

AFRL-IF-RS-TR-2004-64
Final Technical Report
March 2004



SENSOR WEBS OF SMART DUST: DISTRIBUTED SIGNAL PROCESSING/DATA FUSION/INFERENCING IN LARGE MICROSENSOR ARRAYS

University of California at Berkeley

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K253

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-64 has been reviewed and is approved for publication.

APPROVED:

/s/
Richard C. Butler II
Project Engineer

FOR THE DIRECTOR:

/s/
WARREN H. DEBANY, JR.
Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE MARCH 2004	3. REPORT TYPE AND DATES COVERED FINAL Jun 00 – Jun 03	
4. TITLE AND SUBTITLE SENSOR WEBS OF SMART DUST: DISTRIBUTED SIGNAL PROCESSING/DATA FUSION/INFERENCING IN LARGE MICROSENSOR ARRAYS			5. FUNDING NUMBERS G - F30602-00-2-0538 PE - 62302E PR - K253 TA - 25 WU - A1	
6. AUTHOR(S) Kannan Ramchandran Kristofer Pister				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Regents of the University of California (UC, Berkeley), Sponsored Projects Office 336 Sproul Hall Berkeley CA 94720			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFGA 3701 North Fairfax Drive 525 Brooks Road Arlington VA 22203-1714 Rome NY 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2004-64	
11. SUPPLEMENTARY NOTES DARPA Program Manager: Sri Kumar/IPTO/(703) 696-0174 AFRL Project Engineer: Richard C. Butler/IFGA/(315) 330-1888 Richard.Butler@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) The research has, to a large measure, addressed many of the theoretical foundations and algorithmic advances necessary to exploit the capabilities of Sensor Webs. Leveraging the SmartDust infrastructure, this project has been instrumental in developing a framework for real-time distributed/decentralized information processing.				
14. SUBJECT TERMS Distributed Sensor Networks, Sensor Webs, SmartDust, Distributed Position Estimation, Localization, Magnetometers				15. NUMBER OF PAGES 19
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

Introduction	1
Applications, Results, and Discussion.....	1
Distributed Sampling	2
Distributed Estimation	2
Multi-terminal (distributed) Source Coding	3
Duality Theory	3
Security	4
Routing and Adaptation for Improving Energy Efficiency of Sensor Networks	4
Routing with Long Links	5
Energy Metrics for Sensor Networks	5
The Ivy Project	6
Localization in Sensor Networks	7
Effort-1:.....	7
Effort II:	8
Effort-III: distributed position estimation.....	9
Distributed map building and navigation.....	9
Sensor Networks for Multiple Target Tracking	10
Environmental Monitoring.....	11
Interacting particle systems models in sensor networks.....	11
Error exponents for distributed detection and security	12
Game-theory in communication networks	12
Learning theory in Sensor Networks	13
Fundamental role:.....	13
Sensor field model learning with application to blind separation of linearly mixed signals:	13
References.....	15

Introduction

This research has, to a large measure, addressed many of the theoretical foundations and algorithmic advances necessary to exploit the capabilities of Sensor Webs. Leveraging the SmartDust infrastructure, this project has been instrumental in developing a framework for real-time distributed / decentralized information processing that address the following key issues:

- 1) We have developed several constructive, distributed data compression and signal processing algorithms that approach the fundamental theoretical bounds; collaborative, source-channel coding and detection methods that optimize the end- to-end performance metrics.
- 2) Algorithms for reliable distributed information fusion and interpretation, using a learning-theoretic formalism were also successfully investigated.
- 3) Information-theoretic foundations for understanding the design, performance optimization, and fundamental limits of distributed sensory systems were developed.
- 4) Demonstration of these algorithms on a sensor testbed, dubbed the Berkeley Campus Sensor Network (BCSN), was partially executed.

This research effort has gone a long way towards providing a unified approach to the development of theoretical tools, algorithms, and software for Sensor Webs, and to connect this with a variety of application scenarios as explained in detail in the sequel. Our approach was to include a re-examination of a classical view of information theory and practical constructive algorithms inspired by this. 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.

