In *the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*
pp. 2764-2769, EPFL, Switzerland, Semptember 30 - October 4, 2002

A Region-based Approach for Cooperative Multi-Target Tracking in a Structured Environment*

Boyoon Jung and Gaurav S. Sukhatme

University of Southern California, Los Angeles, USA, boyoon|gaurav@robotics.usc.edu

Abstract

This paper addresses the problem of tracking multiple targets using a network of communicating robots and stationary sensors. We introduce a region-based approach which controls robot deployment at two levels. A coarse deployment controller distributes robots across regions using a topological map and density estimates, and a target-following controller attempts to maximize the number of tracked targets within a region. A behavior-based system is presented implementing the region-based approach. Intensive simulations were performed to investigate the correlation between our approach and the degree of occlusion in the environment. The region-based approach shows better performance than a 'naive' local-following strategy when the environment has significant occlusion. We performed real-robot experiments to validate the system. These experiments open up a new line of research, which suggests that an optimal ratio of robots to stationary sensors may exist for a given environment with certain occlusion characteristics.

1 Introduction

Although the target tracking problem has been investigated thoroughly [1, 2, 3, 4, 5, 6], most research in this area has focused on the development of a single accurate tracker. However, in security or surveillance applications, it is often more important to carefully position multiple sensors in order to cover a large structured environment. In this context, mobile robot-based trackers are attractive; they can potentially reduce the overall number of sensors needed in the tracking network, and they can also adapt to the movement of targets or the dynamic changes in an environment by re-positioning themselves in response.

We introduce a scalable, cooperative tracking system that consists of multiple mobile robots and multiple stationary sensors. Prior research [7, 8] on cooperative tracking problem has focused on allocating multiple robots to multiple targets in a bounded environment with no obstacles. These systems were not applied or scaled for dynamic, structured environments, e.g. office environments which consist of many small rooms. We focus on a structured, complex environment instead of a large single space. The key problem is to develop an online, coordinated motion strategy for robot positioning that leads to a high degree of coverage of the environment. Our approach [9] to the problem is to divide the environment into topologically simple constituents called regions. Robot deployment is controlled at two levels; a coarse level of control causes robots to distribute themselves across regions depending on the estimated number of targets in a region relative to the size of a region. Within a region, robots employ other strategies to look for and follow targets. For the coarse level of control, even though sensors interchange their tracking information, each sensor is still independent in the sense that each sensor maintains its topological map and other information independently.

We performed experiment with three different environments which have different degrees of occlusion, and compared performance with a 'naive', non-cooperative algorithm which does not utilize the topological map. Our approach shows better performance when the environment has significant occlusion. We also demonstrate the modularity of our algorithm by integrating stationary sensors into the tracking system without any modification.

2 Related Work

Previous research has focused on building single-robot, reliable trackers. In [1], a single robot tracked multiple targets using a particle filter and Joint Probabilistic Data Association Filters. [2] demonstrated a strategy that maximizes the probability of future visibility in a structured environment. In [3] a real-time visual tracking method was proposed and implemented on an outdoor robot CAIR-2. [4, 5] used a laser rangefinder and [6] used an omnidirectional stereo system to track multiple moving targets.

There is relatively little research on cooperative multi-target tracking using multiple robots. [7] utilized the ALLIANCE architecture to achieve target-assignment; it used implicit communication (*acquiescence* and *impa-*

*This work is supported in part by DARPA grant DABT63-99-1-0015 and NSF grants IIS-0133947, ANI-0082498, and EIA-0121141.

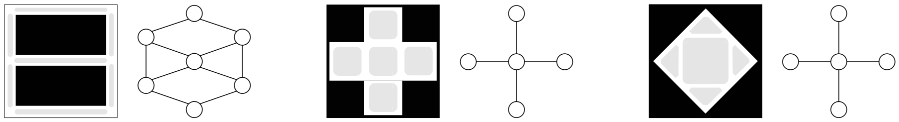

(a) Corridors (b) Open Space (c) Empty Space

Figure 1: *The simulation environments (the black regions are occupied and the rest is free space). The associated topological maps are shown as graphs*

tience levels) for cooperation. In contrast [8] used explicit communication (inhibition signal over the network) for target allocation. [10] described a Variable Structure Interacting Multiple Model (VS-IMM) estimator combined with an assignment algorithm for tracking multiple ground targets using multiple stationary sensors.

As different approaches to the same problem [11, 12, 13] have studied the target tracking problem in the context of distributed immobile sensor networks. [14, 15] introduced the Pursuit-Evasion problem, and analyzed the bounds on the number of necessary pursuers algorithmically.

