

Monte Carlo Sensor Networks

Thomas C. Henderson, Brandt Erickson,
and Travis Longoria
School of Computing
University of Utah
Salt Lake City, UT, USA
Email: tch@cs.utah.edu

Edward Grant, Kyle Luthy,
Leonardo Mattos, and Matt Craver
Electrical and Computer Engineering
North Carolina State University
Raleigh, NC, USA
Email: egrant@eos.ncsu.edu

Abstract—Biswas et al. [1] introduced a probabilistic approach to inference with limited information in sensor networks. They represented the sensor network as a Bayesian network and performed approximate inference using Markov Chain Monte Carlo (MCMC). The goal is to robustly answer queries even under noisy or partial information scenarios. We propose an alternative method based on simple Monte Carlo estimation, and which is more time efficient on low numbers of sensors (e.g., 10-20) which are more commonly found in subnet clusters; our method allows a distributed algorithm, pre-computation of probabilities, a more refined spatial analysis, as well as desiderata for sensor placement in the friendly agent surrounded by enemies problem. In addition, we performed experiments with real microphones and robots to determine the sensor correct response probability.

I. INTRODUCTION

Many advances have been made in sensor network technology and algorithms in the last few years. See [2] for an overview of the state of the art. Work has been done on: architecture [3], systems and security [4], [5], [6], and applications [7]. Our own work has focused on the creation of an information field useful to mobile agents, human or machine, that accomplish tasks based on the information provided by the sensor network [8], [9], [10], [11], [12].

Some drawbacks of sensor networks include the need to conserve power and not run all the nodes all the time (partial data), and sensors are noisy (sometimes return the wrong value). This motivates a statistical approach to decision making. As described in the abstract, Biswas et al. [1] introduced an interesting problem in the context of sensor networks, as well as an MCMC solution. We present an alternative approach here.

A. The Problem

Suppose there is a 2D area of interest (we'll consider the unit square), and a friendly agent located

at a fixed position, LF , in the area. In addition, there are m sensor nodes located throughout the area (these sensor nodes report the presence or absence of enemy agents within some fixed range). Finally, n enemy agents are placed in the area; each enemy agent's coordinates are chosen by independently sampling a uniform distribution for x and y .

Query: Given a set of sensor responses, R_i (each declares there is or is not an enemy within its range), and a set of probabilities of the correctness of the responses, what is the probability, $P(S | R_i)$, that the friendly agent is surrounded by the enemy agents (surrounded means that it is within the convex hull of the enemy locations).

It should be noted that although not cited by Biswas et al., there is a good bit of knowledge about this problem in the stochastic geometry literature. For example, Santaló [13] is a prime source on stochastic geometry and discusses several related problems; e.g., the expected area and perimeter of the convex hull of n random points, the expected value of the number of vertexes, etc. Mathai [14] gives a thorough discussion of the probability that a fixed point P is inside a random triangle inside a given triangle. That discussion is based on the work of Halász and Kleitman [15] who give the details of the estimation formula. Given this, it would have been better if Biswas et al. had selected a triangular area so that the estimates could be checked against theory; we intend to pursue this in the future.

B. Previous Work

The authors cited above set up a Bayesian network with a given a priori probability of correctness of the sensor responses, and conditional probabilities on the locations for the enemies given the responses, and also a conditional probability of being surrounded given the enemy locations. They present results of answering the query for two examples, and investigate

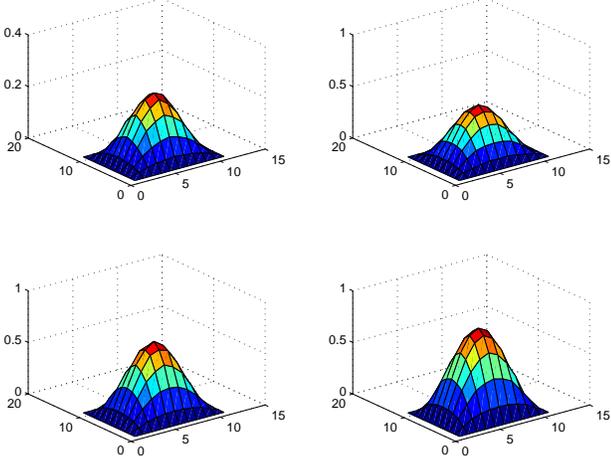


Fig. 1. A Priori Probability that Point in Unit Square is Surrounded (for 3,4,5 and 6 enemies).

the affect of various parameters on the posterior probability; these parameters include: number of sensors, number of enemies, sensing range and MCMC convergence.

