

Experiments in Free-Space Triangulation Using Cooperative Localization

Ioannis Rekleitis¹, Gregory Dudek² and Evangelos Milios³

¹Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA *

²Centre for Intelligent Machines, McGill University, Montreal, Québec, Canada.

³Faculty of Computer Science, Dalhousie University, Halifax, Nova Scotia, Canada
email: rekleitis@cmu.edu, dudek@cim.mcgill.ca, eem@cs.dal.ca

Abstract

This paper presents a first detailed case study of collaborative exploration of a substantial environment. We use a pair of cooperating robots to test multi-robot environment mapping algorithms based on triangulation of free space. The robots observe one another using a robot tracking sensor based on laser range sensing (LIDAR). The environment mapping itself is accomplished using sonar sensing. The results of this mapping are compared to those obtained using scanning laser range sensing and the scan matching algorithm. We show that with appropriate outlier rejection policies, the sonar-based map obtained using collaborative localization can be as good or, in fact, better than that obtained using what is typically considered to be a superior sensing technology.

1 Introduction

In this paper we will present a case study that illustrates the particular trade-offs necessary to achieve both acceptable speed and good accuracy in the context of collaborative exploration [9]. Specifically, we provide detailed specifications of how collaborative implementation can be carried out in practice and we compare the results from laser-based scan matching to those from sonar-based mapping with collaborative exploration. In prior work we have defined collaborative exploration and associated algorithms in which a team of two or more robots coordinate their motion through a potentially unknown environment to jointly estimate one another's position and, in so doing, estimate the layout in the environment of any spatial parameter of interest. This prior work has dealt, primarily, with the

*The experimental data for this paper was acquired while the author was at McGill University, Montreal, Québec, Canada.



Figure 1: The two robots exploring a laboratory. In the lower-left one robot carries a three plane target, on the upper-right a second robot is mapping.

theoretical properties of the methods as opposed to actual performance issues.

The key to collaborative exploration, as we define it, is to have at least one “tracker” sensor that allows a robot to estimate the positions of other robots in the team. This allows inter-robot sensing to compensate for arbitrarily large odometry errors, as well as presenting other advantages [9, 11]. Our specific strategy for collaborative exploration as applied to a pair of robots is to have them take turns moving so that at any time one is stationary and can act as a fixed reference point. In this paper we consider the experimental validation of collaborative exploration, statistically robust modeling of error and map synthesis using (otherwise) uncertain sensing.

We estimate the positions of the robots using a particle filter that combines an open-loop estimate of odometry error with sensor data collected from the tracker, a LIDAR-based laser range finder³ on one robot and a three-plane target mounted on top of the second robot

(alternative implementations have been used in prior work). Figure 1 shows the two robots during the exploration of a laboratory. The three-plane target is visible from any orientation as can be seen in Figure 2 where the sensor data from the laser range finder are recorded over time. Parts of the walls are also mapped by the laser and provide ground truth for the cooperative exploration.

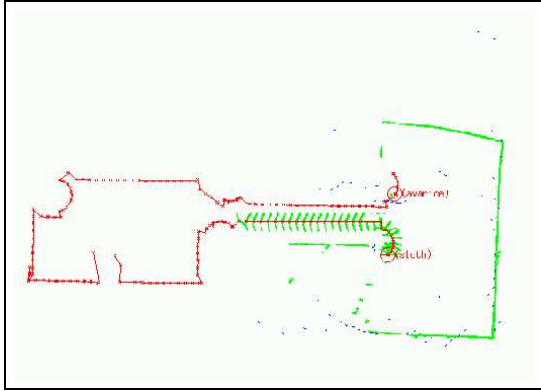


Figure 2: The trajectories of the two robots with the laser data also marked. Note the target pattern detected.

