on learning probabilistic models.

6. *Rockwell* supports research in AI at Berkeley (Russell).

Government Laboratories and Academic Institutions

1. *California Department of Transportation.* In a major program, Caltrans supports the experimental and simulation test bed for intelligent vehicle highway systems at Berkeley.

2. *NASA Ames Research Center.* Berkeley and Stanford have a long working relationship with the Flight Control Systems and the AI groups of NASA Ames and most notably with researchers like Meyer and Hynes. NASA Ames is interested in the activities of the MURI.

3. *IRISA Rennes.* Sastry has many interactions with the group of Benveniste and have had an exchange of software (Signal and Signalea). The program continues with visits.

4. *INRIA Grenoble.* Sastry has extensive collaboration with Sifakis and his group, with long term visitors from Grenoble at Berkeley and software exchange (Lustre).

References

[1] J. B. Kruskal, "Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis," *Psychometrika*, vol. 29, pp. 1–27, 1964.

[2] S. Haykin, *Neural networks: A comprehensive foundation.* Upper Saddle River, NJ: Prentice Hall, 1999.

[3] E. Simoncelli, "Course-to-fine estimation of visual motion," in *IEEE Signal Processing Society 8th Workshop on Image and Multidimensional Signal Processing*, 1993.

[4] B. Lucas and K. T., "An iterative image registration technique with an application to stereo vision," in *Proc. 7th Intl Joint Conf on Artificial Intelligence*, pp. 674–679, 1981.

[5] E. Simoncelli, "Bayesian multi-scale differential optical flow," in *Handbook of Computer Vision and Applications*, pp. 397–422, 1999.

[6] G. Zweig and S. J. Russell, "Speech recognition with dynamic Bayesian networks," in *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, (Madison, Wisconsin), July 1998.

[7] G. Zweig and S. J. Russell, "Probabilistic modeling with bayesian networks for ASR," in *International Conference on Speech and Language Processing*, (Sydney, Australia), 1998.

[8] G. Zweig and S. J. Russell, "Probabilistic modeling with Bayesian networks for automatic speech recognition," *Australian Journal of Intelligent Information Processing Systems*, vol. 5, no. 4, pp. 253–260, 1999.

[9] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. New York: Academic Press, 1988.

[10] Y. Weiss, "Motion segmentation using em - a short tutorial," computer science report, University of California, Berkeley.

[11] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, New Jersey: Prentice-Hall, 1995.

[12] H. Pasula, S. Russell, M. Ostland, and Y. Ritov, "Tracking many objects with many sensors," in *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, (Stockholm, Sweden), Aug. 1999.

[13] M. Jerrum and A. Sinclair, "The Markov chain Monte Carlo method," in *Approximation Algorithms for NP-hard Problems* (D. S. Hochbaum, ed.), Boston: PWS Publishing, 1997.

[14] S. J. Nowlan, *Soft Competitive Adaptation: Neural Network Learning Algorithms based on Fitting Statistical Mixtures*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1991.

[15] G. C. G. Wei and M. A. Tanner, "A Monte-Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms," *Journal of the American Statistical Association*, vol. 85, no. 411, pp. 699–704, 1990.

[16] A. Yao, "Some complexity questions related to distributive computing," *Proc. of the 11th Annual ACM Symposium on the Theory of Computing*, 1979.

[17] A. Orlitsky and A. E. Gamal, "Communication complexity," in *Complexity in Information Theory* (Y. Abu-Mostafa, ed.), Springer-Verlag, 1988.

[18] L. Schulman, "Coding for interactive communication," *IEEE Transactions on Information Theory*, pp. 1745–1756, 1996.

[19] S. Venkatesan and V. Anantharam, "The common randomness capacity of a pair of independent discrete memoryless channels," *IEEE Transactions on Information Theory*, vol. 44, no. 1, pp. 215–224, 1998.

[20] S. Venkatesan and V. Anantharam, "The common randomness capacity of a network of processor," *IEEE Transactions on Information Theory*, 2000.

[21] V. S. Borkar, S. K. Mitter, and S. C. Tatikonda, "Optimal sequential vector quantization of markov sources," *Proceedings of the IEEE International Symposium on Information Theory*, p. 71, 1998.

[22] V. S. Borkar and S. K. Mitter, "Lqg control with communication constraints," in *Communications, Computation, Control and Signal Processing: a tribute to Thmoas Kailath*, pp. 365–374, Kluwer Academic Press, 1997.

[23] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, "Robust semidefinite programming," in *Semidefinite Programming and Applications* (H. W. R. Saigal, L. Vandenberghe, ed.), 1998. To appear.

[24] L. El Ghaoui, F. Oustry, and H. Lebret, "Robust solutions to uncertain semidefinite programs," *SIAM J. Optimization*, vol. 9, no. 1, pp. 33–52, 1998.

[25] A. BEN-TAL and A. NEMIROVSKI, "Robust convex optimization," *Math. Oper. Res.*, vol. 23, no. 4, pp. 769–805, 1998.

