

Multi-Robot Cooperative Localization: A Study of Trade-offs Between Efficiency and Accuracy

Ioannis M. Rekleitis¹, Gregory Dudek¹, Evangelos E. Miliotis²

¹Centre for Intelligent Machines, McGill University, Montreal, Québec, Canada

²Faculty of Computer Science, Dalhousie University, Halifax, Nova Scotia, Canada
contact: {yiannis,dudek}@cim.mcgill.ca, eem@cs.dal.ca

Abstract

This paper examines the tradeoffs between different classes of sensing strategy and motion control strategy in the context of terrain mapping with multiple robots. We consider a larger group of robots that can mutually estimate one another's position (in 2D or 3D) and uncertainty using a sample-based (particle filter) model of uncertainty. Our prior work has dealt with a pair of robots that estimate one another's position using visual tracking and coordinated motion. Here we extend these results and consider a richer set of sensing and motion options. In particular, we focus on issues related to confidence estimation for groups of more than two robots ¹.

1 Introduction

In this paper we discuss the benefits different sensing modalities for *cooperative localization* by a team of mobile robots. The term *cooperative localization* describes the technique whereby the members of a team of robots estimate one another's positions [13]. This type of multi-robot exploration strategy is able to compensate for deficiencies in odometry and/or a pose sensor by combining measurements. Herewith we look at how the expressive power of the sensor relates to the quality of the final pose estimates produced by collaborative exploration. A key aspect of collaborative exploration is the use of a sensor (*robot tracker*) to estimate the pose of a moving robot relative to one or more stationary ones (see section 1.1). Furthermore, we consider the effects of different robot tracker sensors on the accuracy of localization for a moving robot *using only* the information from the rest of the robots (as opposed to observations of the environment). This approach results in an open loop estimate (with respect to the entire team) of the

moving robot's pose without dependence on information from the environment. The experimental results allow us to examine the effectiveness of cooperative localization and estimate upper bounds on the error accumulation for different sensing modalities. To the extent that limited space permits, we also discuss the advantage of using randomized formation control to move the robots.




Figure 1: Two robots, one equipped with laser range finder (right) and the other with a target (left), employing cooperative localization.

1.1 Cooperative Localization

Several different sensors have been employed for the estimation of the pose of one robot with respect to another robot. We restrict our attention to “robot tracker” sensors which return information in the frame of reference of the observing robot (i.e they estimate pose parameters of one robot relative to another robot making the observation). Consequently, for “two-dimensional robots” in a two dimensional environment, or for robots whose pose can be approximated as a combination of 2D position and an orientation, we can express the pose using three measurements; for ease of reference we represent these measurements by the triplet $T = [\rho \ \phi \ \theta]$, where ρ is the distance between the two robots, ϕ is the an-

¹To appear in 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems, EPFL, Switzerland, September 30 - October 4, 2002.

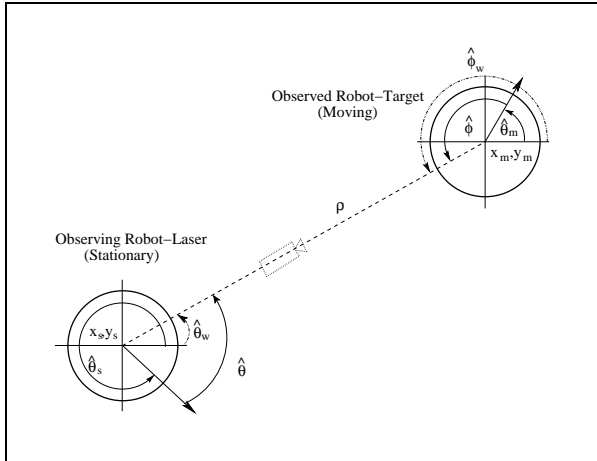


Figure 2: Pose Estimation via Robot Tracker: Observation of the Moving Robot by the Stationary Robot. Note that the “camera” indicates the robot with the Robot Tracker; and $\hat{\theta}_w, \hat{\phi}_w$ are angles in world coordinates.

gle at which the observing robot sees the observed robot relative to the heading of the observing robot, and θ is the heading of the observed robot as measured by the observing robot relative to the heading of the observing robot (Figure 1b). If the stationary robot is equipped with the Robot Tracker, where $\mathbf{X}_m = [x_m, y_m, \theta_m]^T$ is the pose of the moving robot and $\mathbf{X}_s = [x_s, y_s, \theta_s]^T$ is the pose of the stationary robot then Equation 1 returns the sensor output T :

$$\begin{bmatrix} \rho \\ \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \sqrt{dx^2 + dy^2} \\ \text{atan2}(dy, dx) - \theta_s \\ \text{atan2}(-dy, -dx) - \theta_m \end{bmatrix} \quad (1)$$

where $dx = x_m - x_s$ and $dy = y_m - y_s$.