3 The Region-based Approach

Our region-based approach is based on two fundamental assumptions: (1) *the environment can be divided into topologically simple regions using certain landmarks as demarcaters.* Figure 1 shows environments used in our experiments and the associated topological maps shown as graphs. In each graph a node corresponds to a region, and a link corresponds to a landmark. In our implementation the topology (represented as a planar graph) is made available *a priori* to each robot (2) *For two comparably sized regions, more robots should be deployed in the one with the higher number of targets.*

Given a topological map, every robot independently maintains two density estimates for each node (region). The *Robot Density* (Equation 1) is an estimate of the density of robots in a region, and the *Target Density* (Equation 2) is an estimate of the number of targets in a region. Both estimates are normalized by the sensor coverage area of a single robot.

$$D_r(R) = \frac{The\ number\ of\ robots\ in\ region\ R}{Area\ of\ region\ R\ /\ Unit\ coverage} \quad (1)$$

$$D_t(R) = \frac{The\ number\ of\ targets\ in\ region\ R}{Area\ of\ region\ R\ /\ Unit\ coverage} \quad (2)$$

The behavior of the target density estimate $D_t(R)$ is somewhat more complex than the robot density estimate $D_r(R)$. The value of $D_t(R)$ is normally calculated onboard each robot using Equation 2 but is set to -1.0 when the number of targets in region R is 0 and the number of robots in region R is nonzero i.e. $D_r(R) > 0$. This explicitly encodes the case when the region is marked as empty ($D_r(R) > 0$, $D_t(R) = -1$) as opposed to the case where the region is unobserved ($D_r(R) = 0$, $D_t(R) = 0$). Further, when no robots visit region R for a while, $D_t(R)$ is increased slowly over time until it becomes 0, which means that robots forget over time that a region R was explicitly labeled as being empty.

Each robot broadcasts its current position and the tracked targets' positions, and maintains internal estimates of $D_t(R)$ and $D_r(R)$ based on the broadcast packets it receives and its own sensor readings. Due to packet loss the estimates may not accurately measure the actual values. Further, the target density values will necessarily be inaccurate because the robots can only count tracked targets. Lastly, both sets of estimates can (and do) vary from robot to robot.

The Coarse Deployment Strategy

Each robot attempts to track as many targets as possible within its current region, and continually checks if it is available to migrate to a region which more urgently needs robots. When a robot is tracking T targets in the current region R_c, the robot is said to be available if the inequality in Equation 3 is satisfied. (θ_t is a threshold parameter.)

$$\frac{D_t(R_c)}{D_r(R_c)} \times T < \theta_t \quad (3)$$

Once a robot becomes available, it searches for the most urgent region based on its internal map. Given a distance d from its current position to a region r and an average distance d_{avg} between adjacent regions, the urgency value $U(R_i)$ of a region R_i is calculated using Equation 4. A robot selects the largest $U(R_i)$ value, and if the value is greater than a threshold θ_u, it navigates to region R_i; otherwise, it stays within its current region R_c.

$$U(R_i) = uf(D_t(R_i), D_r(R_i)) \times \frac{d_{avg}}{d} \quad (4)$$

where $\alpha = \frac{Area(R_i)}{Unit\ coverage} \times \theta_d$, and

$$uf(D_r, D_t) = \begin{cases} \frac{D_t}{D_r} & D_r \neq 0 \\ D_t \times \alpha & D_r = 0\ \&\ D_t \neq 0 \\ 1.0 & D_r = D_t = 0 \end{cases} .$$

θ_d is a parameter that controls how far a robot willing to traverse; for example, by setting θ_d to a small number, a robot is more likely to choose a closer region with less targets over a father region with more targets, and vice versa.

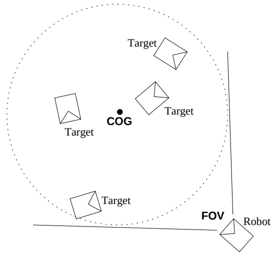

Figure 2: *Following targets within a region*

Target Tracking within a Region

Each robot tries to maximize the number of tracked targets within a region. In order to track multiple targets using a camera with limited field of view (FOV), a robot should be positioned close enough not to lose any target, and far enough to keep all targets within its FOV. In order to approach to this proper position, each robot calculates the center of gravity (COG) of the current tracked targets as shown in Figure 2. The robot attempts to keep at a distance $\frac{d_{max}}{sin(FOV/2)}$ from the COG where d_{max} is the distance between the COG and the target furthest from the COG in the group.

4 System Architecture

Figure 3 shows a behavior-based control architecture for the mobile robots, which implements the deployment strategies described above. Thin arrows indicate information flow (a module can share its internal data with other modules), and thick arrows indicates parameter modification (one module can change other modules' behavior). The controller consists of three layers: *Motor Actuation Layer*, *Target Tracking Layer*, and a *Monitoring Layer*.

