CoCoA: Coordinated Cooperative Localization for Mobile Multi-Robot Ad Hoc Networks

Dimitrios Koutsonikolas Saumitra M. Das Y. Charlie Hu Yung-Hsiang Lu C.S. George Lee
School of Electrical and Computer Engineering, Purdue University, USA
{dkoutson, smdas, ychu, yunglu, csglee}@purdue.edu

Abstract

Mobile robot teams are particularly suited to applications where infrastructure is unavailable or damaged since they can be used to quickly form an infrastructure-less mobile ad hoc network performing tasks in a distributed manner. In such deployment scenarios, localization of individual robots is essential for enabling many applications. In this paper, we propose CoCoA, Co*ordinated* Co*operative localization for mobile multi-robot* A*d Hoc Networks, that is low-cost, quickly deployable, energy-efficient, and provides reasonable accuracy for the applications. In CoCoA, only a subset of the robots in the mobile robot team are equipped with external localization devices (e.g. GPS, laser rangers, vision sensors). Subsequently, while robots perform their tasks, they cooperate to help localize each other without requiring static landmarks to be deployed. In addition, CoCoA coordinates the localization operations of robots to reduce energy consumption. Using detailed simulations and localization models calibrated from experimental data, we find that CoCoA is effective in reducing energy consumption while providing good localization accuracy.*

1. Introduction

In many applications, mobile multi-robot teams need to be deployed quickly in an ad-hoc manner and work without infrastructure support. Examples of such applications are exploring remote terrains, disaster relief or in applications that target low costs. In most such application scenarios, localization is a critical system service required for the mobile robot teams to operate effectively and efficiently. For example, in search and rescue operations, the location of a survivor needs to be indicated so that additional personnel can be dispatched to the area. Localization can also enable efficient communication among the mobile multi-robot team through geographic routing. In this paper, we explore how to provide a localization system for mobile robot teams that work in *infrastructure-less* scenarios.

There are typically two extreme solutions to localization in such systems. On one extreme, all the robots in the mobile robot team could potentially be equipped with localization devices (e.g. GPS, laser rangers, vision sensors) so that no cooperation is necessary for the purpose of localization. However, outfitting such devices on *all* robot nodes increases system costs and energy drain.

On the other extreme, each robot in the mobile robot team could be provided with its initial position when deployed, and rely on odometry to localize itself afterwards. Given an initial position estimate, odometry uses sensors on robot wheels to estimate a new position based on the wheel movement characteristics. However, this technique is known to suffer from accumulated localization errors, and the localization error is likely to be exacerbated by the uneven surfaces encountered in many application scenarios. Improving the performance of this technique would require more precise but very costly odometry devices and there are limits to the accuracy possible even with costly devices.

In this paper, we propose the CoCoA system that takes a middle-ground between these two extremes. In CoCoA, only a subset of nodes in the mobile robot team is equipped with localization devices in order to reduce costs. Subsequently, CoCoA makes the robots *cooperate* in localizing each other using RF beacons through cheap commodity IEEE 802.11 wireless cards. Mobile robot teams in many cases already have such wireless devices for inter-robot or robot-to-operator communication and they can be reused for CoCoA as well. It is well known that wireless transmissions consume energy heavily. Thus, CoCoA also proposes to *coordinate* this RF beaconing in order to reduce the energy usage of the mobile robot team spent on localization. The coordination is achieved using the Mobile Robot Mesh Multicast (MRMM) [1] protocol which itself operates in an infrastructure-less environment. We evaluate CoCoA using detailed simulations and localization models calibrated from experimental data.

[1]This work was supported in part by the National Science Foundation under Grant IIS-0329061.

0-7695-2541-5 /06 $20.00 © 2006 IEEE

2. CoCoA Design

In this section, we describe the design of CoCoA. We first state our assumptions about the mobile robots in the team as well as the RF communication mechanism. We then describe the two key components of CoCoA: (1) the cooperative localization algorithm, and (2) the coordination mechanism for energy efficiency.