1.1 Background

As noted above, we use the collaborative-exploration-based triangulation algorithm developed earlier to build a map [10, 9]. To very briefly recapitulate the principle of this approach, the robots follow the walls of the environment while maintaining visual contact with one another. The wall following robustly traces the environment boundaries, while the line of visual contact between the robots both facilitates localization and also assures no (opaque) obstacles fail to be detected, no matter how good their camouflage (or how bad sensors are). In this algorithm, the robots take turns exploring an environment, such that one remains stationary while the other moves. When a reflex vertex is encountered (i.e. a vertex that hides one robot from the other), the moving robot backs up until the line of sight is reestablished, and then the previously stationary robot starts following the wall it is beside. This algorithm is assured to produce a complete map of the environment with a triangulation of free space as well as a dual graph of the triangulation that can be used for subsequent tasks such as path planning. The pose estimate during exploration is maintained using a particle filter (described in detail elsewhere [12]), a Monte-Carlo Simulation technique [2], that has gained popularity recently under different names in different

fields. The technique we use was introduced as particle filtering by Gordon *et al.* [4]. In mobile robotics particle filtering has been applied successfully by different groups [1, 6, 14] as an alternative to the traditional Kalman Filter estimator [13].

1.2 Wall Following and Mapping

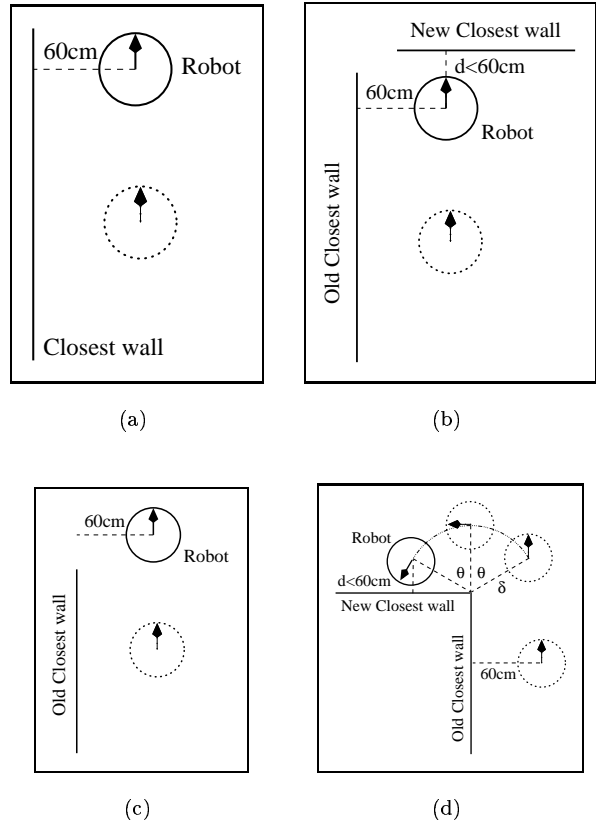


Figure 3: Illustration of the wall following behavior, where the robot is shown by a circle and its past position is shown by a dashed circle. (a) So long as only one wall is visible, follow it at a nominal 60cm distance. (b) When new wall is discovered approach it until it is closer than the one being followed. If it is connected by a non-reflex vertex (an angle under 180 degrees) map the corner and continue the exploration of the new wall. (c) If the wall being followed ends, it must be a reflex vertex (as we assume the world is composed of bounded objects); find the adjacent wall as in (d). (d) Map the reflex vertex by moving around the end-point of the old closest wall. Move by intervals of θ at a distance less than 60cm and more than 30cm from the end-point.




Figure 4: (a) The cooperative localization estimates (*) versus the trajectory resulting from the motion commands (+) given to robot R0 during the exploration. (b) The same for robot R1. (c) The triangulation map produced using only the sonar for mapping. The trajectory of Robot 0 is marked in magenta and the trajectory of Robot 1 is marked black. The walls are displayed in red and their lengths (in cm) is marked next to them. The internal diagonals that define the triangulation are marked as dashed lines (shown as blue on color renditions).