In pose estimation problems such as uncertainty management can be challenging. In order to estimate the probability distribution function (*pdf*) of the pose of the moving robot i at time t ($P(\mathbf{X}_i^t)$) we employ a particle filter (Monte Carlo simulation approach: see [7, 3, 11]). The weights of the particles (W_i^t) at time t are updated using a Gaussian distribution (see Equation 2 where $[\rho_i, \theta_i, \phi_i]^T$ has been calculated as in Equation 1 but using the pose of a single particle “ i ” (\mathbf{X}_{m_i}) instead of the moving robot pose (\mathbf{X}_m)).

$$W_i^t = W_i^{t-1} \frac{1}{\sqrt{2\pi}\sigma_\rho} e^{-\frac{(\rho-\rho_i)^2}{2\sigma_\rho^2}} \frac{1}{\sqrt{2\pi}\sigma_\theta} e^{-\frac{(\theta-\theta_i)^2}{2\sigma_\theta^2}} \frac{1}{\sqrt{2\pi}\sigma_\phi} e^{-\frac{(\phi-\phi_i)^2}{2\sigma_\phi^2}} \quad (2)$$

The rest of the paper is structured as follows. The next Section 2 presents some background work. Section 3 contains an analysis and experimental study of

the primary different classes of sensory information that can be naturally used in cooperative localization. Finally, Section 5 presents our conclusions and a brief discussion of future work.

2 Previous Work

Prior work on multiple robots has considered collaborative strategies when the lack of landmarks made localization impossible otherwise ([4]). A number of authors have considered pragmatic multi-robot map-making. Several existing approaches operate in the sonar domain, where it is relatively straightforward to transform observations from a given position to the frame of reference of the other observers thereby exploiting structural relationships in the data ([10, 5, 1]). One approach to the fusion of such data is through the use of Kalman Filtering and its extensions ([15, 14]).

In other work, Rekleitis, Dudek and Milios have demonstrated the utility of introducing a second robot to aid in the tracking of the exploratory robot’s position ([12]) and introduced the concept of cooperative localization.

Recently, several authors have considered using a team of mobile robots in order to localize using each other. A variety of alternative sensors has been considered. For example, [8] use robots equipped with omnidirectional vision cameras in order to identify and localize each other. In contrast, [2] use a pair of robots, one equipped with an active stereo vision and one with active lighting to localize. The various methods employed for localization use different sensors with different levels of accuracy; some are able to estimate accurately the distance between the robots, others the orientation (azimuth) of the observed robot relative to the observing robot and some are able to estimate even the orientation of the observed robot.

3 Sensing Modalities

As noted above, several simple sensing configurations for a robot tracker are available. For example, simple schemes using a camera allow one robot to observe the other and provide different kinds of positional constraints such as the distance between two robots and the relative orientations.

In this section we consider the effect the group size has on the accuracy of the localization for different classes of sensors. The experimental arrangement of the robots is simulated and is consistent across all the sensing configurations. The robots start in a single line and they move abreast one at a time, first in ascending order and then in descending order for a set number of exchanges. The selected robot moves for 5 steps and after each step cooperative

localization is employed and the pose of the moving robot is estimated. Each step is a forward translation by 100cm. Figure 3 presents a group of three robots, after the first robot has finished the five steps and the second robot performs the fifth step.

3.1 Range Only

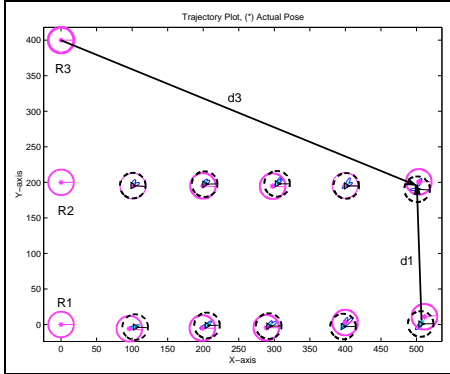


Figure 3: Estimation of the pose of robot R2 using only the distance from robot R1 ($d1$) and from robot R3 ($d3$).

