A New Approach to Range-only SLAM for Wireless Sensor Networks

John R. Spletzer

Lehigh University
Bethlehem, PA, 18015, USA
spletzer@cse.lehigh.edu

Abstract

Range sensors tend to frustrate estimators relying upon linearization techniques. Such approximations often lead to filter divergence, as demonstrated in [1, 2]. To address these shortcomings, we offer a new approach to range-only simultaneous localization and mapping (SLAM) for use in robot augmented wireless sensor networks. Our approach finds its roots in robust estimation, employing an unknown-but-bounded error model for range measurements. However, it also leverages recent advances in convex optimization theory to provide a framework suitable for real-time systems. It offers several advantages over alternate approaches, to include convergence and performance guarantees, as well as resistance to certain measurement outliers.

In this paper, we first extend our previous results in cooperative localization [3] to the range-only SLAM problem through semidefinite programming techniques. We then demonstrate localization performance in the presence of multi-path and signal attenuation errors associated with WSN range sensors. Lastly, we investigate convex approximation techniques for improving localization performance. Sample simulations are provided to support all results.

1 Introduction

Recent advances in sensors, embedded devices, and wireless communication technologies have converged to enable distributed micro-sensing through wireless sensor networks (WSN). In a WSN, each node is capable of sensing phenomena, collecting and processing data, and reporting/routing these data back to a data sink that in turn relays it to a manager node via inter-networks [4, 5, 6].

In order for a WSN to provide meaningful information regarding its environment, node positions must be well localized. This is typically accomplished via relative range measurements collected at each node. To facilitate the localization process, the WSN should be fully connected. Furthermore, each node should have a minimum connectivity degree of three or four (allowing for direct positions estimates), with node connectivity being associative. Such simplifying assumptions often appear in robotics and networking research literature. Unfortunately, they do not survive the tests of reality. Sensor nodes will be distributed in a somewhat haphazard fashion. Connectivity will be limited to within cluster islands of varying sizes. The connectivity for a given node could be 1 or 2 neighbors, and need not be associative - making the position estimation problem significantly more complex.

The solution then is to introduce dynamic capabilities into the network. Mobile robots provide such a medium. Each can be viewed as nothing more than an oversized sensor node, but with significantly greater computational resources, an augmented sensor suite, and the additional degrees of freedom afforded by mobility. By exploiting these, an augmented WSN can achieve otherwise unobtainable levels of robustness and performance.

The main contribution of this paper is a new approach for solving the range-only SLAM problem, whereby a team of mobile robots simultaneous estimate their own positions as well as the positions of WSN nodes. Conceptually, the idea is that range measurements induce constraints on the configuration space of the WSN. Merging these constraints induces a feasible set on the configuration space that represents the set of node positions consistent with all available sensor measurements. Estimates for the uncertainty in the absolute position of a single robot, or the relative positions of two or more nodes can be obtained by projecting this feasible set onto appropriately chosen subspaces of the configuration space.

This scheme offers several advantages over alternate approaches. Range measurements can be modeled exactly by ellipsoidal constraints. As such, it is not subject to failures in convergence associated with techniques based upon linearization. It also provides performance guarantees. At each time-step, we maintain an uncertainty region for the position of each node that is conditioned on all available sensor measurements. This uncertainty region is guaranteed to contain the true node position.

2 Related Work

In our approach, robots must estimate their own position while simultaneously localizing WSN node clusters. This

is known as the SLAM problem within the robotics community. SLAM has been well studied using sensors such as laser range-finders and sonar arrays that can directly estimate the positions of features in the environment [7, 8, 9]. Unfortunately in our network model, such measurements are unavailable. With only relative range information, the resulting unobservability makes range-only SLAM more difficult, and not surprisingly less studied.