Applications, Results, and Discussion

We have developed and mathematically analyzed algorithms for sampling, estimation, distributed compression, secure communication, power-efficient routing, adaptation, localization, tracking, environmental monitoring, and other wireless sensor network applications, with the goal of optimizing the trade-off between communication and computation.

Since most large, wireless sensor networks will most likely be deployed in a random fashion, i.e., by dropping the nodes out of an airplane, we study random networks, in which nodes are randomly (usually uniformly) distributed in a region of deployment.

Distributed Sampling

The problem of deterministic oversampling of bandlimited sensor fields in a distributed communication-constrained processing environment, where it is desired for a central intelligent unit to reconstruct the sensor field to maximum point wise accuracy, was studied. It was shown, using a dither-based sampling scheme, that it is possible to sample fields using minimal inter-sensor communication, with the aid of a multitude of low- precision sensors. The feasibility of having a flexible tradeoff between the oversampling rate and the Analog to Digital (AID) quantization precision per sensor sample with respect to achieving exponential accuracy in the number of bits per Nyquist-period was also shown. Our analysis revealed a key underpinning "conservation of bits" principle, i.e., the bit budget per Nyquist-period can be distributed along the amplitude-axis (NO precision) and space (or time or space-time), using oversampling in an almost arbitrary discrete-valued manner, while retaining the same reconstruction error decay profile. Interestingly, this oversampling is possible in a highly localized communication setting, with only nearest-neighbor communication, making it very attractive for dense sensor networks operating under stringent inter-node communication constraints. It was also shown how the proposed scheme incorporates security as a by-product due to the presence of an underlying dither signal, which can be used as a natural encryption device for security. The choice of the dither function enhances the security of the network.

Distributed Estimation

An information-theoretically achievable rate-error region for an unreliable network of sensors observing a physical process, such as temperature, under symmetric sensor measurement statistics and rate constraints was derived. For independent, jointly Gaussian measurement noise and squared-error distortion, the proposed distributed encoding and estimation framework was found to have the following robustness property: When any k out of n rate- R bits/sec sensor transmissions are received, the central unit's estimation quality can match the best estimation quality that can be achieved from a completely reliable network of k sensors, each transmitting at rate R . Furthermore, when more than k out of the n sensor transmissions are received, the estimation quality strictly improves. When the network has clusters of collaborating sensors, an important question is whether clusters should compress their raw measurements or should they first try to estimate the source from their measurements and compress the estimates instead. For many interesting cases, it was shown that there is no loss of performance in the distributed compression of local estimates over the distributed compression of raw data in a rate-distortion sense, i.e., encoding the local sufficient statistics is good enough.

Multi-terminal (distributed) Source Coding

The concept of successive refinement of information for multiple users was developed and a characterization of the rate versus quality regions for the Gaussian source was determined. The performance of the proposed scheme was shown to be superior to conventional approaches, based on multiplexed solutions of optimal point-to-point, successively refinable transmission strategies. We also developed a more universal approach to multi-user successive refinement, based on the Wyner-Ziv method of coding with side-information, where the source reconstruction based on the base layer is treated as side-information during the refinement phase.

Proposed is a constructive framework for distributed, lossless source coding of two binary sources that have arbitrary correlation structures. The proposed framework accommodates the important special case of the absence of any correlation between the two sources, wherein it becomes an entropy coding scheme for a single source. The proposed framework was developed by combining Low-Density Parity Check (LDPC) codes with the DISCUS framework developed by Sandeep Pradhan. The combined algorithm was found to be sufficiently powerful to attain the Slepian-Wolf bound for two memoryless binary sources as well as the entropy rate for a single memoryless binary source. We are looking into aspects of rate adaptation for multiple (more than two) sources.