II. MONTE CARLO SENSOR NETWORKS

A. Sensor Model

We take a slightly different sensor model than Biswas et al. Let the ground truth be defined as:

$$\hat{d}_i = \begin{cases} 0 & \text{if there's no enemy in range of sensor } i \\ 1 & \text{otherwise.} \end{cases}$$

In the Biswas et al. model:

$$d_i = \begin{cases} \hat{d}_i & \text{with probability } \zeta \\ 0 & \text{with probability } \frac{1-\zeta}{2} \\ 1 & \text{with probability } \frac{1-\zeta}{2} \end{cases}$$

Evaluating these expressions (given that \hat{d}_i is either 0 or 1), leads to a limited model in which d_i is correct at least 50 % of the time (i.e., when $\zeta = 0$). This follows since d_i is correct with probability $\zeta + \frac{1-\zeta}{2}$ which is $\frac{1}{2}$ when $\zeta = 0$.

We propose to let:

$$d_i = \begin{cases} \hat{d}_i & \text{with probability } \rho \\ -\hat{d}_i & \text{with probability } 1 - \rho \end{cases}$$

B. No Sensor Data Available

In the absence of any sensor data, there is an a priori probability that the friendly agent is surrounded. We use Monte Carlo simulation to estimate this probability (see Figure 1 for $n = 3 \dots 6$). Figure 2 shows the a priori probability that the center point (0.5,0.5) is surrounded as the number of enemies goes from 3 to 10.

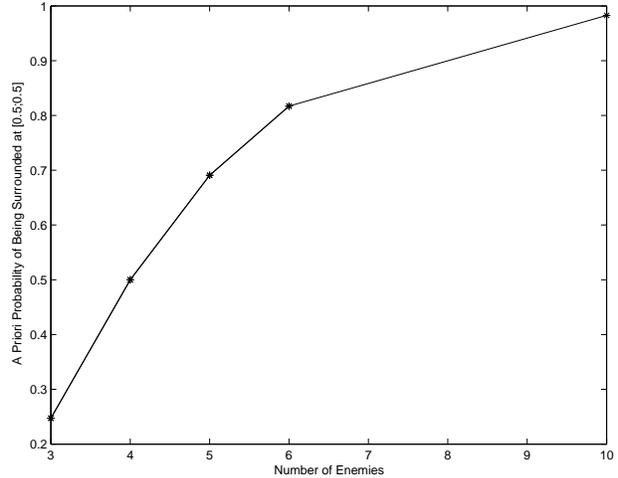


Fig. 2. A Priori Probability that Point in Center of Unit Square is Surrounded (for 3,4,5,6 and 10 enemies).

C. One Sensor Case

To determine the probability when sensor data is available, we use Monte Carlo as follows. Consider the case when 1 sensor responds; i.e., $R := d_1$ is the sensor response and is correct with probability ρ . Then:

$$\begin{aligned} Prob(S | R) = & \rho MC(\text{assume } d_1) \\ & + (1 - \rho) MC(\text{assume } \neg d_1) \end{aligned}$$

where $MC(\cdot)$ means to run the Monte Carlo simulation with condition \cdot set. That is, if d_1 is 1, then fix one enemy at LS_1 (sensor 1's location); if d_1 is 0, then don't use any sample with an enemy in range of LS_1 .