[26] L. El Ghaoui and G. Calafiore, "Deterministic state prediction under structured uncertainty," in *Proc. American Control Conf.*, 1999.

[27] L. El Ghaoui and H. Lebret, "Robust solutions to least-squares problems with uncertain data," *SIAM J. Matrix Anal. Appl.*, vol. 18, pp. 1035–1064, Oct. 1997.

[28] D. Cook and V. Fedorov, "Constrained optimization of experiment design," *Statistics*, vol. 26, pp. 129–178, 1995.

[29] L. Vandenberghe, S. Boyd, and S.-P. Wu, "Determinant maximization with linear matrix inequality constraints," *SIAM Journal on Matrix Analysis and Applications*, vol. 19, pp. 499–533, 1998.

[30] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret, "Second-order cone programming: interior-point methods and engineering applications," *Linear Algebra and its Applications*, 1997. Submitted.

[31] J. N. Tsitsiklis, "Decentralized detection," *Advances in Statistical Signal Processing*, vol. 2, pp. 297 –344, 1993.

[32] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Mathematics of Control, Signals, and Systems*, vol. 1, pp. 167 –182, 1988.

[33] R. G. Gallager, "Finding parity in a simple broadcast network," *IEEE Transactions on Information Theory*, vol. 34, pp. 176 –180, Mar. 1988.

12 Curriculum Vitae

12.1 Berkeley Investigators

Shankar Sastry

Professor of Electrical Engineering and Computer Sciences
and Director, Electronics Research Laboratory
269 M Cory Hall
University of California, Berkeley CA 94720
Tel: (510)642-1857. Fax: (510) 642-1341.
E-mail: sastry@robotics.eecs.berkeley.edu

Field of Specialization
Robotics, Complex systems and hybrid control, Simulation and Visualization

Education
B.Tech. (1977), Indian Institute of Technology, Bombay, India. M.S. EECS (1979), University of California, Berkeley. M.A. Mathematics (1980), University of California, Berkeley. Ph.D. EECS (1981), University of California, Berkeley.

Experience
Professor, University of California, Berkeley, 1988 - present, Visiting Professor, Institute National Polytechnique de Grenoble, 1999, Gordon Mc Kay Professor of Electrical Engineering and Computer Sciences, Harvard University, 1994, Visiting Vinton Hayes Professor of Electrical Engineering, MIT, Fall 1992, Directeur Recherche, Center Nationale Recherche Scientifique (CNRS), Toulouse, France, Summer 1991, Professor A Contratto, Universita di Roma. Summer 1990, 1991, Associate Professor, University of California, Berkeley, 1984 - 1988, Visiting Fellow, Australian National University, Canberra, Summer 1985, Assistant Professor, University of California, Berkeley, 1983 - 1984, Assistant Professor, MIT, Cambridge, 1980 - 1982.

Associate Editor of: *IMA Journal of Mathematical Control and Information*, *Journal of Mathematical Systems, Estimation and Control*, *International Journal of Adaptive and Optimal Control.*
Past Associate Editor of: *IEEE Transactions on Circuits and Systems*, *IEEE Transactions on Automatic Control*, *IEEE Control Magazine*, *Large Scale Systems.*

Awards
President of India Medal, 1977, IBM Faculty Development Grant, 1983, NSF Presidential Young Investigator Award, 1985, IEEE Student Best Paper Award, 1977, Eckmann Award of the American Control Council, 1990. M. A. Arts and Sciences, (honorary, 1994), Harvard University, Cambridge, Fellow IEEE (1995). Distinguished Alumnus Award, Indian Institute of Technology, 1999, David Marr Prize for the best paper at the International Conference in Computer Vision, Corfu, 1999.

Books Published
S.S. Sastry and M. Bodson, *Adaptive Control: Stability, Convergence and Robustness*, Prentice Hall, 1989.
R. Murray, Z. Li and S. Sastry *A Mathematical Introduction to Robotic Manipulation*, CRC Press, 1994.
P. Antsaklis, W. Kohn, A. Nerode and S. Sastry (editors), *Hybrid Systems II*, Springer Verlag, Lecture Notes in Computer Science, Vol. 999, 1995.
P. Antsaklis, W. Kohn, A. Nerode and S. Sastry (editors), *Hybrid Systems IV*, Springer Verlag, Lecture Notes in Computer Science, Vol. 1273, 1997.
T. Henzinger and S. Sastry (editors), *Hybrid Systems: Computation and Control*, Springer Verlag, Lecture Notes in Computer Science, Vol. 1386, 1998.
P. Antsaklis, W. Kohn, A. Nerode and S. Sastry (editors), *Hybrid Systems V*, Springer Verlag, Lecture Notes in Computer Science, Vol. 1567, 1999.
J. Baillieul, S. Sastry and H. Sussmann (editors), *Essays in Mathematical Robotics*, IMA Volumes in Mathematics and Its Applications, Vol. 104, Springer Verlag, 1999.
Shankar Sastry, *Nonlinear Systems: Analysis, Stability and Control*, Springer Verlag, Interdisciplinary and Applied Mathematics Series, Volume 10, 1999.