2.1. Assumptions

We focus on an application scenario where many robots are used to form a mobile robot network. Each robot has a simple sensory ability and limited computational power. This makes it practical to build a large number of such robots. The communication among the robots is based on cheap open-license wireless communication (IEEE 802.11b) at 2.4 Ghz. This reflects our mobile robot testbed [3] in which 8 Pioneer mobile robots [21] are equipped with laptops/PDAs and use Orinoco Wavelan 802.11b wireless cards for communication. In addition, our testbed has 4 laser rangers that can be used through the laptop/PDA for robot applications. Thus, we assume that a subset of nodes are equipped with laser ranger receivers that provide localization through a SLAM (Simultaneous Localization and Mapping) algorithm. The size of this subset affects the tradeoff between cost and energy versus localization accuracy and we study this in the evaluation.

2.2. Cooperative Localization

The first key component of CoCoA is a cooperative beacon-based localization scheme. In CoCoA, robots that have localization devices transmit beacons while performing their tasks and robots that do not have localization devices receive these beacons to localize themselves. Whenever a robot without a localization device receives an RF beacon from another robot, it executes a localization algorithm, in order to refine its position estimate. This algorithm was first proposed by Sichitiu et al. [17] for localization in *static* sensor networks. It uses the Received Signal Strength Indicator (RSSI) for ranging and Bayesian inference to estimate the positions of the unknown nodes. We apply this algorithm to mobile robot nodes.

Before running the algorithm, an offline calibration phase is needed to construct the PDF Table. This table is stored at each node and maps every RSSI value to a Probability Distribution Function (PDF). According to the algorithm, the mobile landmark periodically broadcasts beacon packets as it traverses the deployment area. These packets contain the coordinates of the mobile landmark, which can be obtained by GPS. When a node receives a beacon packet,

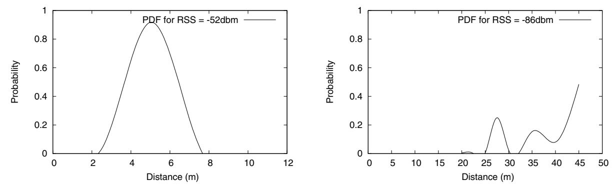

(a) PDF for RSS = -52dbm (b) PDF for RSS = -86dbm

Figure 1. Probability distribution functions (PDFs) for two different RSS values.

it performs a lookup at the PDF Table and obtains the probability distribution function of the distance corresponding to the RSSI of the beacon packet. Using this function, the sensor imposes a constraint on its position estimation.

Bayesian inference is then applied and the new position estimate is computed for the node being localized, based on the old position estimate and the new constraint. This process is repeated for each received beacon packet. Finally, when the node stops receiving any more beacon packets, either because the mobile landmark has moved away, or because a maximum number of beacons has been received, the node computes its position coordinates as a weighted average of its last position estimate over the whole deployment area.

Experimental Verification: The above algorithm assumes that, for each signal strength value, the probability distribution function of this value versus distance is Gaussian. In our outdoor experimental tests with our mobile robot equipment, we have found that this assumption correctly models the real world. One example of this function is shown in Figure 1(a) for RSSI = -52dbm. However, it is interesting to see that this model holds only for signal strength values larger than -80dbm, which correspond to physical distances of up to 40 meters. Beyond this distance the noise in the signal strength measurements fluctuates due to multipath and fading and the probability distribution function of the signal strength versus distance can no longer be approximated by a Gaussian, as shown in Figure 1(b) for RSSI = -86dbm. Note that this reflects our particular hardware and may not be general for all configurations. However, for most 802.11b cards which typically have a transmission range of more than 150m, the Gaussian assumption is likely to hold for up to 40m distances at which the signal strength will be high.