One simple sensing method is to return the relative distance between the robots. Such a method has been employed by [6] in the millibots project where an ultra-sound wave was used in order to recover the relative distance. In order to recover the position of one moving robot in the frame of reference of another, at least two stationary robots (that are not collinear with the moving one) are needed thus the minimum size of the group using this scheme is three robots.

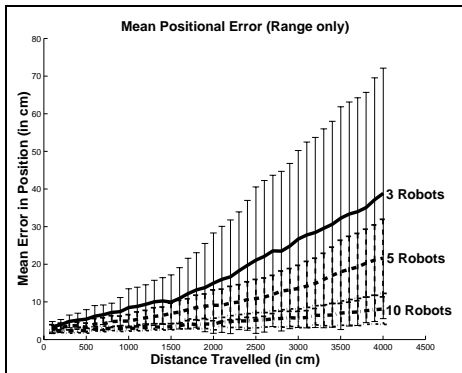


Figure 4: Average error in position estimation using the distance between the robots only (3, 4 and 10 robots; bars indicate standard deviation).

The distance between two robots can be easily and robustly estimated. In experimental simulations, the distance between every pair of robots was estimated and Gaussian, zero mean, noise was added with $\sigma_\rho = 2\text{cm}$ regardless the distance between the two robots. Figure 4 presents the mean error per unit distance traveled for all robots, averaged over

20 trials. As can be seen in Figure 4 with five robots, the positional accuracy is acceptable with an error of 20cm after 40m traveled; for ten robots the accuracy of the localization is very good.

3.2 Azimuth (Angle) Only

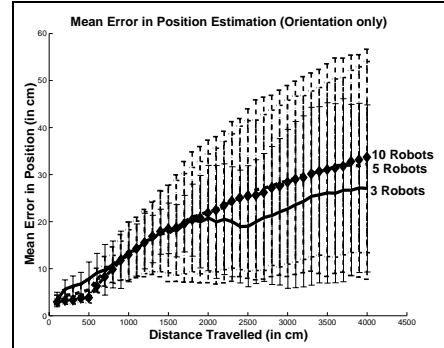


Figure 5: Average error in position estimation using the orientation of the moving robot is seen by the stationary ones.

Several robotic systems employ an omnidirectional vision sensor that reports the angle at which another robot is seen. This is also consistent with information available from several types of observing systems based on pan-tilt units. In such cases the orientation at which the moving robot is seen can be recovered with high accuracy. We performed a series of trials using only the angle at which one robot is observed, with groups of robots of different sizes. As can be seen in Figure 5 the accuracy of the localization does not improve as the group size increases. This is not surprising because small errors in the estimated orientation of the stationary robots scale non-linearly with the distance. Thus after a few exchanges the error in the pose estimation is dominated by the error in the orientation of the stationary robots.

To illustrate the implementation of the particle filter, we present here the probability distribution function (pdf) of the pose of the moving robot after one step (see Figure 6). The robot group size is three and it is the middle robot R2 that moves. The predicted pdf after a forward step can be seen in the first sub-figure (6a) using odometry information only; the next two sub-figures (6b,6c) present the pdf updated using the orientation at which the moving robot is seen by a stationary one (first by robot R1 then by robot R3); finally, the sub-figure 6d presents the final pdf which combines the information from odometry and the observations from the two stationary robots. Clearly the uncertainty of the robot's position is reduced with additional observations.

3.3 Position Only

Another common approach is to use the position of one robot computed in the frame of reference of another (relative position). This scheme has been