Publications
Here is a collection of some recent relevant papers out of over 200 published papers:

1. C. J. Tomlin and S. S. Sastry, "Bounded tracking for nonminimum phase nonlinear systems with fast zero dynamics", *International Journal of Control*, Vol. 68, No.4, 1997, pp. 819-847.

2. J. Lygeros, D. Godbole and S. Sastry, "Verified hybrid controllers for automated vehicles", *IEEE Transactions on Automatic Control*, Vol. 43, No. 4, April 1998, pp. 522-539

3. C. Tomlin, G. Pappas and S. Sastry, "Conflict Resolution in Multi-Agent systems: A case study in air traffic control," in *IEEE Transactions in Automatic Control*, Vol. 43, No. 4, April 1998, pp. 509-521.

4. F. Tendick, S. Sastry, R. Fearing and M. Cohn, "Applications of Micro-Mechatronics in Minimally Invasive Surgery," in *IEEE/ASME Transactions on Mechatronics*, March 1998, Vol. 3, (no.1):34-42.

5. Yi Ma, J. Košecká and S. Sastry, "Vision Guided Navigation for a Nonholonomic Mobile Robot," *IEEE Transactions on Robotics and Automation*, June 1999, Vol. 15 (3), pp. 521-536.

6. Košecká, J.; Tomlin, C.; Pappas, G.; Sastry, S. "Generation of conflict resolution maneuvers for air traffic management," *Proceedings of the 1997 IEEE/RSJ International Conference on Intelligent Robot and Systems*, IROS '97 Grenoble, France, Sept. 1997 IEEE, 1997. pp. 1598-603.

7. G. Lafferiere, G. Pappas and S. Sastry, "Subanalytic Stratifications and Bisimulations," in *Hybrid Systems: Computation and Control*, Springer Verlag Lecture Notes in Computer Science 1386, ed. T. Henzinger and S. Sastry, New York, April 1998, pp. 205-220.

8. F. Tendick, S. Sastry, R. Fearing and M. Cohn, "Applications of Micro-Mechatronics in Minimally Invasive Surgery," in *IEEE/ASME Transactions on Mechatronics*, March 1998, vol.3, (no.1):34-42.

9. G. Pappas, J. Lygeros and S. Sastry, " An approach to verification of the Center TRACON automation system", in *Hybrid Systems: Computation and Control*, Springer Verlag Lecture Notes in Computer Science 1386, ed. T. Henzinger and S. Sastry, New York, April 1998, pp. 289-304.

10. C. J. Tomlin and S. S. Sastry, "Switching through Singularities," *Systems and Control Letters*, Oct. 1998, Vol. 35 (3), pp. 145-154.

11. G. Pappas and S. Sastry, " Straightening Rectangular Inclusions", *Systems and Control Letters*, Sept. 1998, Vol. 35 (2), pp. 79-85.

12. J. Lygeros, C. Tomlin, S. Sastry. "Controllers for Reachability Specifications for Hybrid Systems". *Automatica, Special Issue on Hybrid Systems*, March 1999, Vol. 35 (3), pp. 349-370.

Kristofer Pister

Associate Professor of Electrical Engineering and Computer Sciences
Berkeley Sensor and Actuator Center
497 Cory Hall
Berkeley, CA 94720-1774
(510) 643-6690
pistereecs.berkeley.edu

Education:

1986 B.A., Applied Physics, University of California, San Diego

1989 M.S., Electrical Engineering, University of California, Berkeley

1992 Ph.D., Electrical Engineering, University of California, Berkeley

Academic Experience

1992-96 Assistant Professor of Electrical Engineering, University of California, Los Angeles

1996-Present Associate Professor of EECS, University of California, Berkeley

Industrial Experience

1986 Summer intern, IBM, East Fishkill, NY

1987-88 Hewlett Packard Circuit Technology Research and Development, Palo Alto, CA (full time during summer 87, half time thereafter)

1992-99 Consultant, RAND Corp., Santa Monica, CA, Tanner Research, Pasadena, CA, Neilsen Research Inc., Sunnyvale, CA, Analog Devices, Wilmington, MA, MAXIM Integrated Products, Sunnyvale, CA.

Awards and Honors 1992 Angelakos Award for altruistic assistance to fellow students, Dept. of EECS, UC Berkeley.
1996 NSF CAREER award

Patents 1995 Etchants for use in micromachining of CMOS microaccelerometers and microelectromechanical devices and method of making the same, U.S. patent No. 5,659,195.

Selected Journal Publications

1. Pister, K., Judy, M., Burgett, S., Fearing, S., "Microfabricated Hinges," *Sensors and Actuators,*", Vol. 33, no. 3, pp. 249-256, 1992.

