

Synergetic Localization for Groups of Mobile Robots

Stergios I. Roumeliotis and George A. Bekey
Robotics Research Laboratories
University of Southern California
Los Angeles, CA 90089-0781
stergios|bekey@robotics.usc.edu

Abstract

In this paper we present a new approach to the problem of simultaneously localizing a group of mobile robots capable of sensing each other. Each of the robots collects sensor data regarding its own motion and shares this information with the rest of the team during the update cycles. A single estimator, in the form of a Kalman filter, processes the available positioning information from all the members of the team and produces a pose estimate for each of them. The equations for this centralized estimator can be written in a decentralized form therefore allowing this single Kalman filter to be decomposed into a number of smaller communicating filters each of them processing local (regarding the particular host robot) data for most of the time. The resulting decentralized estimation scheme constitutes a unique mean for fusing measurements collected from a variety of sensors with minimal communication and processing requirements. The distributed localization algorithm is applied to a group of 3 robots and the improvement in localization accuracy is presented. Finally, a comparison to the equivalent distributed information filter is provided.

1 Introduction

Precise localization is one of the main requirements for mobile robot autonomy [6]. Indoors and outdoors robots need to know their exact position and orientation (pose) in order to perform their required tasks. There have been numerous approaches to the localization problem utilizing different types of sensors [7] and a variety of techniques (e.g. [5], [4], [15], [20]). The key idea behind most of the current localization schemes is to optimally combine measurements from proprioceptive sensors that monitor the motion of the vehicle with information collected by exteroceptive sensors that provide a representation of the environment and its signals. Many robotic applications require that robots work in collaboration in order to perform a certain task [8], [16]. Most existing localization approaches refer to the case of a single robot. Even when a group of, say M , robots is considered, the group localization problem is

usually resolved by independently solving M pose estimation problems. Each robot estimates its position based on its individual experience (proprioceptive and exteroceptive sensor measurements). Knowledge from the different entities of the team is not combined and each member must rely on its own resources (sensing and processing capabilities). This is a relatively simple approach since it avoids dealing with the complicated problem of fusing information from a large number of independent and interdependent sources. On the other hand, a more coordinated scheme for localization has a number of advantages that can compensate for the added complexity.

First let us consider the case of a homogeneous group of robots. As we mentioned earlier, robotic sensing modalities suffer from uncertainty and noise. When a number of robots equipped with the same sensors detect a particular feature of the environment, such as a door, or measure a characteristic property of the area, such as the local vector of the earth's magnetic field, a number of *independent* measurements originating from the different members of the group is collected. Properly combining all this information will result in a single estimate of increased accuracy and reduced uncertainty. A better estimate of the position and orientation of a landmark can drastically improve the outcome of the localization process and thus this group of robots can benefit from this collaboration schema.

The advantages stemming from the exchange of information among the members of a group are more crucial in the case of heterogeneous robotic colonies. When a team of robots is composed of different platforms carrying different proprioceptive and exteroceptive sensors and thus having different capabilities for self-localization, the quality of the localization estimates will vary significantly across the individual members. For example, a robot equipped with a laser scanner and expensive INS/GPS modules will outperform another member that must rely on wheel encoders and cheap sonars for its localization needs. Communication and flow of information among the members of the group constitutes a form of sensor sharing and can improve the overall positioning accuracy.

2 Previous Approaches

An example of a system that is designed for cooperative localization is presented in [12]. The authors acknowledge that dead-reckoning is not reliable for long traverses due to the error accumulation and introduce the concept of “portable landmarks”. A group of robots is divided into two teams in order to perform cooperative positioning. At each time instant, one team is in motion while the other remains stationary and acts as a landmark. In the next phase the roles of the teams are reversed and this process continues until both teams reach the target. This method can work in unknown environments and the conducted experiments suggest accuracy of 0.4% for the position estimate and 1 degree for the orientation [11]. Improvements over this system and optimum motion strategies are discussed in [10]. A similar realization is presented in [17], [18]. The authors deal with the problem of exploration of an unknown environment using two mobile robots. In order to reduce the odometric error, one robot is equipped with a camera tracking system that allows it to determine its relative position and orientation with respect to a second robot carrying a helix target pattern and acting as a portable landmark. Both previous approaches have the following limitations: (a) Only one robot (or team) is allowed to move at a certain time instant, and (b) The two robots (or teams) must maintain visual contact at all times.