Duality Theory

The notion of functional duality was developed by Sandeep Pradhan for "one-sided" side-information point-to-point source and channel coding problems and then extended to more instances of multiple-input-multiple-output (MIMO) source and channel coding problems, admitting different scenarios of collaboration among multi-terminal inputs and/or outputs. The collaboration scenarios considered involve those where either the multi-terminal encoders or the multi-terminal decoders can collaborate, i.e. be joint, but not both. (The case of collaboration at both ends degenerates to point-to-point MIMO systems.) Under this one-sided collaboration abstraction, four problems of interest in source and channel coding were addressed: 1) distributed source coding, 2) broadcast channel coding, 3) multiple description source coding with no excess sum-rate, and 4) multiple access channel coding with independent message sets. In 1) and 4), the decoders collaborate, whereas in 2) and 3), the encoders collaborate. These four problems have been studied in the literature extensively. Precise mathematical conditions under which these encoder-decoder mappings are swappable in the two dual MIMO problems have been developed. We have also identified the key roles played by the source distortion and channel cost measures, respectively, in the MIMO source and channel coding problems in capturing this duality. Study of functional duality serves two important purposes:

- (i) it provides new insights into these problems from the different perspectives of source and channel coding, and allows for cross-leveraging of advances in the individual fields;
- (ii) more importantly, it provides a basis for sharing efficient constructions of the encoder and decoder functions in the two problems, e.g., through the use of structured algebraic codes, turbo-like codes, trellis-based codes, etc.

Security

When a data source is to be transmitted across an insecure, bandwidth-constrained channel, the standard solution is to first compress the data and then encrypt it. We examined the problem of reversing the order of these steps, first encrypting and then compressing. Such a scheme could be used in a scenario where the data generator and the compressor are not co-located, and the link between them is vulnerable to eavesdropping. We have been developing constructive encryption-compression schemes based on the DISCUS framework that was developed earlier for distributed compression.

Routing and Adaptation for Improving Energy Efficiency of Sensor Networks

We have developed a routing algorithm for reducing the amount of energy spent in communication setup and control for single-destination, energy-constrained wireless networks. The energy-efficient routing algorithm dubbed, Data Funneling, uses various packet aggregation ideas to provide significant energy savings. Packet aggregation strategies have the added benefit of decreasing the probability of packet collisions when transmitting on a wireless medium. Additional savings were realized through efficient data compression. This was done by encoding information in the ordering of the sensor packets. This coding by ordering scheme compresses data by suppressing certain readings and encoding their values in the ordering of the remaining packets. All these techniques together were found to more than halve the energy spent in communication setup and control. A novel, low-complexity algorithm for reducing energy consumption in sensor networks, using distributed and adaptive signal processing principles, was developed. An adaptive filtering framework was used to continuously monitor and learn the relevant correlation structures in the sensor data. In simulations, sensor nodes were configured for doing "blind", distributed compression of their readings with respect to one another, without the need for explicit and energy-expensive, inter-sensor communication to effect this compression. Simulation results revealed significant energy savings (from 15%-40%) for typical sensor data corresponding to a multitude of sensor modalities.

Routing with Long Links

Multi-hop wireless networks, also known as ad hoc networks, have become increasingly popular in recent years. The nodes in these networks are deployed with no infrastructure and are usually mobile. Examples of ad hoc networks are, but not limited to, sensor networks, wireless LAN, and environmental monitoring. In most of these ad hoc networks, the nodes have limited power, and therefore it is crucial to minimize the power consumption in performing computations and communication. There have been a number of routing protocols proposed to minimize the energy consumption in routing packets from source to destination. However, most of these protocols are very specific, in the sense that they are designed for networks with specific functionalities. Most of these routing algorithms cannot be used in more general cases since the assumptions in designing the protocols are very limiting. The second and more important drawback of these protocols is that they assume a symmetric disk model for transmission range, with all the links inside the disk being 100% reliable. However, it has been observed, based on the empirical data collected from a group of sensors at UC Berkeley that the transmission range is not a symmetric disk, and there exists long range, unreliable links among nodes. We have been working on a routing protocol which is power efficient and, at the same time, tries to decrease the expected time of delivery of packets. We consider a random network in which each pair of nodes is connected according to some probabilistic function of the pair-wise distance from each other. This model gives a more realistic view of the network and accounts for the unreliable, long range links. Currently, we are working on the simulation and implementation of the protocol. Once this step is successfully finished, we intend to perform mathematical analysis of the algorithm to see if the result of the simulation and implementation is consistent with theory.