D. m-Sensor Case

The same approach is used, but now there are 2^m terms in the probability calculation. Let $\mathbf{d} = [d_1, d_1, \dots, d_m]$ be the set of sensor node assertions about the presence of enemies. We must consider all combinations of these being true or false. This is characterized by a vector $\mathbf{b} = [b_1, b_1, \dots, b_m]$ and the contribution of each to the total probability of being surrounded:

$$Prob(S | R) = \sum_{\mathbf{b} \in \mathcal{P}(\{0,1\}^m)} Prob(\mathbf{b}) Prob(S | \mathbf{b})$$

For example, if we assume that all the sensor nodes report that there is an enemy present, and the first $m - 1$ sensor responses are correct but that sensor node m is wrong, then we compute that term in the above summation as:

$$\rho^{m-1} (1 - \rho) Prob(S | Positions)$$

where *Positions* indicates that Monte Carlo is run assuming enemies in positions indicated by nodes 1 to $m - 1$ and no enemy in range of node m .

The calculation of these probabilities can be made quite efficient by using every sample generated. It suffices to keep two counters for each set of correct/incorrect sensor combinations: (1) a counter for the number of times surrounded for a sample from that combination, and (2) the total number of samples for that combination. Thus, the sequence of actions is:

- generate a set of n random enemy locations,
- determine to which combination it corresponds,
- increase total sample count by one for that combination, and
- if the friendly is surrounded by this sample, then increase the surrounded count by one for this combination.

III. EXAMPLES

We give here our results for the experiments of Biswas et al. (We estimate the locations from their figures!) The layout of Experiment 1 is shown in Figure 3. The corresponding table is given in Table

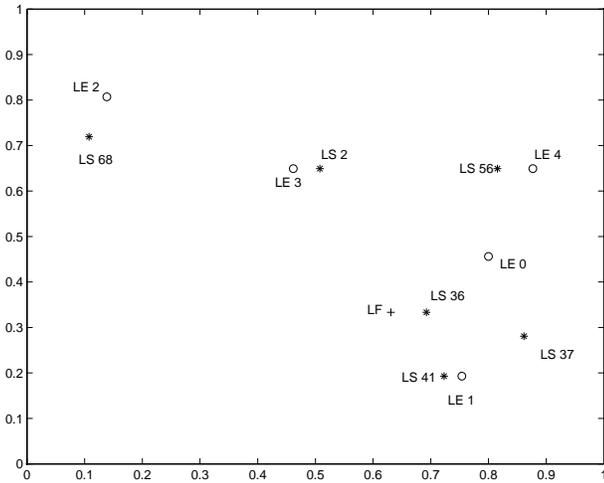


Fig. 3. Biswas et al. Experiment 1 Layout of friendly agent (+), enemy agents, LE (o) and sensors, LS (*).

1; we give both our results (MC) and theirs (MCMC) for comparison.

R	MC	MCMC
56, 68	37	44
none	49	51
2, 41	87	80
41	82	81
41, 56, 68	99	86
36, 37, 41, 56, 68	99	93

Table 1. Results from Experiment 1.

Figure 4 shows the layout of Experiment 2, and Table 2 shows the results obtained from both the MC and MCMC methods.

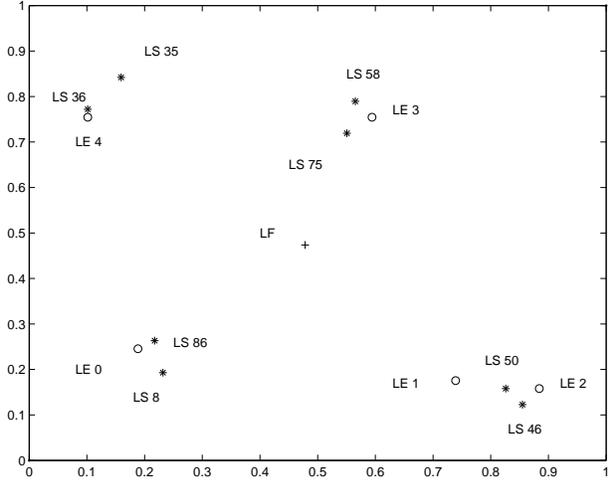


Fig. 4. Biswas et al. Experiment 2 Layout of friendly agent (+), enemy agents, LE (o) and sensors, LS (*).