A different implementation of a collaborative multi-robot localization scheme is presented in [9]. The authors have extended the Monte Carlo localization algorithm to the case of two robots when a map of the area is available to both robots. When these robots detect each other, the combination of their belief functions facilitates their global localization task. The main limitation of this approach is that it can be applied only within known indoor environments. In addition, since information interdependencies are being ignored every time the two robots meet, this method can lead to overoptimistic position estimates.

Although practices like those previously mentioned can be supported within the proposed *distributed multi-robot localization* framework (Section 5), the key difference is that it provides a solution to the most general case where all the robots in the group can move simultaneously while continuous visual contact or a map of the area are not required. In order to treat the group localization problem, we begin from the reasonable assumptions that the robots within the group can communicate with each other (at least 1-to-1 communication) and carry two types of sensors: 1. Proprioceptive sensors that record the self motion of each robot and allow for position tracking, 2. Exteroceptive sensors that monitor the environment for (a) (static) features and identities of the surroundings of the robot to be used in the localization process, and (b) other robots

(treated as dynamic features). The goal is to integrate measurements collected by different robots and achieve localization across all the robotic platforms constituting the group.

The key idea for performing *distributed multi-robot localization* is that the group of robots must be viewed as one entity, the “group organism”, with multiple “limbs” (the individual robots in the group) and multiple virtual “joints” visualized as connecting each robot with every other member of the team. The virtual “joints” provide 3 degrees of freedom ($\Delta x, \Delta y, \Delta \phi$) and thus allow the “limbs” to move in every direction within a plane without any limitations. Considering this perspective, the “group organism” has access to a large number of sensors such as encoders, gyroscopes, cameras etc. In addition, it “spreads” itself across a large area and thus it can collect far more rich and diverse exteroceptive information. When one robot detects another member of the team and measures its relative pose, it is equivalent to the “group organism’s” joints measuring the relative displacement of these two “limbs”. When two robots communicate for information exchange, this can be seen as the “group organism” allowing information to travel back and forth from its “limbs”. This information can be fused by a centralized processing unit and provide improved localization results for all the robots in the group. At this point it can be said that a realization of a two-member “group organism” would resemble the multiple degree of freedom robot with compliant linkage shown to improve localization implemented by J. Borenstein [1], [2], [3]. The main drawback of addressing the cooperative localization problem as an information combination problem within a single entity (“group organism”) is that it requires centralized processing and communication. The solution would be to attempt to decentralize the sensor fusion within the group. The *distributed multi-robot localization* approach uses the previous analogy as its starting point and treats the processing and communication needs of the group in a distributed fashion. This is intuitively desired; since the sensing modalities of the group are distributed, so should be the processing modules. As it will be obvious in the following sections, our formulation differs from the aforementioned ones on its starting point. It is based on the unique characteristic of the multi-robot localization problem that the state propagation equations of the centralized system are decoupled while state coupling occurs only when relative pose measurements become available. Our focus is distributed state estimation rather than sequential sensor processing. Nevertheless, the latter can be easily incorporated in the resulting distributed localization schema. In order to deal with the cross-correlation terms (localization interdependencies) that can alter the localization result [21], the data processed during each *distributed multi-robot localization* session must be propagated among all the robots in the group. While this can happen instantly in groups of 2 robots, in the following

sections we will show how this problem can be treated by reformulating the *distributed multi-robot localization* approach so it can be applied in groups of 3 or more robots.

3 Problem Statement

We state the following assumptions:

1. A group of M independent robots move in an N – *dimensional* space. The motion of each robot is described by its own linear or non-linear equations of motion,
2. Each robot carries proprioceptive and exteroceptive sensing devices in order to propagate and update its own position estimate. The measurement equations can differ from robot to robot depending on the sensors used,
3. Each robot carries exteroceptive sensors that allow it to detect and identify other robots moving in its vicinity and measure their respective displacement (relative position and orientation),
4. All the robots are equipped with communication devices that allow exchange of information within the group.