2. Murray, R., Deno, D., Pister, K., Sastry, S., "Control Primitives for Robot Systems," *IEEE Trans. Systems, Man, and Cybernetics*, Vol. 22, No. 1, pp. 183-193, Jan/Feb 1992.

3. L.Y. Lin, S.S. Lee, K.S.J. Pister, and M.C. Wu, " Micro-Machined Three-Dimensional Micro-Optics for Integrated Free-Space Optical System," *IEEE Photonics Technology Letters*, Vol. 6, No. 12, Dec. 1994, pp. 1445-1447.

4. G. Lin, K.S.J. Pister, and K.P. Roos, "Measurement of Contractile Characteristics from Isolated Heart Cells Using a Minaturized Force Transducer," *Proc. Intl. Conf. of the IEEE Engineering In Medicine and Biological Society*, Baltimore, MD, November 1994, Vol. 2, pp. 1017-1018.

5. G. Lin, K.S.J. Pister, and K.P. Roos, "Novel Microelectromechanical System Force Transducer to Quantify Contractile Characteristics from Isolated Cardiac Muscle Cells," *J. Electrochemical Society*, March 1995, Vol. 142, No. 3, L31-3.

6. Lin, L.Y.; Lee, S.S.; Pister, K.S.J.; Wu, M.C., "Self-aligned hybrid integration of semiconductor lasers with micromachined micro-optics for optoelectronic packaging," *Applied Physics Letters*, 29 May 1995, Vol.66, (no.22): 2946-8.

7. S.S. Lee, L.Y. Lin, K.S.J. Pister, M.C. Wu, H.C. Lee, P. Grodzinski, "Passively aligned hybrid integration of 8*1 micromachined micro-Fresnel lens arrays and 8*1 vertical-cavity surface-emitting laser arrays for free-space optical interconnect.," *IEEE Photonics Technology Letters*, Sept. 1995, Vol.7, (no.9):1031-3.

8. L.Y. Lin, S.S. Lee, K.S.J. Pister, and M.C. Wu, "Three-dimensional Micro-Fresnel Optical Elements Fabricated by Micromachining Technique," *Electronics Letters*, Vol. 30 No. 5, 1995 pp. 448-449.

9. G. Lin, K.S.J. Pister, and K.P. Roos, "Micro Scale Force Transducer System to Quantify Isolated Heart Cell Contractile Characteristics," *Sensors and Actuators A*, Vol. 46-47, No. 1-3, pp. 233 - 236, 1995.

10. Mark Ross and Kristofer Pister, "Micro-Windmill for Optical Scanning and Flow Measurement," *Sensors and Actuators A*, Vol. A47, No. 1-3 pp.576-579, 1995.

11. Gunawan, D.S.; Lin, L.-Y.; Pister, K.S.J., "Micromachined corner cube reflectors as a communication link, *Sensors and Actuators A (Physical)*, March-April 1995, Vol.A47, (no.1-3):580-3.

12. Richard Yeh, Ezekiel J.J. Kruglick, Kristofer S.J. Pister, "Surface Micromachined Components for Articulated Microrobots," *IEEE/ASME J. MEMS*, Vol. 5, N. 1, March 1996, pp. 10-17.

13. Patrick B. Chu, Phyllis R. Nelson, Mark L. Tachiki, and Kristofer S. J. Pister, "Dynamics of polysilicon parallel-plate electrostatic actuators," *Sensors and Actuators A (Physical)*, March-April 1996.

14. B. Tsap, K.S.J. Pister, H.R. Fetterman, "MEMS Orientational Optomechanical Media for Nonlinear Microwave Applications," *IEEE Microwave and Guided Wave Letters*, December 1996.

Venkat Anantharam

Professor of Electrical Engineering
University of California, Berkeley, CA 94720

Field of Specialization

Communication networks, Queueing networks, Stochastic control, Coordination Theory, Information Theory.

Education

B.Tech. in Electrical Engineering (Electronics), Indian Institute of Technology Madras, India, 1980
M. S. in Electrical Engineering, University of California, Berkeley, 1982
M. A. in Mathematics, University of California, Berkeley, 1983
C. Phil in Mathematics, University of California, Berkeley, 1984
Ph. D. in Electrical Engineering, University of California, Berkeley, 1986

Experience

Member of Technical Staff, Bell Communications Research, Summer 1984
Assistant/Associate Professor of Electrical Engineering, Cornell University, July 1986- May 1991-June 1994
Associate/Full Professor of Electrical Engineering, University of California, Berkeley, July 1994- June 1996-present

Societies

Fellow, IEEE. Member, AMS.
Associate Editor of the following journals
Queueing Systems : Theory and Applications; Markov Processes and Related Fields; IEEE Transactions on Information Theory.