R	MC	MCMC
none	68	67
35, 50, 58	66	70
8, 36, 58	66	74
8, 36, 46	74	75
8, 35, 50, 58	87	78
8, 35, 36, 46, 50, 58, 75, 86	96	91

Table 2. Results from Experiment 2.

Figure 5 shows the convergence of the Monte Carlo process as the number of trials is increased. These results are for row 2 of Table 2, and for the number of trials ranging from 10^0 to 10^5 .

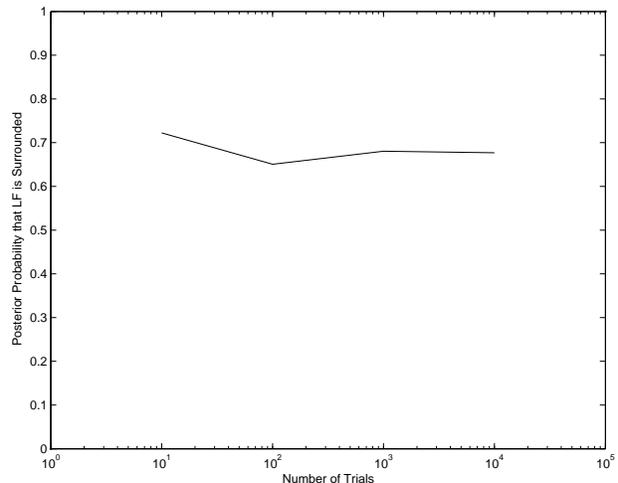


Fig. 5. Convergence of MC for row 2, Table 2.

We also ran a simulation experiment in which the sensor responses' correctness was sampled from the appropriate distribution (uniform, 70% correct; see physical experiments section). We ran one case with the friendly agent surrounded and one case with it not surrounded. (We used Experiment 2 layout for the former, and switched LF and LE 3 for the latter.) Table 3 gives the mean posterior found and the variance (the no sensor case is eliminated).

Surrounded Case (Mean; Variance)	Not Surrounded Case (Mean; Variance)
(0.6776;4.7008e-005)	(0.4216;2.0093e-004)
(0.6904;6.2806e-005)	(0.3985;3.3633e-004)
(0.7298;1.0097e-00)	(0.3019;1.2547e-004)
(0.7629;2.3356e-004)	(0.3571;2.1346e-004)
(0.8418;5.5877e-004)	(0.1093;2.7538e-004)

Table 3. Re-run Experiment 2 with sensor correctness sampling and both with LF surrounded and not surrounded.

IV. PHYSICAL EXPERIMENTS

A set of experiments was undertaken to better understand the physical nature of acoustic detection systems, as well as their probability to produce a correct report of enemy presence. We determined the correctness probability to be around 70%.

For physical testing it was decided to use the EvBot II autonomous mobile robot testbed in the Center for Robotics and Intelligent Machines (CRIM) at North Carolina State University [16]. The autonomous mobile robot platforms in the testbed, the EvBots, are used widely for evolutionary algorithm and distributed sensor experimentation. To test the Mote Carlo sensor network at the heart of this paper a simple microphone circuit was constructed to provide the data for hit or miss decision making. For these experiments the EvBot II robots served as stationary sensor nodes. The enemy in this case is a speaker providing sound within a set range. The microphone circuit consists of a two stage amplifier, the output of which is rectified to yield a DC approximation of the amplitude. A DC approximation removes the need for rapid sampling and frequency determination and gives an integer value after analog to digital conversion. An example of the output of this circuit is shown in Figure 6. There is a slight voltage drop (approximately 0.7V) between the sound source and the rectified signal due to the rectifying diode. This serves to reduce the effective range of the ADC measurement (8 bits) from 1024 to around 930.