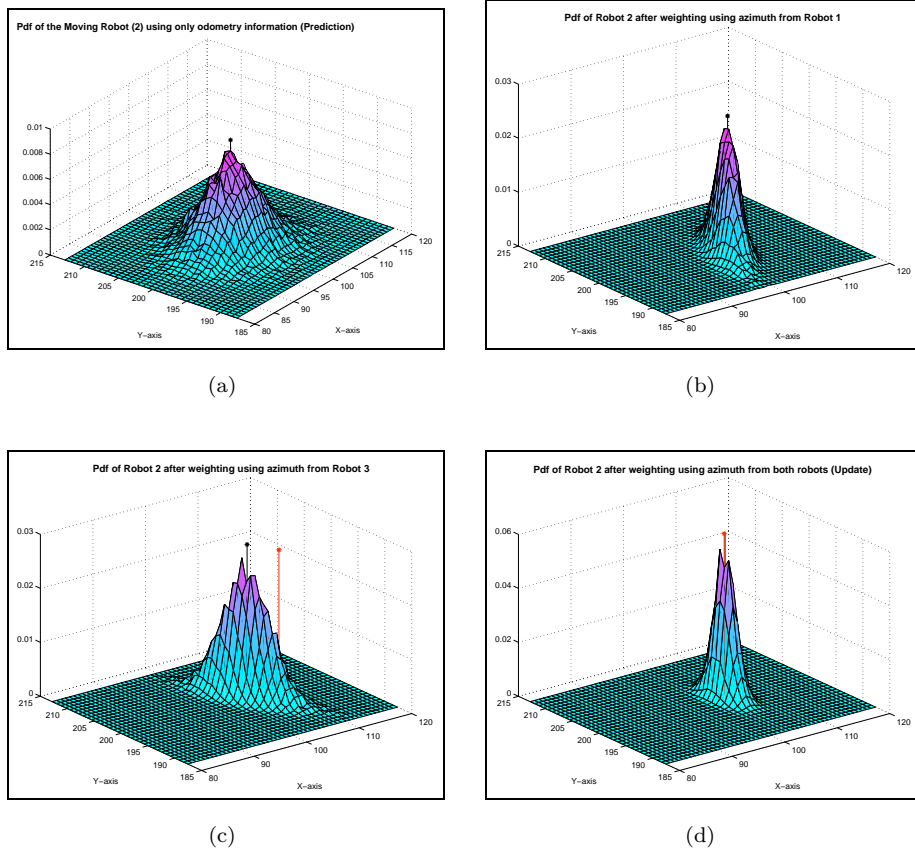


Figure 6: The pdf of the moving robot (R2) at different phases of its estimation: (a) prediction using odometry only; (b) using the orientation from stationary robot R1; (c) using the orientation from stationary robot R3; (d) final pdf.

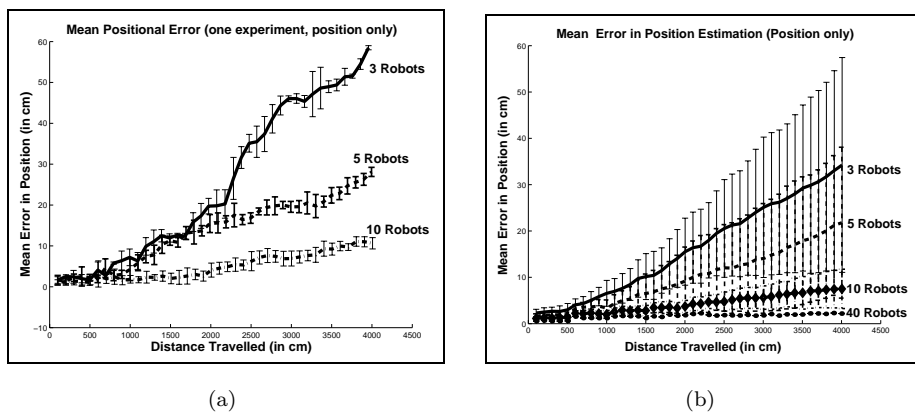


Figure 7: Average error in position estimation using both the distance between the robots and the orientation the moving robot is seen by the stationary ones. (a) Average error in positioning of the team of robots one trial (3,5 and 10 robots). (b) Average error in position estimation over twenty trials (3,5,10 and 40 robots).

employed with two robots (see [1]) in order to reduce the uncertainty. The range and azimuth information $([\rho, \theta])$ is combined in order to improve the pose estimation. As can be seen in Figure 7a even with three robots the error in pose estimation is relatively small (average error 30cm for 40m distance traveled per robot, or 0.75%). In our experiments the distance between the two robots was estimated and, as above, zero-mean Gaussian noise was added both to distance and to orientation with $\sigma_\rho = 2\text{cm}$ and $\sigma_\theta = 0.5^\circ$ respectively. The experiment was repeated twenty times and the average error in position is shown in Figure 7b for groups of robots of size 3,5,10 and 40.

3.4 Full Pose

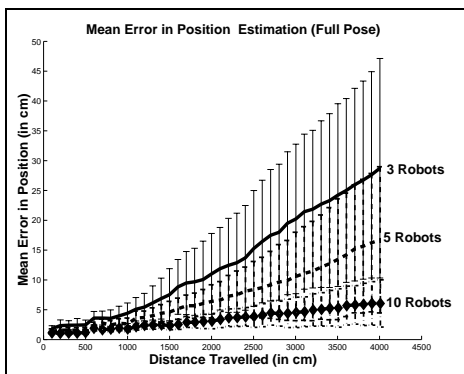


Figure 8: Average error in position estimation using full pose $[\rho, \theta, \phi]$.