Energy Metrics for Sensor Networks

Material published in the International Journal of Parallel and Distributed Sensor Networks, December 2001, in the paper "Energy and Performance Considerations for Smart Dust" was the beginning of our study of how energy is consumed. This paper showed how each of communication, computation, or sensing can dominate energy expenditure in sensor networks. In most scenarios, however, transmission over the wireless channel is the biggest drain on network resources. Research has subsequently focused on quantifying how message passing through the network can be optimized.

Results so far have been both analytical and simulation-based. For example, in a multi-hop network, nodes further from the data-collecting base station are more costly to retrieve data from, than those close by. By penalizing transmission from distant sources, a distortion-minimizing scheme can be developed for a given allowance of message density. This optimization can be performed analytically in an asymptotic setting, but simulation is required to determine how far realistic networks deviate from this ideal.

The overall goal of this work is to find out what determines the total energy cost of data retrieval. With compression possible in the network, what are the qualities of the underlying data-generating field that make some scenarios more costly to monitor than others? By determining how to quantify the complexity of a field being monitored by the sensor network, bounds on performance can be realized. The establishment of accurate definitions of complexity and the corresponding energy metrics is ongoing work.

The Ivy Project

The goal of the Ivy project is to develop and implement algorithms to extend the lifetime of wireless ad-hoc sensor networks. Data collection from sensor nodes has been an energy-intensive endeavor. For example, the "Mica" sensor nodes, developed at UC Berkeley and used for the project, survive for only five days when running the full duty cycle typically used for data collection. Ivy is a framework for reducing the duty cycle of nodes while maintaining a steady stream of data collection in a multi-hop network. With the Mica nodes, the goal is to extend lifetime to a year on a pair of AA batteries. This is an allowance of 1-2% duty. To date, the algorithms have been implemented in small networks of up to 15 nodes and three hops away from the base station. A thorough discussion of the algorithm and its implementation in TinyOS on the Mica nodes has been submitted as a Technical Report in the CS Division at UC Berkeley.

Reducing the duty cycle is accomplished by establishing a time-division multiple access (TDMA) schedule for the radio channel. Nodes are assigned slots during which they can forward data upwards towards the base station or receive data from their child nodes. Slots not assigned to any particular node allow this node to sleep and conserve energy during this period. To meet our energy budget, nodes cannot be scheduled to transmit or receive for more than a small fraction of these slots.

Slots are assigned through a distributed algorithm that runs continually during network operation. There is no dedicated start-up phase for the network: the base station begins by sending "advertisement" messages for empty slots in its schedule. Such advertisement continues as nodes hearing the advertisement probabilistically respond to the base station and secure upstream communication links. The hidden node problem is overcome by a two-step process to secure any slot which allows for slot reuse in disjoint areas of the network. Once a node secures a link with the base station, it begins to advertise its empty slots to nodes further downstream. This process makes joining the network transparent to new nodes as they can simply listen for a frame of the algorithm for advertisements of upstream data collecting slots.

Global synchronization is not required. Local synchronization is accomplished by each child node resetting to the parent's clock immediately following upstream data transmission. This timing message also serves to immediately acknowledge the receipt of the child's packet.

As new nodes join the network, links will fail due to individual node failures and changing RF environments. Link failure is handled by nodes listening for other parents that are advertising empty slots just as is done when first joining the network and is recognized by repeated lack of acknowledgements from a parent node. A node close to the base station that loses connectivity can be costly in terms of time required to reestablish all the traffic that has been disrupted. Mobility at the leaves of the communication structure is anticipated and handled more efficiently.

The next logical step is to implement the Ivy algorithm on a larger network. Trade-offs between data latency and network size will result if the overall duty cycle is to be maintained. Providing that a similar node density results from the increase in node numbers, latency should be the only sacrifice as size scales.