The utility of surface following is often underestimated; its successful achievement is important to the success of this approach. As such, we will discuss its implementation in uncharacteristic detail. When the line of sight between the two robots is uninterrupted, the moving robot explores the environment one triangle at a time by following the closest wall from one corner (end point) to the other. In our implementation of the algorithm the robots follow the walls at a distance Δ using a sonar range finder in order to sense the wall. Lines are fit to the sonar points using recursive split and merge fitting [3, 8], and then the newly sensed lines are merged with the existing map.

An outline of the mapping procedure follows. We presuppose that the environment is a polygon (within some linearization error). The robot follows each wall so long as it remains straight (see Figure 3a, robots drawn with dashed lines represent past positions) and the robot has not moved past its end (the robot’s position projects within the line segment of the wall). If a new wall is detected at distance less than $\Delta = 60\text{cm}$ (see Figure 3b) then a non-reflex vertex must have been reached and the old wall has been fully mapped.

If the closest wall is unchanged but the robot has moved past the end of it (see Figure 3c) then this indicates that the robot has reached a reflex vertex. In order to map the reflex corner the following procedure is used (see Figure 3d). The robot moves in a circular path at a distance $\delta = 0.75\Delta$ from the end point of the closest wall until it finds a new closest wall. The circular motion is done in steps defined by an angle θ , see Figure 3d (in the current implementation $\theta = 15^\circ$). The old closest wall becomes fully mapped and then, if the reflex vertex does not interrupt the line of visual contact between the two robots, the moving robot continues the exploration by mapping the new closest wall. Oth-

erwise the line of visual contact is interrupted and the robots follow the triangulation algorithm ¹.

One limitation of the wall-following algorithm is the mapping of small walls; especially when the robot goes around a reflex corner, the adjacent walls should be minimum 50cm long. This limitation results from our choice of Δ (set distance between robot and wall).

2 Experimental Results

The exploration algorithm used for the mapping of an indoor environment is based on the triangulation of free space by two robots. The line of visual contact is used to “sweep” the space; in other words, if the two robots can observe each other then the space in between them is empty. When one robot is stationary at a corner of the environment and the other robot moves along a wall (without losing visual contact) then a triangle of free space is mapped. By constructing an on-line triangulation of the free space the two robots map the environment completely without any overlaps. Figure 1 presents the two robots at the early stages of the exploration.

The positional error is maintained low throughout the exploration by the use of cooperative localization. Figure 4a,b presents the pose estimates during the exploration when cooperative localization was used (marked as green “+” ²) together with the position of the robot estimated using the recorded motion commands (marked as blue “*”); the map of the environment is drawn in red. The left figure presents the trajectory of Robot 0 and the right figure presents the trajectory of Robot 1. Even though the actual trajectory

¹The complete description of the algorithm is outside the scope of this paper (please refer to previous work[10, 9]).

²We include the colour for the readers of colour reproductions of this paper




Figure 5: The two robots exploring a convex area. Both sub-figures show a composite representation of the succession of particle clouds modeled along the course of the trajectory. Note that the height of each peak represents accuracy (high peak more accurate estimate) (a) The trajectory of Robot 0. (b) The trajectory of Robot 1.

of each robot was kept in an almost straight line and closely corresponds with the cooperative localization estimates, the motion *commands* show a systematic bias (illustrated with ‘blue *’ in Figure 4a,b). The observed drift corresponds to the odometry error during the exploration.

Figure 5 presents experimental results from an indoor environment in the corridors of our building using two SuperScout mobile robots from Nomadic Technologies, Inc. The probability density function (*pdf*) of the pose of each robot is plotted for the entire trajectory. At each step the set of particles has been spatially integrated and then added to the plot (the higher the peak the more accurate the pose estimate). Figure 5a presents the trajectory of Robot 0 which is equipped with the laser range finder. Robot 1 is equipped with the three-plane target. The robot’s pose *pdf* can be seen in Figure 5b. The final map is presented in Figure 4c.

2.1 Sonar vs Laser

In this section we discuss the strengths and weaknesses of the two sensors used (sonar and laser range finders) and how they affected the resulting map.