As we mentioned before, our starting point is to consider this group of robots as a single centralized system composed of each and every individual robot moving in the area and capable of sensing and communicating with the rest of the group. In this centralized approach, the motion of the group is described in an $N \times M$ -dimensional space and it can be estimated by applying Kalman filtering techniques. The goal now is to treat the Kalman filter equations of the centralized system so as to distribute the estimation process among M Kalman filters, each of them operating on a different robot. Here we will derive the equations for a group of $M = 3$ robots. The same steps describe the derivation for larger groups. The trajectory of each of the 3 robots is described by the following equations:

$$\vec{x}_i(t_{k+1}^-) = \Phi_i(t_{k+1}, t_k) \vec{x}_i(t_k^+) + B_i(t_k) \vec{u}_i(t_k) + G_i(t_k) \vec{n}_i(t_k) \quad (3.1)$$

for $i = 1..3$, where $\Phi_i(t_{k+1}, t_k)$ is the system propagation matrix describing the motion of vehicle i , $B_i(t_k)$ is the control input matrix, $\vec{u}_i(t_k)$ is the measured control input, $G_i(t_k)$ is the system noise matrix, $\vec{n}_i(t_k)$ is the system noise associated with each robot and $Q_{di}(t_k)$ is the corresponding system noise covariance matrix.

4 Distributed Localization after the 1st Update

In this section we present the propagation and update cycles of the Kalman filter estimator for the centralized

system *after the first update*.¹ Since there have been introduced cross-correlation elements in the covariance matrix of the state estimate, this matrix would now have to be written as:

$$P(t_{k+1}^-) = \begin{bmatrix} P_{11}(t_{k+1}^-) & P_{12}(t_{k+1}^-) & P_{13}(t_{k+1}^-) \\ P_{21}(t_{k+1}^-) & P_{22}(t_{k+1}^-) & P_{23}(t_{k+1}^-) \\ P_{31}(t_{k+1}^-) & P_{32}(t_{k+1}^-) & P_{33}(t_{k+1}^-) \end{bmatrix} \quad (4.2)$$

4.1 Propagation

Since each of the 3 robots moves independent of the others, the state (pose) propagation is provided by Equations (3.1). The same is not true for the covariance of the state estimate. In [21], we derived the equations for the propagation of the initial, fully decoupled system. Here we will examine how the Kalman filter propagation equations are modified in order to include the cross-correlation terms introduced after a few updates of the system. Starting from:

$$P(t_{k+1}^-) = \Phi(t_{k+1}, t_k) P(t_k^+) \Phi^T(t_{k+1}, t_k) + Q_d(t_{k+1}) \quad (4.3)$$

and substituting from Equation (4.2) we have:

$$P(t_{k+1}^-) = \begin{bmatrix} \Phi_1 P_{11}(t_k^+) \Phi_1^T + Q_{d1} & \Phi_1 P_{12}(t_k^+) \Phi_2^T & \Phi_1 P_{13}(t_k^+) \Phi_3^T \\ \Phi_2 P_{21}(t_k^+) \Phi_1^T & \Phi_2 P_{22}(t_k^+) \Phi_2^T + Q_{d2} & \Phi_2 P_{23}(t_k^+) \Phi_3^T \\ \Phi_3 P_{31}(t_k^+) \Phi_1^T & \Phi_3 P_{32}(t_k^+) \Phi_2^T & \Phi_3 P_{33}(t_k^+) \Phi_3^T + Q_{d3} \end{bmatrix} \quad (4.4)$$

Equation (4.4) is repeated at each step of the propagation and it can be distributed among the robots after appropriately splitting the cross-correlation terms. For example, the cross-correlation equations for robot 2 are:

$$\sqrt{P_{21}(t_{k+1}^-)} = \Phi_2 \sqrt{P_{21}(t_k^+)}, \quad \sqrt{P_{23}(t_{k+1}^-)} = \Phi_2 \sqrt{P_{23}(t_k^+)} \quad (4.5)$$

After a few steps, if we want to calculate the (full) cross-correlation terms of the centralized system, we will have to multiply their respective components. For example:

$$P_{32}(t_{k+1}^-) = \sqrt{P_{32}(t_{k+1}^-)} \sqrt{P_{23}(t_{k+1}^-)^T} = \Phi_3 \sqrt{P_{32}(t_k^+)} (\Phi_2 \sqrt{P_{23}(t_k^+)^T})^T = \Phi_3 \sqrt{P_{32}(t_k^+)} (\sqrt{P_{23}(t_k^+)^T})^T \Phi_2^T = \Phi_3 \sqrt{P_{32}(t_k^+)} \sqrt{P_{32}(t_k^+)} \Phi_2^T = \Phi_3 P_{32}(t_k^+) \Phi_2^T \quad (4.6)$$

This result is very important since the propagation Equations (3.1) and (4.5) to (4.6) allow for a **fully distributed estimation algorithm during the propagation cycle**. The computation gain is very large if we consider that most of the time the robots propagated their pose and covariance estimates based on their own perception while updates are usually rare and they take place only when two robots meet.