An extended Kalman Filter (EKF) based approach for range-only SLAM was examined in [1]. Due to linearization errors, the filter would diverge if beacon positions were not known a priori to some error tolerance. Another recent approach examined range-only measurements for underwater SLAM [2]. The problem was modeled as a non-linear least-squares problem. However, the non-linear problem was solved via repeated solutions to a linearized system. As a result, this approach also failed to converge without reasonable initial estimates of beacon positions. The clear advantage of our method is that it will converge with no prior information as range measurements can be modeled exactly through quadratic constraints.

In contrast to Bayesian methods, our SLAM approach generates uncertainty regions that are guaranteed to contain the true node positions by employing an unknown-but-bounded error model for sensor noise. Such error models are used extensively in robust estimation. Early forays in this field date back to work by Bertsekas & Rhodes [10] and Schweppe [11]. However, most related to our approach is more recent work by Doherty & El Ghaoui [12]. In this, a similar method relying upon linear and semidefinite programming (SDP) techniques was used to estimate the positions of nodes in a wireless network from sensor measurements when the positions of some of the nodes were known *a priori*. This was accomplished by bounding the feasible set with a rectangle. Such an approach can yield an arbitrarily bad estimate of the uncertainty region. Our previous work built upon this by offering a method to recover the true uncertainty region arbitrarily well. Additionally, through the use of inside-outside approximation techniques, we proved a convergence bound [3]. In this paper, we apply these results to a range-only SLAM task by generalizing to convex constraints and incorporating a dynamic model. We also investigate the performance impact of measurement outliers and integrating nonconvex constraints using convex approximation techniques.

Finally, range-based localization strategies have also been proposed simultaneously within the networking community [13, 14]. However, these results are typically applicable to statically emplaced networks. As such, they cannot overcome "degenerate" configurations where node connectivity prohibits closed form solutions.

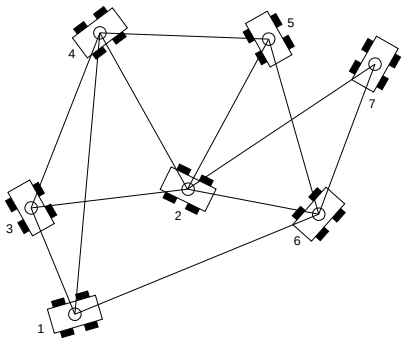

Figure 1: A set of robot or sensor nodes located in a 2D workspace. The edges between the nodes indicate available sensor measurements, that is an edge between nodes i and j indicates that robot i can measure the range and or bearing to node j from its current position.

3 The Localization Approach

The basic elements of the localization framework are diagrammed in Figure 1. Here a set of robots or sensor nodes equipped with range sensors operate on the plane. The edges in the figure indicate available relative measurements, that is an edge between nodes i and j indicates that robot i can measure the range and or bearing to node j from its current position. Note that the edges in this graph will, in general, be directed. Let $\mathcal{C} \subset \mathbb{R}^{2n}$ denote the configuration space of our robot team. Let $\mathbf{x} = \begin{pmatrix} x_1 & y_1 & x_2 & y_2 & \cdots & x_n & y_n \end{pmatrix} \in \mathcal{C}$ denote the current configuration of the ensemble. Note that $\mathbf{x}$ is simply the concatenation of the coordinates of the n robots. Note also that we are purposely vague about the frame of reference to which these coordinates are to be measured. Depending on what parameters we are trying to estimate, various choices will be more or less convenient. Hence we defer this decision until the parameter estimation phase. In this framework, we assume that we can bound the error in the range measurements. Each such measurement is then viewed as a constraint on the possible values of the configuration vector $\mathbf{x}$.

In previous research [3] we noted that if the robots were able to measure their orientation with respect to a common frame of reference, say through the use of a compass, the bearing and range measurements induce a set of linear constraints on the robots configuration $\mathbf{x}$. By concatenating these constraints, we obtain a system of linear inequalities of the form $\boldsymbol{A}\vec{x} \leq b$, where $A \in \mathbb{R}^{l \times (m \cdot n)}$ represents the sensor constraints. Implicitly, the matrix $\boldsymbol{A}$ induces a polytope $\mathbf{P} \subset \mathcal{C}$ on the configuration space of the formation which corresponds to the set of all configurations consistent with the sensor measurements.