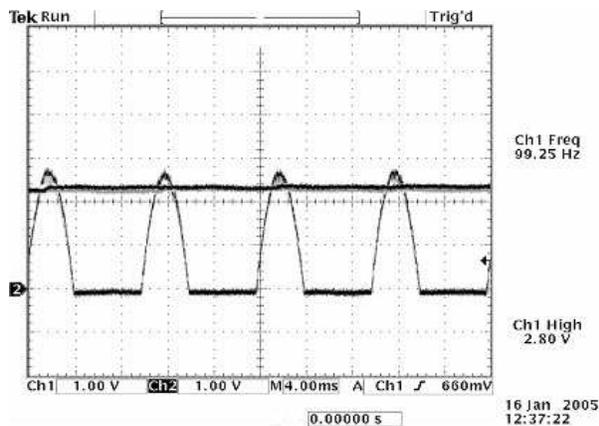


Fig. 6. The Sensed Sound and the Rectified DC Approximation.

The output of the microphone circuit is sampled through an ADC channel that is part of the circuitry of a BasicX 24 microcomputer; which in turn is part of the EvBot II platform. Multiple robots can easily be used to represent a distributed network by utilizing their built in wireless networking. While not necessarily a valid representation of a true distributed sensor network (due to size, processing power, battery life, and cost), they serve as a good test-bed to provide proof of concept. An EvBot fitted with a sensor circuit is shown in Figure 7.

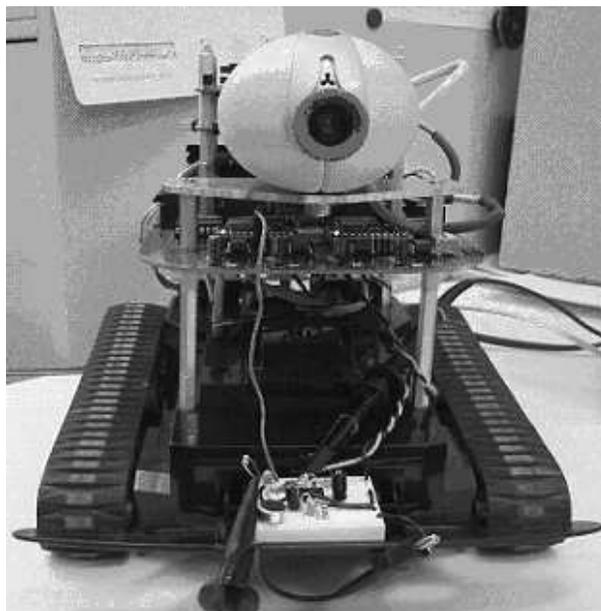


Fig. 7. EvBot II with Affixed Sound Level Detection Circuit.

To provide a listening range for the robot, a 100Hz sound source was placed 22.25 inches from the microphone. The sound level was set at an amplitude

value that equates to a sensor reading of 512, i.e., half the range of the ADC. Any fluctuations observed in the recorded data were considered as being a function of a combination of sensor circuit design and variations in the sound source. The sound was calibrated using a BK Precision model 732 digital sound level meter. For this setup, a level of 107.6dBC equates to approximately half ADC range at 22.25 in. This calibration was performed in the center of the environment to reduce error caused by noise reflection from the walls in the room.

A series of hit or miss experiments was performed in order to characterize the sensor. The EvBot was placed at three positions within the environment, the center, the center to one side, and in a corner, see Figure 8. At each of these positions the sound source was placed at 11.125 inches, 22.25 inches, and 33.375 inches, well within range, putting it just in range (defined previously as 22.25 inches), and well out of range. Likewise, at each range, data from 3 threshold values were recorded; one at 512 (used to define the original range), one at 256, and one at 768. Each measurement location and threshold is recorded 100 times. Furthermore, this experiment is performed on two different sensor/EvBot systems.