4.1 The Motor Actuation Layer

The *Motor Actuation Layer* controls robot deployment, for example, it causes a robot to approach targets, or traverse from one region to another. There are four modules in this layer.

The *RobotMove* module is a repository of basis behaviors [16], and always perform one of these behaviors as selected by the *RobotDispersion* module. The basis behaviors are *RandomWandering*, *WallFollowing*, *TargetFollowing*, *SpotApproaching*, and *TurningAround*. *TargetFollowing* behavior implements the algorithm for target tracking within a region described in Figure 2.

The *RobotDispersion* module disperses robots to regions. It checks a robot's availability by accessing the internal map (*Map*), and drives the robot to the most urgent region by following a path generated by *Map* module.

The *MotorControl* module generates actual motor control signals (Speed & Turnrate) based on input signals from other modules. It has four different modes: *Sum*, *Min*, *Max*, and *One*.

The *ObstacleAvoidance* module is for safe navigation. It acts differently when a robot is tracking targets and when it is not. It executes only speed control when a robot is tracking targets, and leaves turnrate control to *RobotMove* (more specifically *TargetFollowing*). However, it regulates both speed and turnrate control when a robot is not tracking a target.

4.2 The Target Tracking Layer

The modules in the *Target Tracking Layer* detect targets, estimate their positions, and interchange this information among robots so that each robot can compute the density estimates defined earlier.

The *RobotLocalization* module keeps track of a robot's position. Since a robot calculates the positions of tracked targets based on its local coordinate system, knowing the robot's position is important for accurate target position estimation. This module estimates the robot's position based on odometry all the time, and compensates drift error whenever it sees laser beacons [17] which are disseminated in the environment at region boundaries.

The *ColorBlobTracker* module detects targets of certain colors, and calculates their positions. It uses a camera to detect targets and to calculate bearing to them, and uses a laser rangefinder to measure the distance to the targets. More details about the tracker can be found in [18]. The positions of targets in the robot's local coordinates are saved for the *TargetFollowing* behavior, and the positions in global coordinates are broadcast over the wireless network.

The *MapUpdate* module receives the broadcast packets, and updates density information in the *Map* module. The *Map* module is a passive data storage; all its contents are updated by the *MapUpdate* module. However, it has a functionality of a path planner; it can generate topological paths from one region to a specified goal region when there is a request from other modules.

4.3 Monitoring Layer

The *Monitoring Layer* is not directly related to the controller, but it provides a convenient means to observe the internal status of the controller during operation. Currently, this layer is used for debugging only, but it can be used for inter-robot communication for future extension.

5 Experiments

To evaluate our approach, we performed experiments with ActivMedia Pioneer DX-2 robots and the *Player/Stage*[19] software platform.[1] Pioneer robots

[1] *Player* is a server and protocol that connects robots, sensors and control programs across the network. *Stage* simulates a population of Player devices, allowing off-line development of control algorithms. *Player* and *Stage* were developed jointly at the USC Robotics Research Labs and HRL Labs and are freely available under the GNU Public License from *http://playerstage.sourceforge.net*.

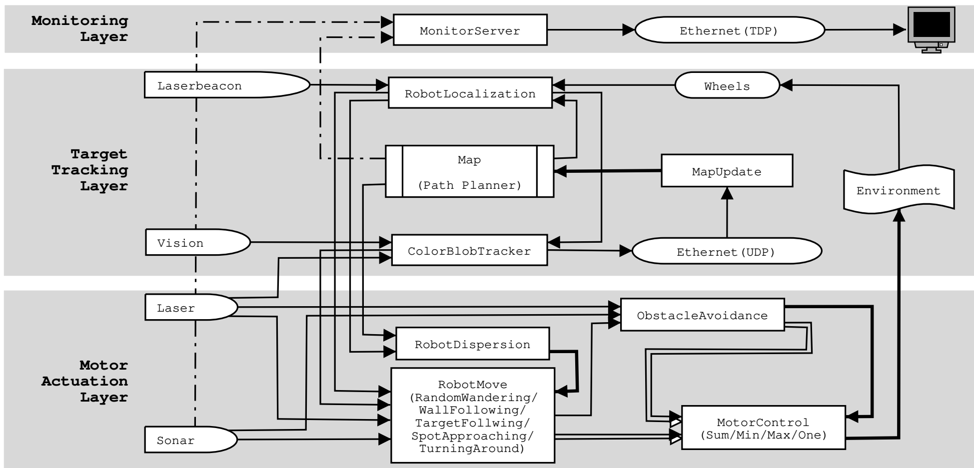

Figure 3: *Behavior-based robot control architecture*

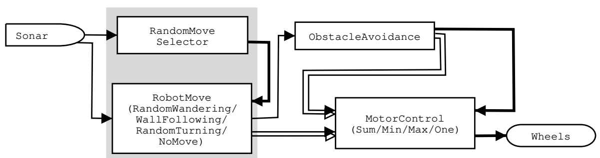

Figure 4: *System architecture for targets*

equipped with a SICK laser rangefinder and a Sony PTZ camera for tracking, and each target was a Pioneer robot carrying a bright-colored cylinder (see Figure 8).