2.3. Energy-efficient Coordination

The second key component of CoCoA is the energy-efficient coordination of the localization algorithm. Wireless transmissions of data are known to consume much more energy than computation. Apart from limiting the number

Proceedings of the 26th IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW'06)
0-7695-2541-5 /06 $20.00 © 2006 IEEE

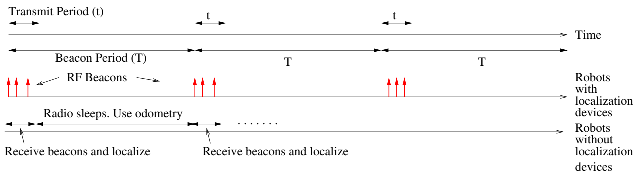

Figure 2. Energy-efficient coordination in CoCoA.

of wireless transmissions, another fundamental technique to conserve energy is to wake and sleep periodically. This is useful because typical 802.11 radios (e.g. the Lucent Wavelan – now Orinoco) consume as much energy being idle as when receiving packets [4]. Thus, significant energy savings are only possible if radios are put in sleep mode instead of idle mode (50mW versus 900mW). However, this requires that nodes be synchronized to some extent. We design CoCoA to require only coarse-grained synchronization achievable through wireless communication among the mobile robots. Once synchronized, the nodes coordinate their wake-sleep periods such that energy is conserved. The time-line of CoCoA operations is depicted in Figure 2. As shown in the figure, time is divided into successive beacon periods (T). Within each beacon period, a small fraction is set aside as a transmit period (t). Both the robots with localization devices and the robots without such devices are coarsely synchronized to follow this time-line. The functions that each type of robot performs are also depicted in Figure 2.

The robots with localization devices transmit k RF beacons during the period t after every beacon period T. k beacons are used for increasing the reliability of beacon delivery. Our evaluation uses k=3. The RF beacon is sent via UDP broadcast. Each beacon, in addition to the IP and UDP headers (20 bytes each), contains the location (x and y coordinates) of the sending robot obtained from the localization devices. The robots without localization devices wake up during each transmit period and use the beacons to execute the localization algorithm presented in the previous section. Following this transmit period, the robots without localization devices use odometry to estimate their position until the next transmit period. At that point the robots that receive at least three beacons throw away their currently estimated positions and find a new position using the new beacons. If certain robots do not receive any beacons, they continue with their old estimated position from the previous beacon period.

To achieve coarse-grained synchronization, we use MRMM (Mobile Robot Mesh Multicast), a protocol specifically designed to provide multicast operation in mobile robot networks, proposed in [1]. This allows us to reuse the 802.11 wireless devices used for localization to also perform synchronization. Due to lack of space we omit a complete description of the protocol and refer the reader to our paper [1]. Briefly, MRMM allows CoCoA to build a virtual mesh of robots from which control packets for synchronization can be delivered to each robot in an energy-efficient manner.

[2] The wake-sleep periods regulated by CoCoA can be easily adapted to accommodate scenarios when the radios need to be awake because of the application tasks.

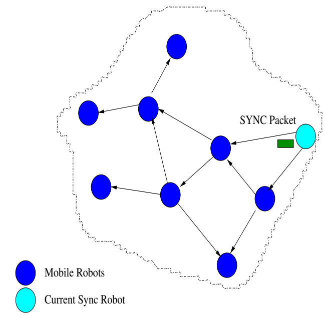

Figure 3. Coarse grained synchronization in CoCoA.

CoCoA synchronization using MRMM is depicted in Figure 3 and works as follows. At the beginning of the deployment, one robot designated as the SYNC robot, begins to send SYNC messages using the data delivery mechanism of MRMM. These SYNC messages are sent at the beginning of every beacon period by the Sync robot. A SYNC message contains the periods T and t. Each node upon receiving a SYNC message, sets up its internal timers for sleeping and waking up the radio according to the specified values of T and t in the packets. This allows a human operator to dynamically adjust these values to depend on the application by notifying the Sync robot to advertise new values. The SYNC message is broadcast down a sparse mesh (depicted in Figure 3), set up dynamically by MRMM, until it reaches every mobile robot. The mesh automatically adjusts to mobility, thus allows CoCoA synchronization to work with mobility as well.