Once a set of constraints has been derived from the available measurements, we can proceed to consider the prob-

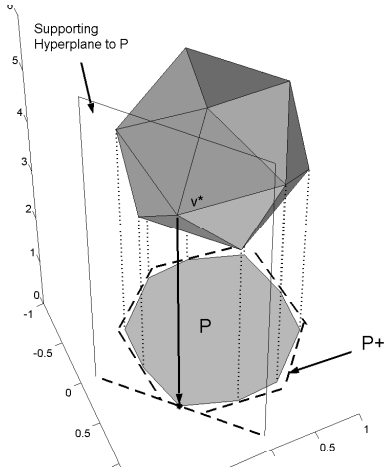

Figure 2: A polytope **P** in $\mathbb{R}^3$ and its $\mathbb{R}^2$ projection $\mathcal{P}$ (shaded polygon). $\mathcal{P}$ can be approximated by solving multiple linear programming problems - each yielding an extreme point x^* on the polytope and its supporting hyperplane. The projections of these correspond to half-spaces in $\mathbb{R}^2$. Intersecting these generates a bounding approximation $\mathcal{P}^+$ to the true projection.

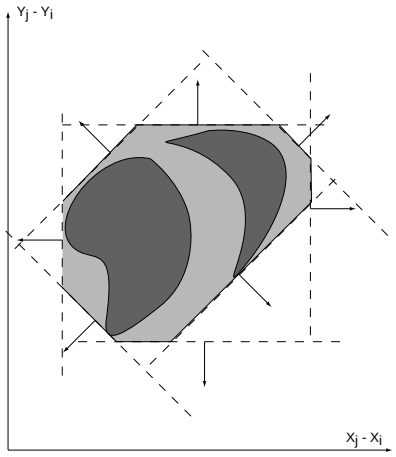

Figure 3: A convex approximation for the feasible set of displacements (the dark region) can be obtained by bounding the its extent along various search directions. The intersection of the resulting half planes defines a convex polytope (the light gray area) which, by construction, encloses the feasible set.

lem of gauging the positions of the robots with respect to each other. Consider for example the problem of estimating the relative position of node j with respect to node i represented by the displacement vector v_{ij}. This vector can be expressed as a linear function of the configuration vector $\mathbf{x}$ as follows: $v_{ij} = (P_j - P_i)\mathbf{x}$, where $P_m \in \mathbb{R}^{2\times 2n}$ denotes the sparse projection matrix that extracts the coordinates of node m from the configuration vector $\mathbf{x}$ in $\mathbb{R}^{2n}$. Hence, we can view the problem of bounding this vector as one of gauging the projection of the set of feasible configurations onto the 2D subspace corresponding to the displacement v_{ij}.

We can bound the uncertainty in this displacement vector by choosing search directions parallel to the subspace of interest and bounding the extent of the feasible region in those directions. This amounts to finding bounds for objective functions of the form shown in Equation 1 given the constraints derived from the measurements.

$$\begin{pmatrix}\cos\theta & \sin\theta\end{pmatrix}^T v_{ij} = \begin{pmatrix}\cos\theta & \sin\theta\end{pmatrix}^T (P_j - P_i)\mathbf{x} \quad (1)$$

Figure 2 depicts the situation when the constraints are all linear and, therefore, convex. Figure 3 depicts the more general situation where the feasible region and its projection need not be convex or even connected. In either case, the goal of the localization scheme is to produce a convex polyhedral approximation which bounds the projection of all feasible configurations and, hence, bounds the displacement vector.

In the case where all the constraints are convex, we have described a procedure for selecting the search directions in a manner which guarantees that the disparity between our convex approximation and the actual projection of the feasible region diminishes at the rate of $\frac{1}{s^2}$, where s corresponds to the number of search directions used to bound the approximation. Interestingly, this convergence bound is *independent of the number of constraints* (see [3] for details). In the general case where the set of feasible configurations is *not* convex, a convex approximation of its projection is readily constructed by simply choosing a set of uniformly spaced search directions.