List of 5 publications most closely related to the research:

1. P. Viswanath, V. Anantharam, and D. Tse, "Optimal Sequences, Power Control, and User Capacity of Synchronous CDMA Systems with Linear MMSE Multiuser Receivers". To appear in *IEEE Transactions on Information Theory.*
2. "Optimal Sequences and the Sum Capacity of Synchronous CDMA Systems". P. Viswanath and V. Anantharam, Proceedings of the 36th Annual Allerton Conference on Communication, Control, and Computing, University of Illinois, Urbana -Champaign, Sep. 23 - 25, 1998, pp. 272 -281. Full paper to appear in *IEEE Transactions on Information Theory.*
3. "Optimal Routing Control: Game Theoretic Approach". Richard J. La and V. Anantharam, *36rd IEEE Conference on Decision and Control*, San Diego, CA, December 1997. (To appear). Available at "http://www.eecs.berkeley.edu/ ananth"
4. "The Common Randomness Capacity of a Network of Processors", *S. Venkatesan* and *V. Anantharam.* To appear in *IEEE Transactions on Information Theory*.
5. "Bits through queues", *V. Anantharam* and *S. Verdu. IEEE Transactions on Information Theory*, Vol. 42, No. 1, pp. 4 -18, January 1996.

List of 5 other significant publications:

1. "Optimal Flow Control Schemes that Regulate the Burstiness of Traffic", T. Konstantopoulos and V. Anantharam, *IEEE/ACM Transactions on Networking*, Vol. 3, No. 4, pp. 423 -432, August 1995.
2. "Networks of Queues with Long-Range Dependent Traffic Streams", V. Anantharam, *Stochastic Networks: Stability and Rare Events*, edited by Paul Glasserman, Karl Sigman, and David Yao, Lecture Notes in Statistics, Vol. 117, pp. 237 - 256, Springer Verlag, 1996.

3. "Achieving 100% Throughput in an Input-Queued switch", N. McKeown, V. Anantharam, and J. Walrand, *Proceedings of the IEEE INFOCOM*, San Francisco, March 24 -28, 1996, Volume 1, pp. 296 -302. 4. "The Stability Region of the Finite User Slotted ALOHA Protocol", *V. Anantharam. IEEE Transactions on Information Theory*, Vol. 37, No. 3, pp. 535 -540, May 1991.
5. "Asymptotically Efficient Adaptive Allocation Schemes for Controlled Markov Chains : Finite parameter space", *R. Agrawal, D. Teneketzis*, and *V. Anantharam. IEEE Transactions on Automatic Control*, Vol. 34, No. 12, pp. 1249 -1259, (1989).

Laurent El Ghaoui

Acting Associate Professor of Electrical Engineering and Computer Sciences
267 M Cory Hall
University of California, Berkeley CA 94720
Tel: (510)642-8351. Fax: (510) 643-8351.
E-mail: elghaoui@robotics.eecs.berkeley.edu

Field of Specialization
Convex optimization, robust optimization, analysis and design of complex systems, worst-case simulation
Education
M.S. (1985), Ecole Polytechnique, France.
Ph.D. Aero/Astro (1990), Stanford University.
Experience
Acting Associate Professor, University of California, Berkeley, 1999 - present, Professor, Ecole Nationale Supérieure de Techniques Avancées and Ecole Polytechnique, France, 1992, Visiting Professor, Université Catholique de Louvain, Belgium, Spring 1998, Research Scientist, Laboratoire Systèmes de Perception, France, 1990.

Associate Editor of: *SIAM Journal of Matrix Analysis and its Applications, American Control Conference.*

Awards: Bronze Medal, Centre National de la Recherche Scientifique, France, 1999.

Books Published
S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics.* SIAM, Philadelphia, PA, June 1994.

L. El Ghaoui and S.I. Niculescu.*Recent Advances on Linear Matrix Inequality Methods in Control,* SIAM, 1999.

Selected Recent Publications

1. A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. Robust semidefinite programming. In H. Wolkowicz R. Saigal, L. Vandenberghe, editor, *Semidefinite Programming and Applications.* Kluwer, 1999. To appear.

2. L. El Ghaoui, F. Oustry and H. Lebret. Robust solutions to uncertain semidefinite programs. *SIAM J. Optimization*, vol.9, no. 1, 1998.

3. S. Boyd and L. El Ghaoui. Method of centers for minimizing generalized eigenvalues. *Linear Algebra and Applications, special issue on Numerical Linear Algebra Methods in Control, Signals and Systems*, 188, July 1993.

4. L. El Ghaoui. State-feedback control of systems with multiplicative noise via Linear Matrix Inequalities. *Syst. Control Letters*, 24(3):223–228, February 1995.

5. L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM J. Matrix Anal. Appl.*, 18(4):1035–1064, October 1997.

6. L. El Ghaoui and G. Scorletti. Control of rational systems using Linear-Fractional Representations and Linear Matrix Inequalities. *Automatica*, 32(9):1273–1284, September 1996.

7. L. El Ghaoui and G. Calafiore. Deterministic state prediction under structured uncertainty. In *Proc. American Control Conf.*, 1999.