Localization in Sensor Networks

Effort-1:

Our goal in studying the problem of localization of nodes in a wireless ad hoc network was to design an algorithm which would: 1) minimize computation so as to be implementable on the current nodes running TinyOS; 2) be distributed (or localized) in the sense that each node estimates its own position based only on local information obtained from its neighbors; 3) be mathematically analyzable, so that we can answer basic questions like "What density of nodes is necessary to achieve a certain degree of accuracy with a given confidence?"

We proposed such an algorithm in [1] and [2] under the assumption that some nodes (called beacons) know their position. Each node communicates to its neighbors, obtains their positions, and estimates its own position by computing the intersection of the communication regions of the neighbors.

We achieved all the goals listed above, but to make the algorithm suitable for mathematical analysis, we had to make a modeling sacrifice. Namely, the communication region of each node is assumed to be a square and is therefore unrealistic. More precisely, we assumed that nodes lie in a grid and that two nodes can communicate with each other if their grid (or Manhattan) distance is smaller than some specified radius. On the positive side, we were able to compute the expected value of the estimate and the probability that the size of the estimate is ideal (i.e., equal to one cell in the grid). This enabled us to give a lower bound on the number of beacons necessary to achieve a given degree of accuracy with a given confidence.

To make the algorithm more realistic we need to extend it to a more general signal attenuation function and include obstacles in the region of deployment. This in turn will make the mathematical analysis and obtaining precise estimates much more difficult. This will be done in future work.

Unfortunately, practical localization has still not been satisfactorily solved. This is due to the lack of robustness of the current ranging methods to random disturbances or attacks on the network, as well as the absence of a sufficiently simple (in the sense of ability to implement on the current sensor network platform), truly distributed algorithm which would be sufficiently accurate.

Effort II:

Nodes in low-power sensor networks are limited in energy expenditure and device cost, and consequently are unable to self-localize using GPS. Several strategies for localizing nodes in a network have been proposed. We proposed a general convex programming approach that was published at Infocom 2001 in a paper entitled "Convex Position Estimation in Wireless Sensor Networks". This work showed how knowledge of pair-wise distant constraints can be combined into a global constraint set. It also illustrated a few cases in which the intersection of such constraints results in easily solvable problems using convex programming. In particular, the intersection of circular regions results in a second-order cone problem that can be efficiently solved. With a small number (~5-8) of known node locations, the remaining locations could be determined to uncertainties much less than the area of the original pair-wise constraints. Because of the convexity requirement however, node estimates are always placed within the convex hull of the known node locations.

The formulation was simplified for a discretized distributed setting in which convergence properties were studied. Instead of intersecting circular regions for example, rectangular regions were used, instead with boundaries lying along grid lines. The intersection of such regions requires nothing more than taking maxima and minima of the potential resultant rectangle. As the density of nodes increases to infinity, it was shown that all nodes can be localized to the finest resolution, i.e. the area of the grid mesh used to define the rectangles. This methodology also is more satisfying in the realm of sensor networks as collection of the data at a centralized location for the solution of the global convex problem can be costly in terms of energy and bandwidth consumption.

An undergraduate project which took a different approach to a centralized solution of the problem was also supervised. Assuming that pair-wise distances were known, a multi-dimensional scaling (MDS) algorithm can place nodes accurately to within a rotation and a translation of their actual positions. While this system proved robust to significant (but unbiased) variation in each distance measurement and a handful of outlier data, the downside is that it requires all pairs of distances to be known. To overcome this obstacle, the algorithm was run with several missing data points with some success, but analysis of functionality, when the majority of distances were unavailable, was not carried out. A discretized approach in which the distances used were only hop-numbers between nodes allowed for a surprisingly accurate reconstruction of the global map.

The continuation of work in localization has been waiting for technological advancement in ranging. RF time-of-flight systems have appeared which would lend themselves to an

MDS-like approach or a solution using relaxation of non-convex constraints needed to "push" nodes away from each other in the global reconstruction. A system for measuring angles between mobile robots is being developed at UC Berkeley which might also profit from our work in bounded-angle convex constraints. A custom low-computation solution using linear programming is underway.