3.1 Modeling Range Constraints

Figure 4 depicts the annulus of feasible configurations induced by range constraints on the range between nodes i and j. The constraints induced by a bounded error range measurement can be expressed quite simply [12] in terms of the configuration vector $\mathbf{x}$. Here we again note that the vector displacement between nodes i and j, v_{ij} can be expressed as a linear function of the global configuration vector $\mathbf{x}$. *i.e.* $v_{ij} = (P_j - P_i)\mathbf{x}$. If r_{ij}^+ and r_{ij}^- respectively denote the upper and lower bounds on the range measurement, we can derive two quadratic constraints as follows:

$$v_{ij}^T v_{ij} = \mathbf{x}^T (P_j - P_i)^T (P_j - P_i)\mathbf{x} \quad \leq \quad (r_{ij}^+)^2 \quad (2)$$
$$v_{ij}^T v_{ij} = \mathbf{x}^T (P_j - P_i)^T (P_j - P_i)\mathbf{x} \quad \geq \quad (r_{ij}^-)^2 \quad (3)$$

Note that the upper bound constraint shown in Equation 2 is a convex function of $\mathbf{x}$ while the lower bound constraint in Equation 3 is concave. The former can be integrated directly into our localization framework [3] by replacing the linear programming problem with a more generalized semidefinite program. Exploiting the latter constraint is less trivial. However, in [15] we observed that

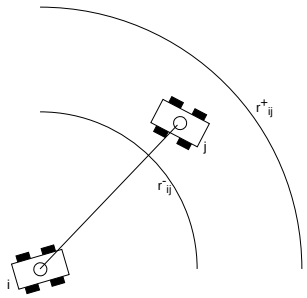

Figure 4: Range measurements with bounded error constrain the distance between two nodes i and j.

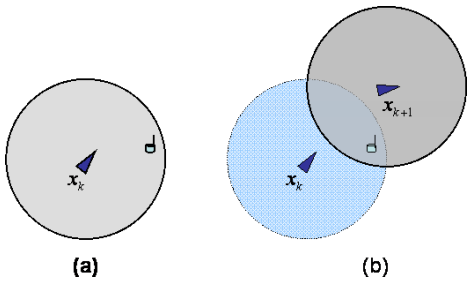

Figure 5: Localization of a sensor node across two time steps. Bounds on robot motion error allows the uncertainty in node position to be refined over time.

convex relaxations exists that could provide useful bounds on the feasible solutions. This relaxation essentially corresponds to tackling the Lagrangian dual of the original optimization problem. These convex relaxations can also be formulated as semidefinite programs which means that we can still take advantage of the efficient solution schemes developed for this class of problems.

3.2 Extending to Dynamic Operations

The motivation for range-only SLAM is the unobservability of sensor positions from a limited number of measurements. In an attempt to provide intuition into the approach, consider the example in Figure 5 in which a robot (blue triangle) attempts to localize the position of a single sensor node. With one observation (using the convex constraint), the node position is only constrained to lie within a circular uncertainty region (a). When the distance between the robot and sensor node is large, such an estimate by itself provides little information for approximating the node's position. However, let us assume that the robot moves (for now) without motion error and receives a new range measurement. A new constraint set which implicitly bounds the node's position can then be formed by merging the posterior position estimate with this new sensor measurement (b). Recovering the projection estimate at time $k+1$ then provides a refined bounds on the node's position. The process then recursively iterates, eventually yielding useful bounds for the position of the sensor node.

Now consider the more realistic case where robot position estimates are also subject to dead reckoning errors. If these errors can be bounded across a single time step, the corresponding uncertainty can also be modeled by a finite set of convex constraints that relate the robot's position x_k with its future position at time x_{k+1}. These motion constraints can then be merged with the posterior and new range measurements to refine the projection estimates for both the robot and sensor nodes over time. What is compelling about such an approach is that the new position estimates are *simultaneously conditioned* on the uncertainty in prior estimates, robot motion, and current sensor estimates. For a more formal treatment of the filtering algorithm, the reader is referred to [15].