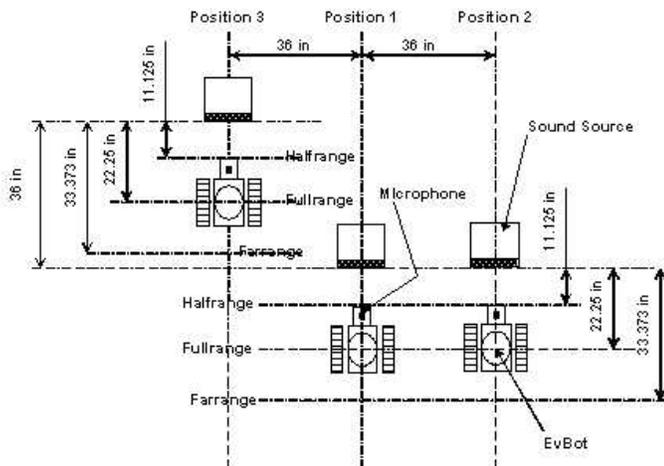


Fig. 8. Experimental Setup.

The sound level was measured at the sensor while at standard range in each of the three positions in the environment. While in the center, the amplitude was measured as 93.3dBC. At the side, this number is higher at 94.2dBC, and similarly, in the corner, 94.6dBC is measured. All three systems are 22.25 inches from the same source. This is an indication of sound level variation due to location, most likely surface reflections.

Examination of the data yielded the variation in hit and miss results due to the microphone sensor changing between the positions, ranges, and detection thresholds. To ensure that any variation in the sensor system is accounted for experimentally, two sets of experiments were carried out using two separate robots, i.e., EvBot 4 and EvBot 7. Whether or not there is a hit or a miss is determined by:

- 1) The position of the sensor within the environment (are there reflections off of walls?)
- 2) The threshold level used to determine whether or not an enemy is in the predefined range (a higher threshold will miss enemies in range, a lower will identify enemies out of range)
- 3) Variations between sensor systems.

The experiment proved that the percentage of correctness of hit or miss depends upon a variety of factors. When the system is set up such that the threshold is mid-range then the decision of hit or miss is predictable in cases where the sensor system is half-range and far-range, these give 100% hit and 100% miss, respectively. At the distance associated with the threshold, i.e., full-range, the percentage correctness obtained varies because of the variations in sound source, sensor, or ambient noise within the testbed. The experiments also proved that the position in the environment also plays an important role in the correctness of the results obtained, e.g., sound levels increase near a wall, due to sound reflection. This was verified using the BK sound meter.

V. DISCUSSION AND CONCLUSIONS

We have shown that a simple Monte Carlo approach gives good results in the presence of noise and partial information, and is comparable to the MCMC approach. In addition, our method is more time efficient if the number of sensors considered is kept low (e.g., 10-20) as is usually the case in a sensor network sub-cluster. Moreover, the Monte Carlo calculations can be performed off-line (once the positions of the sensors are known) and downloaded to them and a table lookup used to answer queries once the sensor responses are known. Alternatively, the computation is distributed among the sensor nodes if performed in real-time. Higher accuracy can be achieved by increasing the number of trials (note that this can be done incrementally – e.g., run ten 1000 trial experiments, then later run another ten, and keep averaging $P(S | R)$).

One assumption that we make in running the Monte Carlo experiments is that if multiple sensor nodes report the presence of an enemy agent, and if these



Fig. 9. 100 Telos B Mote System.

sensor nodes have a non-empty intersection, then we assume that there is just one enemy agent and it is placed in the intersection area. We intend to explore the significance of better area estimates for enemy location in overlapping sensor areas. This and other slight differences between the MC and MCMC methods accounts for the differences in specific values in the common experiments.

Our results indicate that it is also possible to get a good level of confidence by taking separate singleton or pairs of sensor responses, and use the max probability of those to answer the query. Moreover, it is possible, given the friendly agent's location, to determine sensor placements that decrease the number of responses needed to produce a high-quality answer.