The laser range finder was used during the exploration only for tracking the other robot; thus the map produced from the triangulation algorithm is constructed solely by the sonar sensor data calculated using the corrected poses of the two robots. To further validate our approach, the laser data were collected and pro-

cessed off-line and compared with the resulting map from the sonar data of the wall following in the triangulation algorithm. We fused the recorded laser data using the scan matching algorithm of Lu and Milios, a technique based on least-squared minimization of the distance among all observations aligned based on robot motions [7]. The scan matching was accomplished using Gutmann’s Scanstudio [5] implementation. Figure 6a presents the laser data that were observed and Figure 6b presents the same data after scan matching.

Figure 6 also illustrates a weakness of the laser sensor: the robot observes inconsistent object locations, even after the alignment of scans. In particular, the upper wall on the corridor appears to have a (fictitious) opening in it (see section 3 for further discussion). Because the laser range finder senses in a plane parallel to the wheels, in a perfectly flat terrain it would map all the obstacles at its height. In practice, even office floors are not perfectly flat and measurements out of the horizontal plane are a serious issue, especially for distant objects.

A weakness of the sonar sensor can be seen in Figure 7. The sonar data used for wall following were collected during the exploration. Figure 7 presents the sonar points (in blue): the left column presents the data gathered by Robot 0 while the right column presents the data gathered by Robot 1; the trajectories of the two robots are marked in red, and the positions from where the sonar scans were taken are marked with ‘*’. During the exploration the robots follow the walls at distance $\Delta = 60cm$. Figures 7a,b and 7c,d present the

sonar points detected in less than 65cm and 120cm respectively. Figure 7c,d is the data actually used during the exploration. As we saw in section 1.2, the sonar points are fit with line segments and then the lines are merged together; Figure 7c,d show that the sonar data filtered at 120cm are aligned with the walls. Figures 7e,f and 7g,h present the sonar scans filtered at 250cm and 400cm .

Clearly, the sonar data obtained are very noisy and, if the robots believed the occupancy of space from them, navigation would be impossible. It is worth noting some straight lines formed inside open space (see Figure 7g). Such lines correspond to small anomalies on the floor at the borders of the tiles in our laboratory (see Figure 1 for the appearance of the floors). As can be seen in Figures 7a-d there is virtually no distortion of the data. This is due to our cooperative localization approach which maintains an accurate position for the robots.

Finally, Figure 4c presents the map created by the multi-robot triangulation algorithm using only sonar data to measure the walls. Most significantly the map obtained from cooperative localization and sonar sensing has correct measurements that are as good as those from the scan matching with laser data, but with fewer discrepancies and outliers. This is illustrated in the maps of Figures 6b and 4c, in which the non-outlier data is consistent. Of course, the tracker sensor used for the collaborative mapping is based on the same LIDAR sensor used to obtain the scan matching data, but, as noted earlier (and demonstrated in prior work) this is just one of many possible implementations of the tracker.

2.2 Map Quality

The resulting map of the exploration can be seen in Figure 4c. The trajectory of Robot 0 is marked in magenta and the trajectory of Robot 1 is marked black. The walls are displayed in red and their length (in cm) is indicated. The internal diagonals that define the triangulation are marked as blue dashed lines. Moreover the lengths of the walls were measured manually (by measuring tape) and the results are presented in table 1 together with the estimated length from the triangulation map and the difference between the two measurements. The mean error was 4.6cm per wall. The perimeter of the environment mapped was measured to 42.71m while the perimeter of the resulting map was 42.63m . The angle of the two walls of the upper right corner was measured to 98° and the map estimate is 97° .