Effort-III: distributed position estimation

We are developing distributed algorithms on how to perform position estimations for sensor networks. We consider each sensor in a sensor network has an on-board communication module so that it can establish local communication connectivity with a set of neighboring sensors. If an unknown sensor is able to receive communication signals from a nearby beacon, it must lie in a disc centered at that beacon with the radius of the maximum communication range. On the other hand, if this sensor can receive the position information of some other beacons in its neighborhood, it must lie in the intersection of all these discs. Therefore, an outer-approximation of this intersection could be used as an estimation of the position of the unknown sensor. Every unknown sensor is capable of performing position estimation algorithms with its own computational power by using the received accurate positions of its neighboring beacons, and the estimated position can be stored in its own memory. The position estimations of the whole sensor network can thus be done in such a distributed fashion. We then use a bounding polytope to approximate the intersection of these discs. Position estimation algorithms are performed in a sequential manner. To be more specific, we first find a series of polytopes to cover the intersection of discs pairwise. Then, for the obtained polytopes, we utilize a new series of polytopes to outer-approximate the intersection of these. By iterating this procedure, we can finally obtain a single polytope that outer-approximates the intersection of all the discs. Due to the nature of the iteration, the unknown sensor must lie in this polytope. The advantage of the sequential outer-approximation procedure is that it avoids dealing with all the discs simultaneously, which significantly reduces the computational loads.

Distributed map building and navigation

In this research we try to study how the sensor network formed by the sensors aboard a group of mobile robots can help the robots to navigate and build a map of an unknown environment that may contain obstacles of arbitrary shape at unknown locations. Traditionally, the map building task is performed by a single robot controlled by a central controller. By employing a group of robots with on-board sensors that can communicate with each other, it is hoped that one can not only speed up the map building process, but can also improve its accuracy and robustness with respect to mechanical and communication failures. We have designed a distributive algorithm, coordinating the motions of robots to ensure that they are collision-free and that under-explored regions are explored with priority. The algorithm has been tested through computer simulations with satisfactory performance. Being a distributed algorithm, the performance easily scales with the number of robots; thus it is suitable for real-time implementation. Our

task in the next stage is to test its real-world effectiveness on the Unmanned Ground Vehicles (UGV) platform, currently available for our group at the Richmond Field Station.

Sensor Networks for Multiple Target Tracking

We are trying to use sensor networks for tracking multiple moving objects on a terrain. The idea is to use this information for improving pursuit-evasion scenarios, where multiple pursuers try to catch multiple evaders. Current algorithms rely on probabilistic map-building of the terrain and the position of the evaders, using information from the sensors, such as cameras and ultrasound range finders, which are onboard pursuers. The problem with this approach is that it is very computationally expensive and does not provide deterministic performance guarantees. Our idea is to try to estimate the positions of moving objects by deploying a large numbers of devices with magnetometers that can detect their presence if they are sufficiently close. This approach removes uncertainty from the map-building process and can significantly reduce the time-to-capture of evaders. Over the course of this project, we built software simulations for such scenarios and developed several algorithms to collect the measurements, aggregate the data, and transmit it to a base station, where the motion of objects is reconstructed. We also built a physical testbed consisting of a 10 X 10 grid of magnetometers spaced approximately 1.5 meters from each other. There are several problems that need to be solved to make this approach efficient. Some of these are theoretical and others are practical. For example, the information given by the magnetometers is very crude and not as informative as that obtained by cameras. In fact, the information provided is just the magnetometer's signal strength. Only in combination with the information of several adjacent sensors can one try to reconstruct the motion of the object by performing triangulations. This approach however, fails in the presence of two or more objects since it is not possible to disambiguate the sensor readings.

Multiple Hypothesis Tracking (MHT) theory provides an algorithm that addresses this issue, but it has exponential complexity. We have developed heuristics to reduce the complexity of the problem. A centralized (greedy) version has been implemented which is working quite well and is very fast. While the centralized algorithms we developed successfully deal with the problem of multiple targets, it is computationally expensive and we are currently trying distributed implementations to make it real-time and energy-efficient for the pursuit game. Other problems arise from non-idealities of hardware. The magnetometers we use are off-the-shelf, inexpensive devices which suffer from calibration problems, i.e., sensors placed at the same location give quite different sensor readings. Even when initially calibrated, drifts skew the reading over time and the estimation performance degrades dramatically. Also, some sensors failed and started sending false readings that biased the position estimates of the object. Another major problem is related to the timing information. Delays in packet delivery from the sensor to the base station greatly influence routing protocols, which are very difficult to model and simulate. While our object position estimates are accurate, the delay makes

that information useless from a practical viewpoint, because the evader would have already moved somewhere else and therefore would be difficult to chase.