3. Performance Evaluation

3.1. Simulation Methodology

We use the Glomosim simulator [20] to evaluate the performance of CoCoA. Glomosim is a widely used mobile wireless network simulator with a detailed and accurate physical signal transmission model. We implemented the robot movement model as well as an odometry model (described below) in Glomosim.

Movement-Odometry Model: The movement model used for the robots is as follows. As the simulation starts, each robot is given a random command to move to a random destination in the given area and starts moving towards the chosen destination with a speed chosen uniformly between 0.1 and v_{max} meters/second. We vary v_{max} as 0.5m/s and 2.0m/s to study its impact on the performance of CoCoA.

Proceedings of the 26th IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW'06)
0-7695-2541-5 /06 $20.00 © 2006 IEEE

Once a robot reaches the destination, it is given a new random command (to move to a new destination with a newly chosen speed). This effectively models the movement of a group of robots performing tasks: each robot moves towards a particular area, performs a task, and then moves to the next position. We implemented an odometry model in Glomosim to simulate the robots actual movement with errors in position estimation. We assume odometry displacement error to be zero-mean Gaussian with variance 0.1m/s and assume the angular odometry error to also be zero-mean Gaussian with variance of $10°$.

Energy Model: Since our wireless communication is based on IEEE 802.11, we adopt the energy model and measurements of IEEE 802.11 wireless cards in [4]. We use a wireless network interface with a 2 Mbps bandwidth. Our simulations model the energy spent in transmission, reception, idling and sleeping along with the energy spent in powering the card on and off.

Metrics: The following metrics are evaluated for the CoCoA system: (1) Localization error – The distance between the real position of the robot and where the robot estimates itself to be; and (2) Energy consumption – The total amount of energy consumed by the the robot team for communication and localization. This includes the energy spent during sending and receiving both data and control packets as well as when the wireless device is idle or in sleep mode.

3.2. Simulation Results

In this section we evaluate the localization performance of CoCoA and investigate the impact of the different configuration parameters on energy saving. We assume a total of 50 mobile robots operating in an area of size of $40000m^2$. Unless otherwise stated, half of the robots are equipped with localization devices and the localization error is not reported for those robots. The transmission range is set to 40m. The robots move according to the movement model defined above. The simulated time is 30 minutes.

3.2.1. Localization Error

In this section, we evaluate the performance of CoCoA, which combines RF localization and odometry and compare it against two extreme approaches. One of them is the case when all the robots are provided with their initial coordinates, but during the whole simulation they only use odometry to maintain an estimate of their position (e.g., no other localization devices are used). The other one is the case when the robots without localization devices use the localization algorithm described in Section 2.2 to get an estimate of their positions in every period (T sec), but this estimate remains the same between two successive periods (e.g., no odometry is used). For this comparison we selected T = 100 sec (shown in Section 3.2.2 to be optimal). The evaluation results are shown in Figures 4(a), 4(b) for the two different maximum speeds.

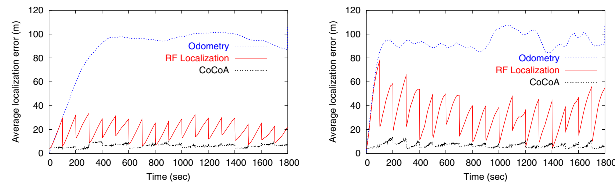

(a) maximum speed 0.5 m/s (b) maximum speed 2.0 m/s

Figure 4. Localization error over time for T = 100 seconds using (i) only odometry, (ii) only the RF localization algorithm and (iii) CoCoA (using both).