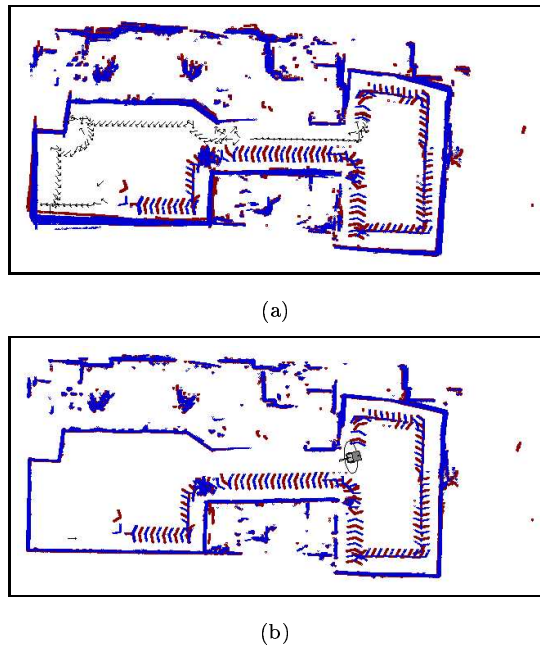


Figure 6: (a) The laser data collected during the exploration. (b) Map obtained after correcting data using scan matching. In both cases, compare the data from this “almost ideal” sensor to that obtained using sonar (Fig. 4c).

3 Discussion and Conclusions

In this paper we demonstrated the practical feasibility of collaborative exploration in mapping an unknown environment. The practical realization of this theoretical algorithm involved several design choices upon which its feasibility depended. These include the mechanism for wall following, the use of a heuristic wall-synthesis mechanism and an outlier rejection policy. Statistical estimation of the robot pose using particle filtering (described in [12]) was also an important ingredient.

The experimental results verify the improvement in the map accuracy over what could be obtained without cooperative localization. Areas of $13\text{m} \times 5\text{m}$ and $12\text{m} \times 9\text{m}$ were mapped completely, with a mean error less than 5cm . The perimeters of the environments mapped were of the order of $42 - 44\text{m}$. These results demonstrate the practical feasibility of the algorithms used, and illustrate some of the performance characteristics that had been predicted.

In these experiments a general-purpose laser-based sensor was used, but in an application context a special