To summarize, the theory and practice turned out to be different due to a series of unexpected problems that had to deal with hardware non-idealities, sensor failure, and data aggregation, routing, and delay of received packets. More accurate models of routing and sensor non-idealities need to be studied and their effect on the estimation algorithms need to be assessed. This is the direction in which we are currently proceeding: designing robust, real-time, scalable and failure-tolerant algorithms.

Environmental Monitoring

We focused on estimation of the gradient of a scalar field present in the region of deployment of a random sensor network. Why this particular problem? The significance of having a good estimate of the gradient of a scalar variable is that it reveals the rate of its change. In the example of monitoring temperature, this means knowing the heat flow. This is of great importance in detecting and fighting forest fires. We proposed a distributed algorithm for gradient estimation in [3] and [4]. In the algorithm, each node talks to its neighbors and computes the direction of the greatest increase of the scalar field V , given the information gathered by the neighbors. This is the direction of the gradient of V . We were able to estimate the confidence with which the algorithmic error is smaller than a given threshold and give a lower bound on the number of nodes sufficient for such confidence. The downside of our analysis is that it does not take into account node failures and noise. This will be addressed in future work.

Interacting particle systems models in sensor networks

One of the applications sensor networks are useful for is tracking. The network as a whole can be interrogated periodically to recover the state of the object being tracked. While the object is being tracked, the sensors, being computationally limited, may lose track of it, but lock on the target can be refreshed by messages from neighboring sensors. Since there is a cost to communication, the problem of designing the architecture of refresh messages is an important one. We studied this problem using a particle systems model called the contact process. Our problem becomes one of optimal design of a contact process. A key conceptual discovery was made in this work. We demonstrated the role of phase transitions in the optimal solution. The optimal solution is spatially inhomogeneous, because of a phenomenon of symmetry breaking.

Error exponents for distributed detection and security

We have been pursuing an investigation of the error exponents for distributed detection by an array of sensors, when there are uncertainties at the fusion center as to the kind of sensors the data is coming from. We are interested in contrasting this with the known form of the error exponents, when the sensor types are known to the fusion center. Partial results have been obtained and we are continuing to work on extending this result to build up a comprehensive body of results.

We have made significant progress on the study of the error exponent of timing channels. This problem is of significant interest in establishing security guarantees in networks, and specifically in sensor networks. This is because it is possible to carry out covert communications through timing channels in the system. The reliability exponent translates directly into the block sizes needed to communicate at a given data rate (the higher the reliability, the smaller the needed block length). Since tests for the presence of covert channels require watching the channel for periods of time (they are generally based on sequential probability ratio approaches), the reliability exponent is what governs the data rates at which covert communication is a real possibility in an environment with security checks. We have determined the reliability exponent of the timing channel associated to the exponential server queue, at zero rate. This work has involved developing some novel techniques for estimation from point process observations. Somewhat surprisingly, one of the corollaries is that the exponential server queue is more reliable as a timing channel than the Poisson intensity modulated channel. We have used this to give a straight line upper bound for the reliability exponent at all rates (this requires interpolating between the zero rate error exponent and the sphere packing bound).

Game-theory in communication networks

We have studied a cooperative game theoretic formulation of the rate allocation problem in a Gaussian multi-access channel. We find that there is a unique allocation, which is feasible and in the core of the game, which satisfies certain natural envy freeness assumptions. The multi-access channel is a basic model for wireless up link communication. This work provides a deeper understanding of fairness issues in such uplinks.