Michael Jordan

Professor of Electrical Engineering and Computer Sciences
and Professor of Statistics
739 Soda Hall
University of California, Berkeley CA 94720
Tel: (510) 642-3806. Fax: (510) 642-5775.
E-mail: jordan@cs.berkeley.edu

Field of Specialization

Estimation and inference in graphical models, Pattern recognition, Nonlinear multivariate statistics, Time series analysis

Education

B.S. (1978), Louisiana State University. M.S. Mathematics (1980), Arizona State University. PhD Cognitive Science (1985), University of California, San Diego.

Experience

Professor, Department of Electrical Engineering and Computer Science, Department of Statistics, University of California, Berkeley, 1998 – present, Professor, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 1997 – 1998, Associate professor with tenure, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 1994 – 1997, Associate professor, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 1992 – 1994, Assistant professor, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 1988 – 1992, Postdoctoral researcher, Department of Computer and Information Science, University of Massachusetts, Amherst, 1986 – 1987.

Associate Editor of: *Journal of the American Statistical Association*, *Machine Learning*, *Neural Computation*, *Cognition*, *International Journal of Neural Systems*, *Neural Networks.*

Awards

Presidential Young Investigator Award, 1991 – 1996. Best Paper Award, American Control Conference, 1991. MIT Class of 1947 Career Development Award, 1992 – 1995.

Books Published

Mozer, M. C., Jordan, M. I., & Petsche, T. (Eds.). (1997). *Advances in Neural Information Processing Systems 9*, Cambridge MA: MIT Press.

Jordan, M. I., Kearns, M. J. & Solla, S. A. (Eds.). (1998). *Advances in Neural Information Processing Systems 10*, Cambridge MA: MIT Press.

Jordan, M. I. (Ed.). (1999). *Learning in Graphical Models*, Cambridge, MA: MIT Press.

Domeniconi, C., & Jordan, M. I. (in press). *Discorsi sulle Reti Neurali e l'Apprendimento.* Milan: Franco Angeli Editore.

Jordan, M. I., & Sejnowski, T. J. (Eds.). (in press). *Graphical Models: Foundations of Neural Computation.* Cambridge MA: MIT Press.

Recent Publications (out of more than 100 published papers)

1. Jacobs, R. A. & Jordan, M. I. (1993). Learning piecewise control strategies in a modular neural network architecture. *IEEE Transactions on Systems, Man, and Cybernetics, 23*, 337–345.

2. Jordan, M. I. & Jacobs, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, *6*, 181-214.

3. Jordan, M. I., Flash, T., & Arnon, Y. (1994). A model of the learning of arm trajectories from spatial targets. *Journal of Cognitive Neuroscience*, *6*, 359-376.

4. Jordan, M. I. (1995). The organization of action sequences: Evidence from a relearning task. *Journal of Motor Behavior*, *27*, 179-192.

5. Wolpert, D., Ghahramani, Z., & Jordan, M. I. (1995). An internal forward model for sensorimotor integration. *Science*, *269*, 1880-1882.
6. Jordan, M. I., & Xu, L. (1995). Convergence results for the EM approach to mixtures-of-experts architectures. *Neural Networks*, *8*, 1409–1431.
7. Saul, L. K., Jaakkola, T., & Jordan, M. I. (1996). Mean field theory for sigmoid belief networks. *Journal of Artificial Intelligence Research*, *4*, 61–76.
8. Cohn, D., Ghahramani, Z., & Jordan, M. I. (1996). Active learning with statistical models. *Journal of Artificial Intelligence Research*, *4*, 129–145.
9. Ghahramani, Z., & Jordan, M. I. (1997). Factorial Hidden Markov models. *Machine Learning*, *29*, 245–273.
10. Jordan, M. I., Ghahramani, Z., & Saul, L. K. (1997). Hidden Markov decision trees. In M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems 9*. Cambridge MA: MIT Press.
11. Houde, J., & Jordan, M. I. (1998). Adaptation in speech production. *Science*, *279*, 1213–1216.
12. Jaakkola, T., & Jordan, M. I. (1999). Variational probabilistic inference and the QMR-DT network. *Journal of Artificial Intelligence Research, 10*, 291–322.

Jitendra Malik

Professor of Electrical Engineering and Computer Science
University of California, Berkeley, CA 94720.

Field of Specialization

Computer Vision, Computational Modeling of Human Vision

Education

Ph.D. in Computer Science, Stanford University, December 1985
.B.S. in Electrical Engineering, Indian Institute of Technology, Kanpur, 1980.

Experience

Vice-Chair of Graduate Matters, EECS, 1995-1998.
Professor, EECS, UC Berkeley, since July 1996.
Associate Professor, EECS, UC Berkeley, July 1991-June 1996.
Assistant Professor, EECS, UC Berkeley, Jan 1986-June 1991.
Member, Groups on Cognitive Science and Vision Science, UC Berkeley.