4.2 Update

If now we assume that robots 2 and 3 are exchanging relative position and orientation information, the residual covariance matrix:

$$S(t_{k+1}) = H_{23}(t_{k+1}) P(t_{k+1}^-) H_{23}^T(t_{k+1}) + R_{23}(t_{k+1}) \quad (4.7)$$

¹Due to space limitations the propagation and update equations of the Kalman filter before and up to the first update are omitted from this presentation. The interested reader is referred to [21] for a detailed derivation.

is updated based on Equation (4.2), for $H_{23}(t_{k+1}) = \begin{bmatrix} 0 & I & -I \end{bmatrix}$, as:

$$\begin{aligned} S(t_{k+1}) &= P_{22}(t_{k+1}^-) + P_{33}(t_{k+1}^-) \\ &- P_{32}(t_{k+1}^-) - P_{23}(t_{k+1}^-) + R_{23}(t_{k+1}) \end{aligned} \quad (4.8)$$

where $R_{23}(t_{k+1})$ is the measurement noise covariance matrix associated with the relative position and orientation measurement between robots 2 and 3. In order to calculate matrix $S(t_{k+1})$, only the covariances of the two meeting robots are needed along with their cross-correlation terms. All these terms can be exchanged when the two robots detect each other, and then used to calculate the residual covariance matrix S . The dimension of S is $N \times N$, the same as if we were updating the pose estimate of one robot instead of three. (In the latter case the dimension of matrix S would be $(N \times 3) \times (N \times 3)$). As we will see in Equation (4.9), this reduces the computations required for calculating the Kalman gain and later for updating the covariance of the pose estimate. The Kalman gain for this update is given by:

$$\begin{aligned} K(t_{k+1}) &= P(t_{k+1}^-) H_{23}^T(t_{k+1}) S^{-1}(t_{k+1}) = \\ &\begin{bmatrix} (P_{12}(t_{k+1}^-) - P_{13}(t_{k+1}^-)) S^{-1}(t_{k+1}) \\ (P_{22}(t_{k+1}^-) - P_{23}(t_{k+1}^-)) S^{-1}(t_{k+1}) \\ -(P_{33}(t_{k+1}^-) - P_{32}(t_{k+1}^-)) S^{-1}(t_{k+1}) \end{bmatrix} = \begin{bmatrix} K_1(t_{k+1}) \\ K_2(t_{k+1}) \\ K_3(t_{k+1}) \end{bmatrix} \end{aligned} \quad (4.9)$$

The correction coefficients (the matrix elements $K_i(t_{k+1})$, $i = 2, 3$, of the Kalman gain matrix) in the previous equation are smaller compared to the corresponding correction coefficients calculated during the first update [21]. Here the correction coefficients are reduced by the cross-correlation terms $P_{23}(t_{k+1}^-)$ and $P_{32}(t_{k+1}^-)$ respectively. This can be explained by examining what is the information contained in these cross-correlation matrices. As it is described in [21], the cross-correlation terms represent the information common to the two meeting robots acquired during a previous direct (robot 2 met robot 3) or indirect (robot 1 met robot 2 and then robot 2 met robot 3) exchange of information. The more knowledge these two robots (2 and 3) already share, the less gain can have from this update session as this is expressed by the values of the matrix elements of the Kalman filter (coefficients $K_i(t_{k+1})$, $i = 2, 3$) that will be used for update of the pose estimate $\hat{\mathbf{x}}(t_{k+1}^+)$. In addition to this, by observing that $K_1(t_{k+1}) = (P_{12}(t_{k+1}^-) - P_{13}(t_{k+1}^-)) S^{-1}(t_{k+1})$ we should infer that robot 1 will be affected by this update to the extent that the information shared between robots 1 and 2 differs from the information shared between robots 1 and 3.

Finally, as it is shown in [21], the centralized system covariance matrix calculation can be divided into $3(3 + 1)/2 = 6$, $N \times N$ matrix calculations and distributed among the robots of the group.²

²In general $M(M + 1)/2$ matrix equations distributed among M robots, thus $(M + 1)/2$ matrix calculations per robot.