4 Experimental Results

To characterize the performance of our SLAM approach, several flavors of simulations were conducted: using purely convex constraints, using convex constraints in the presence of measurement outliers, and using both convex and nonconvex constraints. In each, a single robot was charged with estimating its own position while simultaneously localizing a set of WSN nodes randomly distributed on the plane. To accomplish this, the robot explored the WSN by making a figure-eight trajectory once through the environment. Dead reckoning error bounds for each time step were based upon a velocity error $dv = \pm 0.05v$ and heading error $d\theta = \pm 5^\circ$. These were modeled by a set of four linear constraints. Filter measurement updates occurred at a rate of 1 *Hz*.

4.1 Range-only SLAM Using Convex Constraints

In the initial experiments, our objective was to demonstrate the viability of range-only SLAM using solely convex constraints. Range errors were generated from a uniform distribution with a magnitude of ± 0.2 times the true range, and projection estimates (corresponding to uncertainty regions in node positions) were approximated by considering 12 search directions on the plane.

A representative trial is shown at Figure 6, and reflects the evolution of the uncertainty regions over time. The initial robot/WSN pose (a) also reflects the available measurement links between sensor nodes. The low degree of connectivity prevents a closed form solution for estimating the node positions. However, over time the uncertainty regions develop as the robot explores the environment (b-c). After a single exploratory loop (d), the uncertainty regions are well defined and are capable of providing reasonable point estimates for node positions.

While these results are strongly tied to the simulation

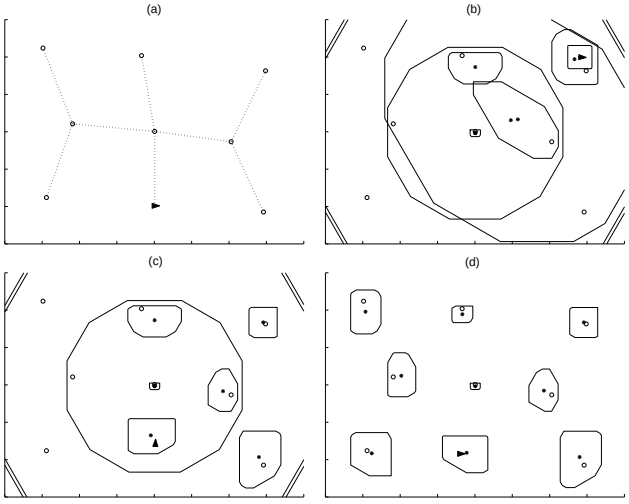

Figure 6: A simulation trial showing the initial network configuration (a), and the evolution of SLAM estimates over time (b-d). Polyhedra bound the uncertainty regions guaranteed to contain the robot/node positions. Actual node positions are denoted by circles, and the polyhedra centroids by stars.

parameters, they clearly demonstrate the ability of our approach to provide useful position estimates as well as quantifying the degree of uncertainty associated with each estimate.

4.2 Convergence in the Presence of Measurement Outliers

To demonstrate convergence in the presence of outliers, simulation trials were also conducted where range measurements suffered from the effects of signal attenuation and multi-path errors. Each measurement was chosen as an outlier with a probability of 0.25. Those so designated were *overestimated* by 50-100% (as taken from a uniform distribution). Unmodeled errors of this magnitude have significant impacts on traditional filter estimates, and potentially lead to filter divergence. In our approach, such outliers are accommodated gracefully.

A total of 10 simulations were conducted where a single robot was charged with estimating its own position while simultaneously localizing a set of 6-12 WSN nodes randomly placed on the plane. In addition to outliers, range measurements were corrupted with random errors of ± 0.2 times the actual range. Projection estimates for the robot/nodes were constructed by considering 8-12 search directions in the plane.