Some robot tracker sensors provide accurate information for all three parameters $[\rho, \theta, \phi]$ and they can be used to accurately estimate the full pose of the moving robots (see [9, 13]). In the experimental setup the robot tracker sensor was characterized by Gaussian, zero mean, noise with $\sigma = [2\text{cm}, 0.5^\circ, 1^\circ]$. By using the full Equation 2 we weighted the *pdf* of the pose of the moving robot and performed a series of experiments for 3, 5 and 10 robots. As can be seen in Figure 8 the positional accuracy is consistently lower than in the case of range only, orientation only and position only measurements.

In addition, experiments were conducted for larger group sizes and for longer distances traveled. Figure 9 presents the mean error over thirty experiments for 3,5,10,15,20 and 30 robots. The mean positional error was calculated as a function of the group size in order to examine the contribution of each additional robot to localization. Two different functions were used in order to model the error with respect to the group size (N) (a) $\mathcal{E}_a(N) = \alpha N^\beta + \gamma$ and (b) $\mathcal{E}_b(N) = \alpha e^{\beta N} + \gamma$. Using cross-validation² $\mathcal{E}_a(N)$

²The two functions were fitted for robot group sizes of 3-10,15,20 and 30 (11 group sizes in total), each time omitting one group size and then calculating the difference between the observed error value and the function response.

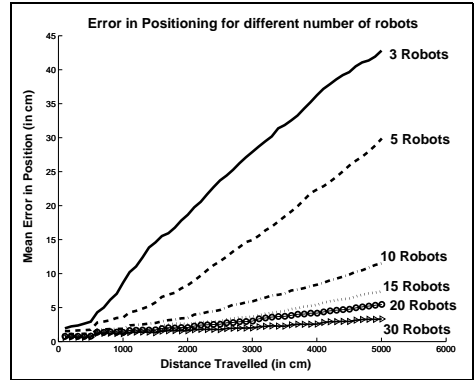


Figure 9: Average error in position estimation using full pose $[\rho, \theta, \phi]$ for different number of robots.

was selected because it had smaller mean squared error. For a fixed distance traveled (50m) the error function is given in Equation 3. As expected the incremental benefit of each additional robot is a function decreasing asymptotically to zero.

$$\mathcal{E}_a(N) = 126.866N^{-0.948} \quad (3)$$

4 Trajectory variation

In this section we outline results regarding the effects of formation control on the accuracy of collaborative exploration – that is, the way the motion pattern of the robots relates to pose errors. In prior work we have considered the geometric optimization of the trajectory of a pair of robots to minimize the effort in covering space, and then estimated the net pose error that accrues.

An alternative viewpoint is to consider the optimization of the robot formation (that is the combination of robot positions) to minimize the accrued pose error. This can be achieved by describing the motion control problem as a variation problem. Unfortunately, an analytical treatment of this problem is both outside the scope of this paper and of limited utility. Instead, we present here a dichotomy between two different classes of formation: the fixed deterministic robot formation described earlier, and a randomized variant of the fixed formation where each robot moves forward according to a stochastic schedule and each robot steps forward by a random step ($step_{rand}$) following a Gaussian distribution with mean equal to the individual steps of the deterministic algorithm ($step_{det}$) and standard deviation equal to 10% to the distance traveled: $step_{rand} = N(step_{det}, 0.1step_{det})$.

In 14 simulated trials with 6 robots we have observed mean errors in pose were substantially reduced with randomized formations where the variance of the individual steps was 1/3 the average step size. These results are illustrated in Figure 10. We believe that this improvement in performance results from the

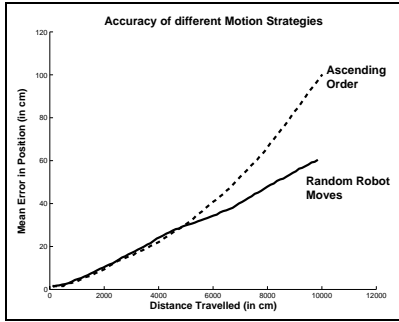


Figure 10: Average error in position estimation using full pose $[\rho, \theta, \phi]$ over 16 trials. Two different motion strategies of 6 robots. Dashed line: robots move in ascending order. Solid line: robots move in random order.

more varies arrangements of the robots when pose estimates are taken. Pose estimation is subject to several geometric degeneracies that can lead to error and by using a randomized motion strategy is appear that these degeneracies are efficiently avoided.