5 Observability Study

5.1 Case 1: At least one of the robots has absolute positioning capabilities

In this case the main difference is in matrix H . If we assume that robot 1 for example has absolute positioning capabilities then the measurement matrix H and the observability matrix M_{DTI} would be:

$$H = \begin{bmatrix} I & 0 & 0 \\ I & -I & 0 \\ 0 & I & -I \\ -I & 0 & I \end{bmatrix}$$

$$M_{DTI} = \begin{bmatrix} I & I & 0 & -I & | & I & I & 0 & -I & | & I & I & 0 & -I \\ 0 & -I & -I & I & | & 0 & -I & I & I & | & 0 & -I & I & 0 \\ 0 & 0 & 0 & 0 & | & 0 & 0 & -I & I & | & 0 & 0 & -I & I \end{bmatrix}$$

The rank of the M_{DTI} matrix is 9 and thus the system is observable when at least one of the robots has access to absolute positioning information (e.g. by using GPS or a map of the environment).

5.2 Case 2: At least one of the robots remain stationary

If at any time instant at least one of the robots in the group remains stationary, the uncertainty about its position will be constant and thus it has a direct measurement of its position which is the same as before. This case therefore falls into the previous category and the system is considered observable. Examples of this case are the applications found in [12], [11], [10], [17], [18].

6 Experimental Results

The proposed *distributed multi-robot localization* method was implemented and tested for the case of 3 mobile robots. The most significant result is the reduction of the uncertainty regarding the position and orientation estimates of each individual member of the group.

The 3 robots start from 3 different locations and they move within the same area. Every time a meeting occurs, the two robots involved measure their relative position and orientation³. Information about the cross-correlation terms is exchanged among the members of the group and the distributed modified Kalman filters update the pose estimates for each of the robots. In order to focus on the effect of the *distributed multi-robot localization* algorithm, no absolute localization information was available to any of the robots. Therefore the covariance of the position estimate for each of them is bound to increase while the position estimates will drift away from their real values.

³The experiments were conducted in a lab environment with an overhead camera tracking the absolute poses of the 3 robots. The relative pose measurements were provided by the camera while white noise was added to each of them. The accuracy of the relative measurements was +/- 30cm for the relative position and +/- 17 degrees for the relative orientation.




Figure 1: Distributed multi-robot localization results: The covariances of the position x estimates for each of the three robots in the group.

At time $t=100$ robot 1 meets robot 2 and they exchange relative localization information. At time $t=200$ sec robot 2 meets robot 3, at $t=300$ sec robot 3 meets robot 1, and finally at $t=400$ sec robot 1 meets robot 2 again. As it can be seen in Figure 1, after each exchange of information, the covariances, representing the uncertainty of the position x estimates, of robots 1 and 2 ($t=100$ sec), 2 and 3 ($t=200$ sec), 3 and 1 ($t=300$ sec), and 1 and 2 ($t=400$ sec) is significantly reduced.

7 Discussion

At this point it is worth mentioning that a decentralized form of the Kalman filter was first presented in [22] and later revisited in its inverse (*Information filter*) formulation in [13] for sequential processing of incoming sensor measurements. These forms of the Kalman filter are particularly useful when dealing with asynchronous measurements originating from a variety of sensing modalities (an application of this can be found in [19]). The Information filter has certain advantages compared to the Kalman filter for specific estimation applications ([14]). For the case of the distributed multi-robot localization the Kalman filter is significantly better due to the reduced number of computations. The single matrix inversion required is of the residual covariance matrix $S(t_{k+1})$ (3×3) and this oc-

curs only when a relative pose measurement is available. The Information filter requires large matrix inversions at each propagation step. More specifically the information matrix propagation equation is:

$$P^{-1}(t_{k+1}^-) = M(t_{k+1}) - M(t_{k+1})G_d(t_k)$$

$$\left[G_d^T(t_k)M(t_{k+1})G_d(t_k) + Q_d^{-1}(t_k) \right]^{-1} G_d^T(t_k)M(t_{k+1}) \quad (7.10)$$

where

$$M(t_{k+1}) = \Phi^T(t_k, t_{k+1})P^{-1}(t_k^+) \Phi(t_k, t_{k+1}) \quad (7.11)$$

For a group of M robots, the matrix $G_d^T(t_k)M(t_{k+1})G_d(t_k) + Q_d^{-1}(t_k)$ of dimensions $(M \times 3) \times (M \times 3)$ has to be inverted during each propagation step and for a large group of robots this becomes computationally inefficient. In addition, the information filter produces estimates of $\hat{y}(t_{k+1}^+) = P^{-1}(t_{k+1}^+) \hat{x}(t_{k+1}^+)$ instead of $\hat{x}(t_{k+1}^+)$ and therefore the information matrix $P^{-1}(t_{k+1}^+)$ (of dimensions $(M \times 3) \times (M \times 3)$) must also be inverted in order to get the estimates of the poses of all the robots in the group.