Figures 4(a), 4(b) show that the average localization error when only odometry is used increases significantly over time and it is much larger, compared to the other two approaches. The reason for this is that the odometry error is accumulated over time. Odometry is useful only for a small period of time, after which, robots need to update their position estimates using external devices.

Compared to using odometry, the localization accuracy is significantly improved by using RF localization. We observe that the localization error becomes minimum in the beginning of every period T, and it increases with time, as the position estimates become stale.

To keep the error low even between two successive beacon periods, CoCoA combines RF localization with odometry. In figures 4(a), 4(b) we observe that CoCoA offers significantly higher accuracy, compared to the other two approaches, since it combines the advantages of both. For example, the average localization error over time for a maximum speed 2 m/s is only 6.5m with CoCoA, while it is about 33m when only the localization algorithm is used and about 90m when only odometry is used.

Figure 5 shows the Cumulative Distribution Function (CDF) of the localization error in CoCoA for three different time instances (in the end of a beacon period just before the start of the next transmit period, in the end of a transmit period after localization has been completed, and in the middle of a beacon period while the radio is sleeping, i.e., $\frac{T}{2}$ seconds after the transmit period is over). As expected, the localization is best right after beacons are received at time 804s. Since T=100s, the time instant 804s is right after a transmit period has occurred. The CDF also shows that the location accuracy does deteriorate over time but not significantly. These results demonstrate that CoCoA not only improves the overall average localization error but also reduces the localization error of a large fraction of the nodes. For example, Figure 5 shows that more than 90% of the robots have a localization error lower than 10m.

Proceedings of the 26th IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW'06)
0-7695-2541-5 /06 $20.00 © 2006 IEEE

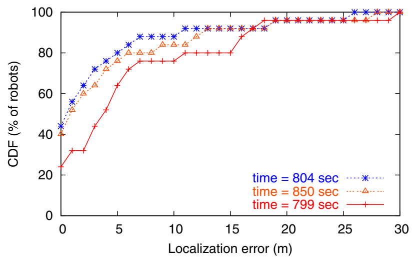

Figure 5. CDF for the localization error at three time instances.

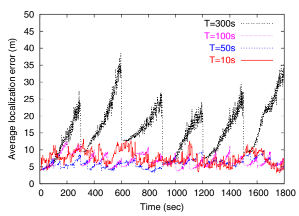

(a) Impact of T on localization error

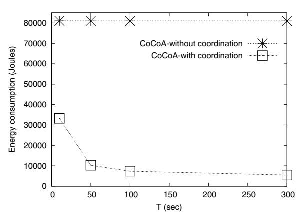

(b) Impact of T on energy consumption

Figure 6. Impact of beacon period (T) on CoCoA localization error and energy consumption. 50% of the nodes are equipped with localization devices.

3.2.2. Parameter tuning

In addition to improving localization accuracy, CoCoA reduces the energy consumption by putting the radios of the robots to sleep mode during each period T and waking them up in the beginning of the next period. In the following, we evaluate the impact of the beacon period T and the number of robots equipped with localization devices on localization accuracy and energy efficiency.

Impact of Beacon Period In this section, we evaluate the impact of the beacon period parameter T and discuss how this should be chosen to provide a good tradeoff between localization accuracy and energy consumption. The evaluation results are shown in Figures 6(a), 6(b).

Figure 6(a) shows the localization error over time, when 50% of the robots are equipped with localization devices, for four different beacon periods: 10, 50, 100 and 300 sec. In general, a small T improves the localization accuracy over time, since it gives the robots the chance to update frequently their position estimates. However, we surprisingly observe that for very small T (T = 10 sec) the average error over time becomes worse (it is about 7m for T = 10 sec, 5m for T = 50 sec, and 6.6m for T = 100 sec). This shows that the position updates due to the localization algorithm are not always accurate. Some "bad" beacons, e.g., beacons received from long distances, may worsen the performance of the algorithm. For this reason, the algorithm should not be executed arbitrarily frequently, but only when the error due to odometry has become quite large.