As expected, no instances of filter divergence were observed during these trials. Since each measurement update in the SLAM filter is effectively a set intersection, such outliers cannot adversely affect the current position estimates. Instead, they can cause the filter to converge more slowly. In an attempt to quantify this, each simulation was run with and without measurement outliers. A representative trial from these is shown at Figure 7. The net effect of the outliers was to increase the mean area of the uncertainty regions by 10% over the course of simulations.

It must be mentioned that while our approach offers protection against outliers which overestimate range measurements, unmodeled underestimates could result in filter divergence. In the presence of such outliers, additional tests for robustness such as chi-squared techniques for individual measurement rejection, or a RANSAC based approach for model rejection [16] would need to be incorporated.

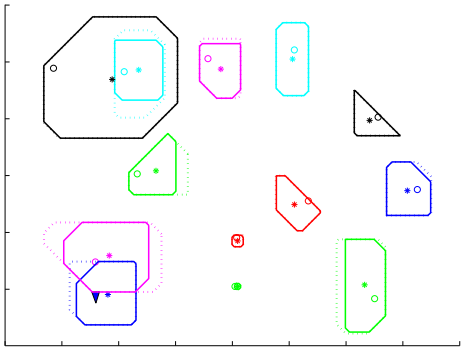

Figure 7: A simulation trial showing SLAM performance with (dotted lines) and without (solid lines) measurement outliers. Outliers failed to cause filter divergence, but resulted on average in an 11% increase in the size of the uncertainty regions.

4.3 Integrating Non-convex Constraints

In the final simulation phase, we attempted to empirically quantify the effects of incorporating nonconvex constraints on SLAM performance. For purposes of comparison, each trial was run using purely convex as well as convex and non-convex range constraints combined. Furthermore, three different bounds on range errors were examined: $dr = 0.01, 0.10$, and 0.20 times the actual range.

A total of 30 trials were conducted, with additional parameters consistent with those used in previous simulations. To compare the results quantitatively, we computed the areas of the projection approximations returned by the procedure for each set of simulation parameters. These areas provide an indication of how effective the constraints are at narrowing down the set of feasible configurations. Larger areas correspond to greater uncertainty. A representative trial for dr=0.01 is shown at Figure 8. A summary analysis is provided at Table 1.

While preliminary, these results imply that significant performance improvements from approximating the nonconvex constraints will only be seen for small range error bounds. This contrasts with our previous findings in static localization performance, where improvements of

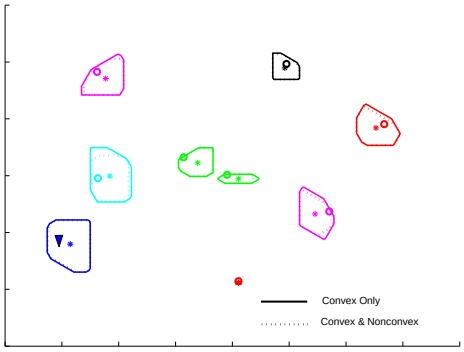

Figure 8: Simulation results contrasting SLAM performance using purely convex (solid lines) and both convex and nonconvex constraints (dotted lines). In this trial, the mean projection area was reduced by 12% in the latter case

30% were the norm [15]. However, in these trials two or more node positions were assumed known *a priori*. Such an assumption is invalid for the SLAM problem. In hind-

Error Bound	**Area(Convex/Nonconvex)**
dr=0.01	1.121
dr=0.10	1.020
dr=0.20	1.002

Table 1: The impact integrating nonconvex constraints. For these simulations, a significant improvement in performance was only seen in the case of small range errors.

sight, these results are not completely unexpected. As minimum range constraints are completely concave, convex approximations of these can provide no initial improvements in localization performance. It is only when the uncertainty regions for the node positions begin to form that the intersection of the corresponding linear constraint sets with the concave constraints can benefit from a convex approximation. Unfortunately, by this time uncertainty in the robot's position has accumulated and tempers their utility.