5 Conclusions

In this work we examined the effect of the size of the team of robots and the sensing paradigm on cooperative localization (see Table 1 for a synopsis). Also, preliminary results from experiments with varying odometry error have shown that cooperative localization is robust even with 10-20% odometry errors. The cost-benefit tradeoff seems to be maximized for small teams of robots. While these results are not definitive, being based on several domain-specific assumptions, they seem to illustrate a general relationship.

In addition, it appears that a randomized motion strategy can outperform a deterministic one. For small teams of robot it seems likely that there are even better purely deterministic strategies, although computing these may become complicated as the team-size grows. While this bears further examination it seems likely that for teams of more than two or three robots randomized formation control may provide an appealing alternative to deterministic methods. In future work we hope to further

Number of Robots	3	5	10
Range (ρ)	38.80	21.63	8.13
Azimuth (θ)	27.06	32.20	33.72
Position (ρ, θ)	34.25	21.79	7.50
Full Pose (ρ, θ, ϕ)	28.73	16.71	6.05

Table 1: The mean error in position estimation after 40m travel over 20 trials.

extend the uncertainty study for different group configurations and motion strategies. An interesting ex-

tension would be for the robots to autonomously develop a collaborative strategy to improve the accuracy of localization. Given a large group of robots, an estimate of the effects of team size on error accumulation would allow the group of be effectively partitioned to accomplish sub-tasks while retaining a desired level of accuracy in positioning.

References

- [1] Wolfram Burgard, Dieter Fox, Mark Moors, Reid Simmons, and Sebastian Thrun. Collaborative multi-robot exploration. In *Proc. of the IEEE Int. Conf. on Robotics & Automation*, pages 476–481, 2000.
- [2] A.J. Davison and N. Kita. Active visual localisation for cooperating inspection robots. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, v. 3, pg 1709–1715, Takamatsu, Japan, 2000.
- [3] F. Dellaert, W. Burgard, D. Fox, and S. Thrun. Using the condensation algorithm for robust, vision-based mobile robot localization. In *IEEE Comp. Soc. Conf. on Computer Vision & Pattern Recognition*. 1999.
- [4] Gregory Dudek, Michael Jenkin, Evangelos Miliotis, and David Wilkes. A taxonomy for multiagent robotics. *Autonomous Robots*, 3:375–397, 1996.
- [5] D. Fox, W. Burgard, and S. Thrun. Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 1998. To appear.
- [6] Robert Grabowski and Pradeep Khosla. Localization techniques for a team of small robots. In *Proc. of the 2001 IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pg 1067–1072, v. 3 2001.
- [7] Patric Jensfelt, Olle Wijk, David J. Austin, and Magnus Andersso. Feature based condensation for mobile robot localization. In *IEEE Int. Conf. on Robotics & Automation (ICRA)*, pg 2531–2537, 2000.
- [8] K. Kato, H. Ishiguro, and M. Barth. Identifying and localizing robots in a multi-robot system environment. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, v. 2, pg 966–971, South Korea, 1999.
- [9] R. Kurazume and S. Hirose. Study on cooperative positioning system - optimum moving strategies for cps-iii. In *Proc. IEEE Int. Conf. on Robotics & Automation*, v. 4, pg 2896–2903, 1998.
- [10] John J. Leonard and Hugh F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Trans. on Robotics & Automation*, 7(3):376–382, 1991.
- [11] Jun S. Liu, Rong Chen, and Tanya Logvinenko. A theoretical framework for sequential importance sampling and resampling. In *Sequential Monte Carlo in Practice*. Springer-Verlag, 2001.
- [12] I. Rekleitis, G. Dudek, and E. Miliotis. Multi-robot collaboration for robust exploration. In *Proc. of Int. Conference in Robotics & Automation*, pg 3164–3169, 2000.
- [13] I. Rekleitis, G. Dudek, and E. Miliotis. Multi-robot collaboration for robust exploration. *Annals of Mathematics and Artificial Intelligence*, 31(1-4):7–40, 2001.

- [14] S. Roumeliotis and G. Bekey. Bayesian estimation and kalman filtering: A unified framework for mobile robot localization. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pg 2985–2992, 2000.
- [15] S. Roumeliotis and G. Bekey. Collective localization: A distributed kalman filter approach to localization of groups of mobile robots. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pg 2958–2965, 2000.