References

- [1] J. Borenstein. Control and kinematic design of multi-degree-of freedom mobile robots with compliant linkage. *IEEE Transactions on Robotics and Automation*, 11(1):21–35, Feb. 1995.

- [2] J. Borenstein. Internal correction of dead-reckoning errors with a dual-drive compliant linkage mobile robot. *Journal of Robotic Systems*, 12(4):257–273, April 1995.
- [3] J. Borenstein. Experimental results from internal odometry error correction with the omnimate mobile robot. *IEEE Transactions on Robotics and Automation*, 14(6):963–969, Dec. 1998.
- [4] J. Borenstein and L. Feng. Gyrodometry: A new method for combining data from gyros and odometry in mobile robots. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, pages 423–428, 1996.
- [5] J. Borenstein and L. Feng. Measurement and correction of systematic odometry errors in mobile robots. *IEEE Transactions on Robotics and Automation*, 12(6):869–880, Dec. 1996.
- [6] I. J. Cox. Blanche-an experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Transactions on Robotics and Automation*, 7(2):193–204, April 1991.
- [7] H. R. Everett. *Sensors for Mobile Robots*. AK Peters, 1995.
- [8] M.S. Fontan and M.J. Mataric. Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation*, 14(5):815–822, Oct. 1998.
- [9] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. Collaborative multi-robot localization. In *In Proc. of the 23rd Annual German Conference on Artificial Intelligence (KI)*, Bonn, Germany, 1999.
- [10] R. Kurazume and S. Hirose. Study on cooperative positioning system: optimum moving strategies for cps-iii. In *Proceedings of the 1998 IEEE International Conference in Robotics and Automation*, volume 4, pages 2896–2903, Leuven, Belgium, 16-20 May 1998.
- [11] R. Kurazume, S. Hirose, S. Nagata, and N. Sashida. Study on cooperative positioning system (basic principle and measurement experiment). In *Proceedings of the 1996 IEEE International Conference in Robotics and Automation*, volume 2, pages 1421–1426, Minneapolis, MN, April 22-28 1996.
- [12] R. Kurazume, S. Nagata, and S. Hirose. Cooperative positioning with multiple robots. In *Proceedings of the 1994 IEEE International Conference in Robotics and Automation*, volume 2, pages 1250–1257, Los Alamitos, CA, 8-13 May 1994.
- [13] E.M. Nebot M. Bozorg and H.F. Durrant-Whyte. A decentralised navigation architecture. In *In Proceedings of the IEEE International Conference on Robotics and Automation*, volume 4, pages 3413–3418, Leuven, Belgium, 16-20 May 1998.
- [14] A.G.O. Mutambara and M.S.Y. Al-Haik. State and information space estimation: A comparison. In *Proceedings of the American Control Conference*, pages 2374–2375, Albuquerque, New Mexico, June 1997.
- [15] C.F. Olson and L.H. Matthies. Maximum likelihood rover localization by matching range maps. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pages 272–277, Leuven, Belgium, 16-20 May 1998.
- [16] L.E. Parker. Alliance: An architecture for fault tolerant multirobot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, April 1998.
- [17] I.M. Rekleitis, G. Dudek, and E.E. Milios. Multi-robot exploration of an unknown environment, efficiently reducing the odometry error. In M.E. Pollack, editor, *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, volume 2, pages 1340–1345, Nagoya, Japan, 23-29 Aug. 1997.
- [18] I.M. Rekleitis, G. Dudek, and E.E. Milios. On multi-agent exploration. In *Visual Interface*, pages 455–461, Vancouver, Canada, June 1998.
- [19] S. I. Roumeliotis, G. S. Sukhatme, and G. A. Bekey. Sensor fault detection and identification in a mobile robot. In *Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 1383–1388, Victoria, BC, Canada, 13-17 Oct. 1998.
- [20] S.I. Roumeliotis and G.A. Bekey. Bayesian estimation and kalman filtering: A unified framework for mobile robot localization. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, pages 2985–2992, San Francisco, CA, April 24-28 2000.
- [21] Stergios I. Roumeliotis. *Robust Mobile Robot Localization: From single-robot uncertainties to multi-robot interdependencies*. PhD thesis, University of Southern California, Los Angeles, California, May 2000.
- [22] H. W. Sorenson. *Advances in Control Systems*, volume 3, chapter Kalman Filtering Techniques. Academic Press, 1966.