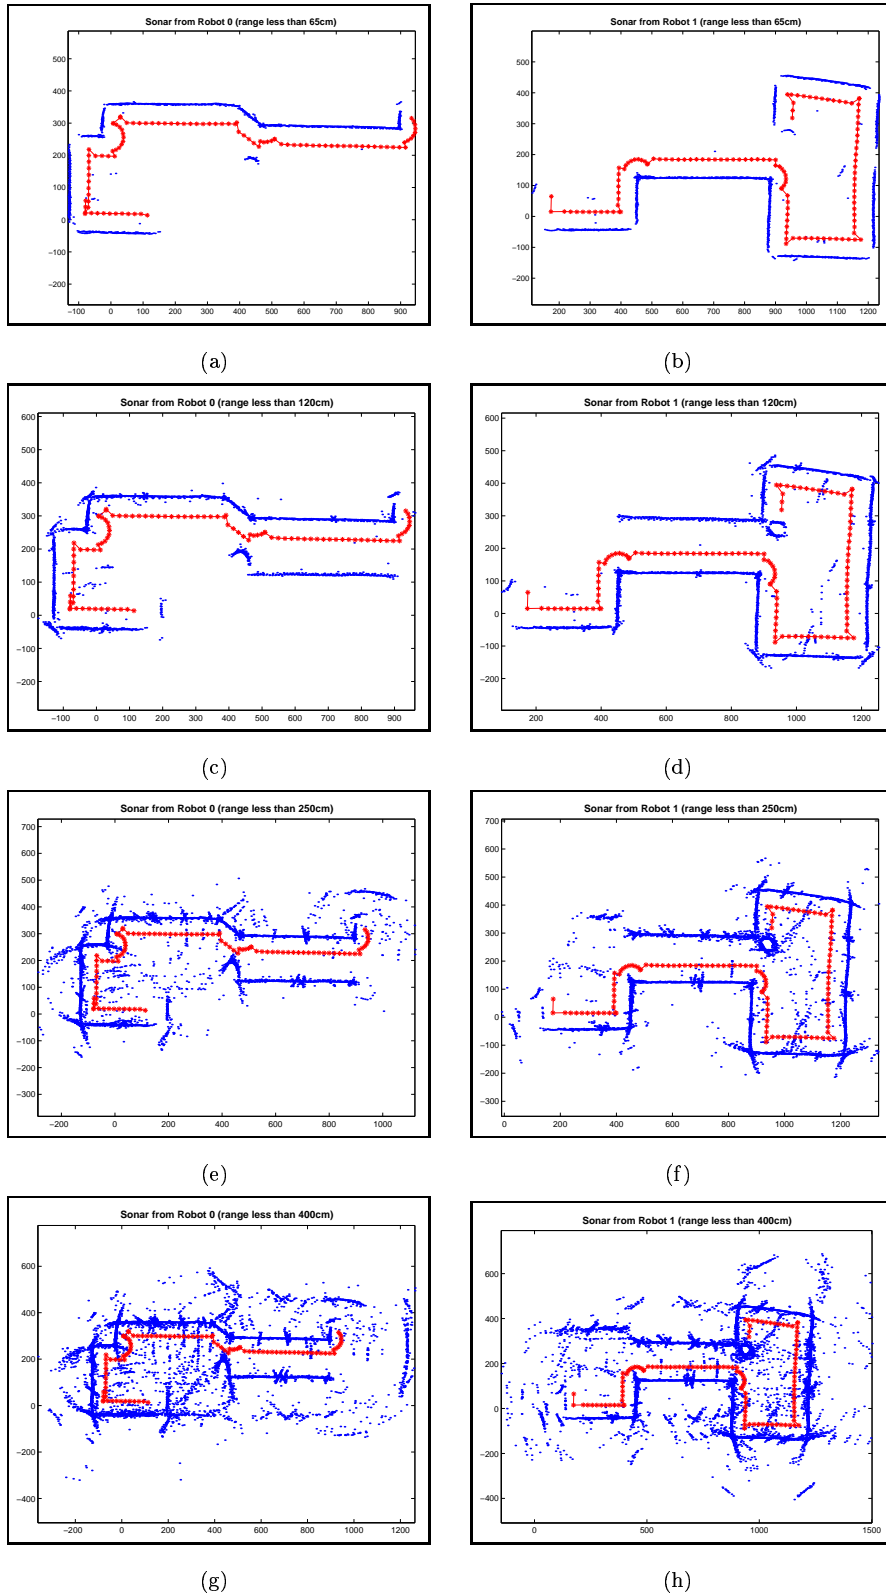


Figure 7: Sonar data collected from Robot 0 (on the left) and Robot 1 (on the right), filtered at different ranges: (a,b) 65cm; (c,d) 120cm; (e,f) 250cm; (g,h) 400cm. The sonar points are marked blue (color reproductions) and the path of the robot is marked red.

Manual	Map	Error	Manual	Map	Error
588	577.5	0.5	179	178.0	1.0
305	296.5	8.5	341	333.4	7.6
99	96.0	3.0	545	548.0	3.0
102	102.6	0.6	343	343.5	0.5
410	403.1	6.9	241	249.9	8.9
99	114.4	15.4	432	427.9	4.1
419	419.7	0.7	168	173.0	5.0

Table 1: The length of the walls measured with tape and from the triangulation map (first two columns) together with the error (the unit is cm). See Figure 4c for the correspondence between walls and lengths.

purpose laser-based tracker could be made substantially more economically. Further, several alternative implementations of the tracker are possible using vision, “active” sonar or other technologies. The insight is that the collaborative approach allows robustness and accuracy to be focused on the design of a suitable tracker, which is a constrained measurement problem between two controllable environment-independent devices, as opposed to having to design a robust and accurate sensor for arbitrary environmental structures and surface properties.

The laser sensor was valuable as part of the robot tracker sensor but, since it has a planar scanning sensor, it was not possible to utilize it at its full range because of the floor inclination and, perhaps, the conjecture that the robot itself was not perfectly level. During one series of experiments the target was not detected at a distance of 7m. Moreover, the laser range finder would miss any obstacle above or below the scanning plane and thus has limitations when used alone for navigation. While true volume scanning would be attractive, it remains a prohibitive option for many applications. To compensate for the shortcomings of the laser sensor, a hybrid system of vision and laser could be used to first locate the other robot, after which the laser sensor would be used to accurately estimate the pose.