Honors and Awards

Rosenbaum Fellow, Isaac Newton Institute, University of Cambridge, 1993.
Presidential Young Investigator Award 1989.
IBM Faculty Development Award 1986-88.
IBM Fellowship for Doctoral Study in Computer Science 1983-85.
Best Graduating Student in Electrical Engineering, IIT Kanpur 1980.
One of the top ten students in the Indian School Certificate Examination 1974.

Selected Invited Talks, Last 5 Years

Smart Cars and Smart Roads, British Machine Vision Conference, 1995 (plenary).
Finding Objects In Large Collections Of Images, Snowbird Conference on Machines that Learn, 1997.
Visual Grouping, Graph Partitioning, And Eigenvalue Problems,MTNS, Padova, 1998
Computational Models of Visual Grouping, German-American Frontiers of Science Symposium, 1999.
Region-based Image Retrieval, DAGM 99 (German Conf. on Pattern Recognition), 1999 (plenary).

Ten Selected publications (from more than 80):

1. J. Malik and P.Perona, "Preattentive texture discrimination with early vision mechanisms," *Journal of Optical Society of America A*, **7** (2), May 1990, pp. 923-932.

2. P. Perona and J. Malik, "Scale space and edge detection using anisotropic diffusion," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, **12** (7), July 1990, pp. 629-639.

3. D. Jones and J. Malik, " Computational framework for determining stereo correspondence from a set of linear spatial filters," *Image and Vision Computing* **10**(10), December 1992, pp. 699-708.

4. J. Weber and J. Malik, "Robust computation of optical flow in a multi-scale differential framework," *International Journal of Computer Vision*, 14(1), Jan 1995

5. J. Malik and R. Rosenholtz, "Computing local surface orientation and shape from texture for curved surfaces," *International Journal of Computer Vision*, 23(2), June 1997, pp. 149-168.

6. D. Forsyth, J. Malik and R. Wilensky, "Searching for Digital Pictures," *Scientific American*, 276(6), June 1997, pp. 88-93.

7. J. Shi and J. Malik, "Normalized Cuts and Image Segmentation," *Proc. of IEEE CVPR, Puerto Rico*, June 1997, pp 731-737.

8. B. Coifman, D. Beymer, P. McLauchlan and J. Malik, "A Real-time Computer Vision System for vehicle tracking and traffic surveillance," *Transportation Research C*, 6 (1998) 271-288.

9. Camillo J. Taylor, Jana Kosecka, Robert Blasi and Jitendra Malik, "A Comparative Study of Vision-Based Lateral Control Strategies for Autonomous Highway Driving," *IJRR*, 18(5), May 1999, pp. 442-453.

10. J. Malik, B.L. Anderson and C.E. Charowhas, "Stereoscopic occlusion junctions," *Nature Neuroscience*, 2(9), Sept. 1999, pp.840-843

Stuart Russell

Professor, Computer Science Division, UC Berkeley

Field of Specialization

Artificial intelligence: machine learning, probabilistic reasoning, real-time decision making, sensor-based reasoning, autonomous vehicles.

Experience

Professor, Computer Science Division, UC Berkeley, 1996–
Associate Professor, Computer Science Division, UC Berkeley, 1991–96
Assistant Professor, Computer Science Division, UC Berkeley, 1986–91
Visiting Professor, Oxford University, 1990, 1992
SERC Visiting Professor, Strathclyde University, 1992

Education

B.A. Physics (1st class honours), Oxford University, 1982.
PhD, Computer Science, Stanford University, 1986.

Principal Honours and Awards

Distinction, British Mathematical Olympiad, 1977
Major Scholar, Wadham College, Oxford University, 1979–82
NATO Scholar, 1982–85
NSF Presidential Young Investigator, 1990–95
IJCAI Computers and Thought Award, 1995
Miller Professorship, UC Berkeley, 1996
Fellow of the American Association for Artificial Intelligence, 1997
Distinguished Paper Prize, Int'l Joint Conf. on Artificial Intelligence, 1997
Forsythe Lecturer, Stanford University, 1998
Jon Postel Distinguished Lecturer, UCLA, 1999
Best Industry-Related Paper (Honorable Mention), Int'l Conf. Patt. Rec., 1994.
Distinguished Paper Prize, Int'l Joint Conf. on AI, 1997.
Best Paper Prize (Honorable Mention), Int'l Conf. on Machine Learning, 1999.

Invited/Keynote Speaker at Discovery Science 99, Tokyo; Automatic Speech Recognition and Understanding 99, Keystone; ACM Conf. on Computational Learning Theory, Madison, 1998; Theoretical Aspects of Rationality and Knowledge, Evanston, 1998; Artificial Intelligence, Vancouver, 1998; Vancouver Cognitive Science Conference, 1998; Neural Information Processing Systems, Denver, 1997; European Conference on Machine Learning, Prague, 1997; Reconnaissance des Formes et Intelligence Artificielle, Rennes, 1996; Int'l Joint Conf. on Artificial Intelligence, Montreal, 1995; Euro. Conf. Economics and Game Theory, Marseille, 1994; American Philosophical Ass'n, Chicago, 1992.