On the other hand, a very large value for T reduces energy consumption, since the radios are in sleep mode most of the time, and there is a limited number of transmissions/receptions which consume much energy. This is shown in Figure 6(b). In this figure we also show the energy consumption for CoCoA without coordination, i.e., when the radios do not go to the sleep mode, but they remain idle when they do not send or receive beacons. We observe that in that case the energy consumption is 2.6 to 8 times larger than the case with coordination, depending on the beacon period T.

From Figure 6(a) we observe that the localization error does not change a lot when T changes from 50 sec to 100 sec. Also from Figure 6(b) we observe that the energy consumption changes very slowly for T larger than 50 sec. Hence, the values between 50 and 100 sec are the best for the beacon period T, offering both high accuracy and low energy consumption.

Impact of Number of Localization Devices In this section, we evaluate the performance of CoCoA with different fractions of nodes equipped with localization devices. Reducing the number of such nodes reduces costs but can lead to inaccuracies in localization due to fewer beacons in the system. A small number of robots equipped with localization devices may affect the localization accuracy in two ways. First, some robots may not be able to receive any beacons in a beacon period, and they will have to localize themselves using only odometry. Second, a robot may receive very few beacons, all from long distances or from one side of the robot, in which cases the localization algorithm will have poor accuracy.

The evaluation results are shown in Figure 7. For this experiment, we used a total of 50 robots, and we varied the number of robots with localization devices from 5 to 35. Figure 7 shows that the average error over time does not increase a lot, when we reduce the robots with localization devices from 35 (5.2m) to 25 (5.9m). Thus, by installing localization devices on only half (25) of the robots, we can achieve an average localization error over time less than 6m and a maximum error less than 10m, by significantly reducing the cost. Depending on the application, we can reduce the cost even more, by using only 15 robots with localization devices, which gives an average error over time of about 8m and maximum error less than 12m.

4. Related Work

Dead reckoning (odometry) is the simplest approach for localization when a robot has some initial estimate about its position (i.e., pose maintenance). However, the odometry errors are added to the initial position estimate, and the total error is accumulated over time. Hence, for long

Proceedings of the 26th IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW'06)
0-7695-2541-5 /06 $20.00 © 2006 IEEE

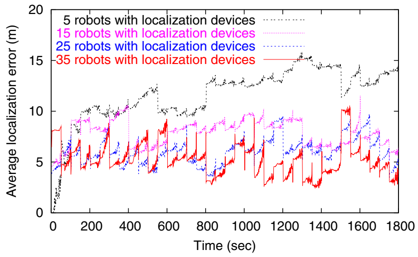

Figure 7. Impact of number of robots with localization devices on CoCoA localization error.

term localization, the robot has to periodically update its position using external references, e.g. GPS-equipped landmarks [12, 18, 2, 8, 14, 7, 10, 5, 19, 11].

Multi-robot localization, where robots cooperate to localize each other, has also been studied during the last 10 years [12, 6, 16, 9, 15]. The common feature in all these approaches is that they require all the robots to be equipped with expensive devices (cameras, laser range-finders, etc.), in order to sense each other. Most of them also require the existence of external landmarks.

In contrast to these previous works, CoCoA is a general purpose architecture which allows mobile robots to cooperate and coordinate in localizing each other, without the need of external landmarks, and without the need of expensive localization devices installed on all robots. Moreover, CoCoA is not tied to a specific localization technique.

The only other work that uses cheap wireless devices for localization is [13]. However, it is proposed for indoor environments, where the distribution of the signal strength vs. distance is not Gaussian. Hence an accurate signal strength map is required and static external landmarks are used, as opposed to CoCoA which is proposed for use in infrastructure-less scenarios.