Professor Anantbram presented an invited plenary talk at the Wireless Optimization conference (WiOpt03) where he proposed several ideas for the investigation of sensor networks using asymptotic control theory techniques. Several ideas, such as symmetry breaking, mean field methods, and common randomness, are beginning to be explored towards this end. Other ideas related to the role of common randomness in sensor networks, including developing a notion of distributed game theory for the design of sensor networks in adversarial situations, was proposed at an invited talk at a symposium in Bielefeld in honor of the eminent Information Theorist Professor Rudolph Ahlswede.

Several problems related to vector multi-access channels have also been investigated. This included work on CDMA, as well as multi-antenna channels. One PhD. Thesis, as well as several conference papers and journal papers, resulted from this work. This work is of fundamental interest in understanding the performance of wireless sensor networks. We have also been working on several problems related to high performance coding and decoding that are of broad interest and are also likely to be of interest in the design of sensor networks. This includes a generalization of the Bayesian belief propagation paradigm that led to several novel algorithms for estimation. Also, several suboptimal decoding algorithms were designed with architectural implementation constraints in view. This work resulted in one PhD. thesis, is of considerable interest to industry, and has also been presented in several conferences and publications.

Learning theory in Sensor Networks

Fundamental role:

Localization, environmental monitoring, and almost any other application of sensor networks can be viewed as an instance of supervised learning, or learning from examples. A sensor node trying to estimate a scalar function defined in the region of deployment is just trying to "learn" this function from examples represented by the sensor data collected by the node itself and its neighbors. The advantage of this point of view is that the theory of machine learning has developed powerful and well understood algorithms which can potentially be employed in the context of sensor networks. This is the main thesis of [5]. There, we interpret a well-known algorithm from learning theory (due to T. Poggio and others) in the context of localization, environmental monitoring, plume tracking, and tracking of moving objects. Two main challenges are optimizing the amount of computation to be distributed to the nodes, (as opposed to it being done centrally), and merging the local estimates into a global one. These challenges have not yet been solved in a satisfactory way and are the topic of our future work.

Sensor field model learning with application to blind separation of linearly mixed signals:

Independent Component Analysis (ICA) is a technique for finding a linear transformation that makes the data components as independent as possible, with as few assumptions as possible, on the signals. It has been successfully applied to many problems where it can be assumed that the data is actually generated as linear mixtures of independent components, such as audio blind source separation or biomedical imagery. However, new application areas (e.g. music, aerospace imagery) are emerging that require a relaxation of the assumption of independence, while keeping the linear mixing assumption.

In order to allow for dependence between the recovered components, we have developed a generalization of ICA, where instead of looking for a linear transform that

makes the data components independent, we look for a transform that makes the data components closely fit to a tree-structured graphical model. See reference [6]. To estimate the linear transform and the tree structure, we have successfully adapted classical ICA estimation techniques to this new model. In particular, TCA allows the underlying graph to have multiple connected components and thus, the method is able to find "clusters" of components such that components are dependent within a cluster and independent between clusters.

References

- [1] S. N. Simic and S. Sastry, Distributed Localization in Wireless Ad Hoc Networks, Technical Memorandum No. UCB/ERL M02/26, UC Berkeley, 2002.
- [2] S. N. Simic and S. Sastry, A Distributed Algorithm for Localization in Random Wireless Networks, submitted to Discrete Applied Mathematics, 2002.
- [3] S. N. Simic and S. Sastry, Distributed Gradient Estimation Using Random Sensor Networks, Technical Memorandum No. UCB/ERL M02/25, 2002
- [4] S. N. Simic and S. Sastry, Distributed Environmental Monitoring Using Random Sensor Networks, Proceedings of the 2nd International Workshop on Information Processing in Sensor Networks, IPSN'03, PARC, Palo Alto, CA, April 2003
- [5] S. N. Simic, A Learning Theory Approach to Sensor Networks, invited paper to the special issue of IEEE Pervasive Computing on Sensor and Actuator Networks, D. Estrin and G. Sukhatme (editors), July 2003
- [6] Beyond Independent Components: Trees and Clusters. F. R. Bach and M. I. Jordan. In press: Journal of Machine Learning Research, 2003.