Professional activities

Editor, Prentice Hall Series in Artificial Intelligence
Area Editor for Artificial Intelligence, Journal of the ACM
Associate Editor, Journal of Artificial Intelligence Research
Editorial Board, Artificial Intelligence Journal
Editorial Board, Machine Learning Journal
Editor, AI Communications
Advisory Editor for AI and Computer Science, MIT Press Encyclopaedia of the Cognitive Sciences
Executive Council, American Association for Artificial Intelligence
Advisory Committee, International Joint Conference on Artificial Intelligence
Program Chair, Int'l. Conf. on Machine Learning, 1995
Program Area Chair, AAAI 90, AAAI 94, NIPS 96
Chair, AAAI Symposium on AI and Limited Rationality, 1989

Selected Recent Publications

S. Russell *The Use of Knowledge in Analogy and Induction.* London: Pitman, 1989.
S. Russell and E. Wefald *Do the Right Thing: Studies in Limited Rationality.* MIT Press, 1991.
S. Russell and E. Wefald "Principles of Metareasoning". *Artificial Intelligence* **49**, 1991.
S. Russell and D. Subramanian "Provably bounded-optimal agents." *J. Artif. Intell. Research*, **2**, 1995.
S. Russell and P. Norvig *Artificial Intelligence: A Modern Approach.* Prentice Hall, 1995.
S. Zilberstein and S. Russell "Optimal composition of real-time systems." *Artificial Intelligence*, **82**, 1996.
J. Binder, D. Koller, S. Russell, K. Kanazawa, "Adaptive Probabilistic Networks with Hidden variables." *Machine Learning*, **29**, 1997.
N. Friedman and S. Russell, "Image Segmentation in Video Sequences." In *Proc. UAI-97.*
R. Parr and S. Russell, "Reinforcement Learning with Hierarchies of Machines." In *Proc. NIPS-97.*
T. Huang and S. Russell, "Object Identification: A Bayesian Analysis with Application to Traffic Surveillance." *Artificial Intelligence*, **103**, 1998 (invited paper).
G. Zweig and S. Russell, "Speech Recognition with Dynamic Bayesian Networks." In *Proc. AAAI-98.*
Hanna Pasula, Stuart Russell, Michael Ostland, and Ya'acov Ritov, "Tracking many objects with many sensors." In *Proc. IJCAI-99.*

George Washington University Investigators
Jagdish Chandra

Education:

Ph.D., 1965, Mathematical Sciences, Rensselaer Polytechnic Institute,Troy, New York
M.A., 1957, Mathematics, Osmania University, India
B.A., 1955, Mathematics, Osmania University, India.

Professional Experience:

Senior Research Scientist, George Washington University, 1999-present
Deputy Director, IS & T Directorate/Director Atmospheric Research, Army Research Laboratory, 1997-1999, (SES-4)
Director, Mathematical & Computer Sciences Division, Army Research Office, 1975-1997.(Member of the Senior Executive Service since July 1979, SES-4)
Associate Director, Mathematical Sciences Division, Army Research Office, 1973-1975
Chief, Applied Mathematics Program, Army Research Office, 1970-1973
Research Mathematician, U.S. Army Watervliet Arsenal, 1966-1970
Instructor of Mathematics, Rensselaer Polytechnic Institute, 1965-1966

Concurrent Positions:

Adjunct Professor of Computer Science, Johns Hopkins University, 1998-present
Adjunct Professor of Mathematics, Duke University, 1980-1997
Adjunct Associate Professor of Mathematics, Duke University, 1972-1980
Adjunct Assistant Professor of Mathematics, Rensselaer Polytechnic Institute, 1966-1970
Acting Director, Electronics Division, ARO, 1988-1990 and 1994-1995.

Areas of Scientific/Technical Expertise:

Nonlinear Analysis, Intelligent Systems. Systems and Control Theor y, Simulation and Modeling, Stochastic Analysis, Computational Modeling.

Professional Activities: Member, AMS, SIAM, SEA
Senior Member, IEEE
Editor, SIAM Review
Associate Editor, International Journal of Nonlinear Analysis.

Recent Awards:

SIAM Commendation for Public Service for Contributions to Applied Mathematics and C omputing, 1999
SES Exceptional Performance Awards, 1989, 1990, 1991, 1994, 1995
SDIO Special Achievement Award, 1989
Department of the Army Commendation, 1989, 1995
Superior Civilian Service Award, 1997, 1999.

Some Recent Publications/Presentations:

Of the more than 40 publications in refereed national and international journals: 1. " Simulation and Modeling Issues for Future Army Command and Control Systems," (Keynote Address), Proc. 1998 Conference on Simulation Methods and Applications: Parallel and Distributed Simulation, p. 3-11,

2. "Integration of Data Fusion and Network Control," (Keynote presentation), Proc. Workshop on Data Compression Processing Techniques for Missile Guidance Data Links, p. 411-440, Special Report MG-99-5 (March 1999)