5. Conclusion and Future Work

We proposed CoCoA, an architecture for low-cost, quickly-deployable, and energy-efficient localization for mobile robot teams in infrastructure-less dynamic environments. CoCoA equips a few robots with localization devices so that they help to localize other robots while performing their own tasks. CoCoA also coordinates the localization process of robots for energy-efficiency. CoCoA provides reasonably accurate locations; the average localization error is about 8m when only one third of the robots are equipped with localization devices. This can be useful for many applications, such as search and rescue operations.

There are several avenues for further investigation. We are interested in determining how transmission power control can be used to increase the distance that nodes in the CoCoA architecture can cooperate. It is also interesting to investigate the noise distributions of RF beacons when operating over special hardware that supports power control. While our simulation model is based on data from real wireless signal measurement tests and hence is expected to accurately model the real environment, we plan to validate our findings using a real testbed.

References

[1] S. M. Das, Y. C. Hu, C. Lee, and Y.-H. Lu. An effi cient group communication protocol for mobile robots. In *Proc. of IEEE ICRA*, 2005.

[2] A. Davison and N. Kita. 3d simultaneous localisation and map-building using active vision for a robot moving on undulating terrain. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, December 2001.

[3] DEAR: Distributed Energy-Effi cient Autonomous Robots Project. Home page `http://www.engineering.purdue.edu/ResearchGroups/DEAR`.

[4] L. M. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *Proc. of IEEE INFOCOM*, April 2001.

[5] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte carlo localization: Effi cient position estimation for mobile robots. In *Proc. of AAAI*, July 1999.

[6] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8, 2000.

[7] D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artifi cial Intelligence Research*, 11:391–427, 1999.

[8] J. Guivant and E. Nebot. Optimization of the simultaneous localization and map building algorithm for real time implementation. *Journal of Robotic Reasearch*, 17(10):565–583, 2000.

[9] A. Howard, M. Mataric, and G. Sukhatme. Localization for mobile robot teams using maximum likelihood estimation. In *Proc. of the IEEE International Conference on Intelligent Robots and Systems*, 2002.

[10] K. Konolige and K. Chou. Markov localization using correlation. In *Proc. of IJCAI*, 1999.

[11] B. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal on Robotics and Automatic Systems*, 8:47–63, 1991.

[12] R. Kurazume, S. Nagata, and S. Hirose. Cooperative positioning with multiple robots. In *Proc. of the IEEE International Conference in Robotics and Automation*, 1994.

[13] A. Ladd, K. Bekris, A. Rudys, L. Kavraki, and D. Wallach. On the feasibility of using wireless ethernet for indoor localization. *IEEE Transactions on Robotics and Automation*, 20, 2004.

[14] J. F. Leonard and H. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions Robotic and Automation*, 7(3):376–382, 1991.

[15] M. D. Marco, A. Garulli, A. Giannitrapani, and A. Vicino. Simultaneous localization and map building for a team of cooperating robots: a set membership approach. *IEEE Transactions on Robotics and Automation*, 19, 2003.

[16] S. Roumeliotis and G. Bekey. Distributed multirobot localization. *IEEE Transactions on Robotics and Automation*, 18, 2002.

[17] M. L. Sichitiu and V. Ramadurai. Localization of Wireless Sensor Networks with a Mobile Beacon. In *Proc. of MASS*, September 2004.

[18] R. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *Journal of Robotic Reasearch*, 5(4):56–68, Winter 1987.

[19] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust monte carlo localization for mobile robots. *Artifi cial Intelligence Journal*, 2001.

[20] X. Zeng, R. Bagrodia, and M. Gerla. Glomosim: A library for parallel simulation of large-scale wireless networks. In *Proc. of PADS Workshop*, May 1998.

[21] Activmedia robotics. http://www.activrobots.com/.

Proceedings of the 26th IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW'06)
0-7695-2541-5 /06 $20.00 © 2006 IEEE