Again, we emphasize that these findings are preliminary and extremely dependent upon the simulation parameters chosen. We will continue to investigate this topic more thoroughly in the near future. At this point however, it appears that the additional computational resources expended in integrating non-convex constraints would be more wisely spent by refining the convex projection estimates with additional searches.

5 Discussion and Conclusions

In this paper, we extended our results in cooperative multi-robot localization to the range-only SLAM problem for use in robot augmented wireless sensor networks. The approach possesses several desirable attributes, to include convergence and performance guarantees, as well as resistance to signal attenuation and multi-path errors. Based on empirical data provided in [12], we expect that a single robot could employ it to localize a 100 node WSN at several hertz using current PC technology.

While all simulations were limited to a single robot, the approach is distributed and extends readily to any number of mobile agents. Additionally, through the use of a bounded error model interdependencies in the uncertainty estimates are handled in a principled manner.

Significant future work is planned on this topic. To this point all of our results have been limited to Matlab simulations. We are currently working with a small set of RF based sensor motes, and will augment these with experimental work in the coming months.

Performance improvements from integrating nonconvex constraints were not as significant as hoped. We plan a more formal analysis in an attempt to better leverage these additional constraints. It is also apparent that localization performance is a strong function of the robot's trajectory. As such, we also plan to investigate control strategies for improving the quality of the position estimates.

Lastly, we are currently working on extending this approach to the problem of localizing nodes in $\mathbb{R}^3$. We envision that such a scheme would be useful for localizing sensor motes and teams composed of both ground and air vehicles.

References

[1] D. Kurth, G. Kantor, and S. Singh, "Experimental results in range-only localization with radio," in *IEEE/RSJ International Conf. on Intelligent Robots and Systems*, Las Vegas, NV, Oct 2003.

[2] P. Newman and J. Leonard, "Pure range-only subsea slam," in *IEEE International Conference on Robotics and Automation*, Taiwan, China, Sep 2003.

[3] J. Spletzer and C.J. Taylor, "A bounded uncertainty approach to multi-robot localization," in *Proceedings of the Conference on Intelligent Robots and Systems (IROS 2003)*, 2003.

[4] J. Nemeroff, L. Garcia, D. Hampel, and S. DiPierro, "Application of sensor network communications," in *IEEE MILCOM*, 2001.

[5] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey on sensor networks," in *IEEE Communications Magazine*, August 2002.

[6] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly resilient, energy efficient multipath routing in wireless sensor networks," in *Mobile Computing and Communications Review (MC2R)*, 2002.

[7] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, "Robust monte carlo localization for mobile robots," *Artificial Intelligence*, vol. 128, no. 1-2, pp. 99–141, 2000.

[8] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probablistic approach to collaborative multi-robot localization," *Autonomous Robots: Special Issue on Heterogeneous Multi-Robot Systems*, vol. 8, no. 3, pp. 325–344, 2000.

[9] G. Dissanayake, P. Newman, H. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building," *IEEE Transactions on Robotics and Automomation*, vol. 17, no. 3, pp. 229–241, 2001.

[10] D. Bertsekas and L. Rhodes, "Recursive state estimation for a set membership representation of uncertainty," in *IEEE Transactions on Automatic Control*, July 1971, vol. 16(2).

[11] F. Schweppe, *Uncertain Dynamic Systems*, Prentice Hall, 1973.

[12] L. Doherty, K. Pister, and L. El Ghaoui, "Convex position estimation in wireless sensor networks," in *Proceedings of IEEE Infocom*, April 2001.

[13] A. Savvides, C. Han, and M. Strivastava, "Dynamic fine-grained localization in ad-hoc networks of sensors," in *Mobile Computing and Networking*, 2001.

[14] C. Savarese, J. Rabay, and K. Langendoen, "Robust positioning algorithms for distributed ad-hoc wireless sensor networks," in *USENIX Technical Annual Conference*, June 2002.

[15] J. Spletzer and C.J. Taylor, "Convex approximation strategies for multi-robot localization," in *Submitted to WAFR*, 2004.

[16] M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," in *Communications of the ACM*, 1981.