The most striking result of this study was the observation that the map constructed from sonar data in the context of collaborative exploration was not only highly accurate, but was apparently even better and more use-

ful than a map of the same environment obtained from laser range (LIDAR) data and scan matching. The above observation does not mean that the laser data is particularly errorful, but rather that we fully utilized the strengths of sonar and laser sensors while we used each sensor to compensate for the weaknesses of the other. These experiments demonstrate that cooperative localization via laser robot tracker allows for dense accurate pose estimation while sonar sensing in close proximity to obstacle boundaries along with a very conservative outlier rejection policy provides a highly robust sensing methodology.

References

- [1] Frank Dellaert, Wolfram Burgard, Dieter Fox, and Sebastian Thrun. Using the condensation algorithm for robust, vision-based mobile robot localization. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE Press, June 1999.
- [2] Arnaud Doucet, Nando De Freitas, and Neil Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, Series Statistics for Engineering and Information Science, January 2001.
- [3] Gregory Dudek and Paul MacKenzie. Model-based map construction for robot localization. In *Proceedings of Vision Interface 1993*, Toronto, ON, May 1993.
- [4] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings For Radar and Signal Processing*, 140(2):107–113, April 1993.
- [5] J.-S. Gutmann, W. Burgard, D. Fox, and K. Konolige. An experimental comparison of localization methods. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1998.
- [6] Patric Jensfelt, Olle Wijk, David J. Austin, and Magnus Andersso. Feature based condensation for mobile robot localization. In *IEEE International Conference on Robotics & Automation (ICRA)*, pages 2531–2537, San Francisco, CA, USA, April 2000.
- [7] Feng Lu and E. Milios. Robot pose estimation in unknown environments by matching 2d range

scans. *Journal of Intelligent and Robotic Systems*, pages 249–275, 1998.

- [8] Paul MacKenzie and Gregory Dudek. Precise positioning using model-based maps. In *Proceedings of the International Conference on Robotics and Automation*, San Diego, CA, 1994. IEEE Press.
- [9] Ioannis M. Rekleitis, Gregory Dudek, and Evangelos Miliotis. Graph-based exploration using multiple robots. In *5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, pages 241–250, Knoxville, Tennessee, USA, October 4-6 2000. Springer. <http://www.cim.mcgill.ca/~yiannis/Publications/darsU131.pdf>.
- [10] Ioannis M. Rekleitis, Gregory Dudek, and Evangelos Miliotis. Multi-robot collaboration for robust exploration. *Annals of Mathematics and Artificial Intelligence*, 31(1-4):7–40, 2001. <http://www.cim.mcgill.ca/~yiannis/Publications/journal.pdf>.
- [11] Ioannis M. Rekleitis, Gregory Dudek, and Evangelos Miliotis. Multi-robot cooperative localization: A study of trade-offs between efficiency and accuracy. In *IEEE/RSJ/ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, October 2002. IEEE/RSJ. <http://www.cim.mcgill.ca/~yiannis/Publications/iros02pdf.pdf>.
- [12] Ioannis M. Rekleitis, Gregory Dudek, and Evangelos Miliotis. Probabilistic cooperative localization and mapping in practice. In *Proceedings of International Conference in Robotics and Automation*, page to appear, Taipei, Taiwan, May 2003. <http://www.cim.mcgill.ca/~yiannis/Publications/icra03.pdf>.
- [13] S.I. Roumeliotis and G.A. Bekey. Distributed multi-robot localization. *IEEE Transactions on Robotics and Automation*, 18(5):781–795, Oct. 2002.
- [14] Nikos Vlassis, Bas Terwijn, and Ben Krose. Auxiliary particle filter robot localization from high-dimensional sensor observations. In IEEE, editor, *IEEE International Conference in Robotics and Automation*, volume 1, pages 7–12, May 2002.