In future work, we intend to explore these issues in a 100-node sensor network testbed now under construction (see Figure 9; the EvBots will provide the mobile agent platforms for use as friendly and enemy agents. Moreover, we are exploring the application of sensor networks to snow and avalanche monitoring[17], and in this case the goal might be to surround (e.g., the signal of a buried person) with mobile agents or people – thus, sometimes it may be of interest to ensure that an object of interest is surrounded.

We also intend to explore the use of this approach in large-scale forest fire scenarios. In such situations, the situation changes rapidly and unpredictably, and it may be that it is possible to produce a statistically accurate model for whether a fire fighter is surrounded by the fire, or will likely be surrounded by the fire in the near future.

Another possibility is to try to develop a model for lost persons, and use the method to attempt

to maximize the likelihood that the lost person is indeed surrounded, and then narrow the area of the search; see an accompanying paper in this volume on precision localization using Monte Carlo sensor networks.

Finally, as mentioned in the introduction, we intend to compare triangular area analytic results with both the Monte Carlo Sensor Network method and MCMC sensor networks.

REFERENCES

- [1] R. Biswas, S. Thrun, and L. Guibas, "A probabilistic approach to inference with limited information in sensor networks," in *LCSN*, (Berkeley, CA), April 2004.
- [2] F. Zhao and L. Guibas, "Preface," in *Proc of IPSN 2003*, (Palo Alto, CA), pp. v–vi, LNCS, April 2003.
- [3] J. Hill and D. Culler, "A wireless embedded sensor architecture for system-level optimization," *ece*, UC Berkeley, October 2002.
- [4] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly resilient, energy efficient multipath routing in wireless sensor networks," *Mobile Computing and Communications Review*, vol. 1, no. 2, 2002.
- [5] A. Perrig, R. Szewczyk, V. Wen, D. Culler, and J. D. Tygar, "SPINS: Security protocols for sensor networks," *Wireless Networks*, vol. 8, pp. 521–534, Sept 2002.
- [6] L. Zhang, "Simple protocols, complex behavior," in *Proc. IPAM Large-Scale Communication Networks Workshop*, March 2002.
- [7] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson, "Wireless sensor networks for habitat monitoring," in *WSNA 2002*, (Atlanta, GA), September 2002.
- [8] T. C. Henderson, M. Dekhil, S. Morris, Y. Chen, and W. B. Thompson, "Smart sensor snow," *IEEE Conference on Intelligent Robots and Intelligent Systems*, October 1998.
- [9] Y. Chen, "Snets: Smart sensor networks," Master's thesis, University of Utah, Salt Lake City, Utah, December 2000.
- [10] Y. Chen and T. C. Henderson, "S-nets: Smart sensor networks," in *Proc International Symp on Experimental Robotics*, (Hawaii), pp. 85–94, Dec 2000.
- [11] T. C. Henderson, "Leadership protocol for s-nets," in *Proc Multisensor Fusion and Integration*, (Baden-Baden, Germany), pp. 289–292, August 2001.
- [12] T. C. Henderson, J.-C. Park, N. Smith, and R. Wright, "From motes to java stamps: Smart sensor network testbeds," in *Proc of IROS 2003*, (Las Vegas, NV), IEEE, October 2003.
- [13] L. A. Santaló, *Integral Geometry and Geometric Probability*. Reading, MA: Addison-Wesley, 1976.
- [14] A. Mathai, *An Introduction to Geometrical Probability*. Amsterdam, The Netherlands: Gordon and Breach Science Publishers, 1999.
- [15] S. Hal'asz and D. Kleitman, "A note on random triangles," *Studies in Applied Mathematics*, vol. LIII, no. 3, pp. 225–2237, 1974.
- [16] L. Mattos, "The EvBot-II: An enhanced evolutionary robotics platform equipped with integrated sensing for control," Master's thesis, North Carolina State University, Raleigh, NC, May 2003.
- [17] T. C. Henderson, E. Grant, K. Luthy, and J. Cintron, "Snow monitoring with sensor networks," in *IEEE Workshop on Embedded Networked Sensors*, (Tampa, FL), IEEE, November 2004.