Tracking performance was evaluated using Equation 5. Following [7] we define the Observation as

$$Observation = \sum_{t=0}^{T} \frac{m}{M} \times \frac{1}{T} \times 100 \qquad (5)$$

where M is the total number of targets, m is the number of targets being tracked at time t and the experiment runs over time T.

5.1 Target Modeling

Pioneer robots were used as moving targets. Each target robot uses sonar sensors for obstacle avoidance, and carries a bright-colored cylinder that is easily detected by a vision system. As shown in Figure 4, targets have four different behaviors: *Random-Wandering*, *Wall-Following*, *Random-Turning*, and *No-Move*. Each behavior is chosen randomly with probabilities (0.3, 0.3, 0.2, 0.2) in order. This target model showed sufficiently unpredictable, complex movement during experiments.

5.2 Simulation

The simulations were performed with various configurations in different environments. After having observed several simulations, we realized that the characteristics of the environment affect the system performance. The shape of the environment, in particular how obstructed it is, seems to be significant. In order to investigate this correlation between our region-based approach and characteristics of environment, we chose three distinct environments shown in Figure 1. The first environment, *Corridors*, consists of long, narrow regions which cause many occlusions. Each region is narrow enough that a single target may cause serious occlusion to the trackers. The second environment, *OpenSpace*, consists of several relatively big regions with few occlusions. The last environment, *EmptySpace*, is a single empty space which has no occlusion at all. Occlusions can be caused only by targets or other robots. The areas of these three environments are equal.

Two different strategies were compared: The *Region-based Strategy* and a *Local-following Strategy*. The *Region-based Strategy* uses the system described in Section 4, and the *Local-following Strategy* uses the same architecture without the *RobotDispersion* module.

We fixed the number of robots at 2, and changed the number of targets from 4 to 20 in steps of 2. Each configuration ran for 10 minutes a total of 9 times, and the average performance was taken as the final result (Figure 5).

In the *Corridors* environment, the *Region-based Strategy* showed better performance, especially when the number of targets was large. Since the regions are narrow and many occlusions are caused, it is almost impossible for a single robot to track more than two targets at a time, which means cooperation among robots is indispensable. On the other hand, the performance of the *Region-based Strategy*, in the *EmptySpace* environment, was not as good as that of *Local-following Strategy*. This can be explained by the fact that a robot is able to track targets in other regions without traversing to those regions in the empty environment because there is zero occlusion. This benefit is maximized for the *Local-following Strategy*, but not for the *Region-based Strategy*. Both strate-

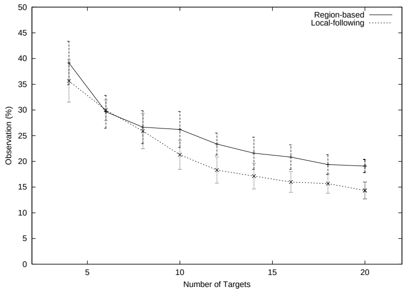

(a) Corridors

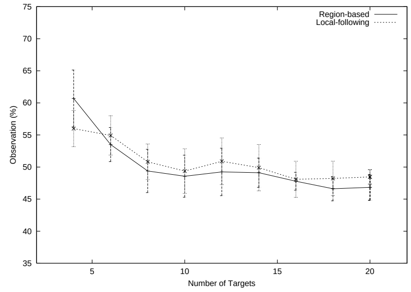

(b) Open Space

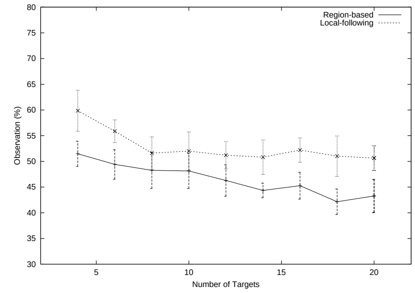

(c) Empty Space

Figure 5: *Simulation results comparing the performance of the two strategies*

Table 1: *Significance values from T-test*

# targets	Corridors	OpenSpace	EmptySpace
4	0.112957	0.022510	0.000117
6	0.808498	0.334473	0.000120
8	0.660541	0.369265	0.062597
10	0.007183	0.638818	0.046621
12	0.000549	0.378994	0.003367
14	0.006264	0.608437	0.000124
16	0.000445	0.786089	0.000042
18	0.001258	0.184465	0.000059
20	0.000008	0.062607	0.000093

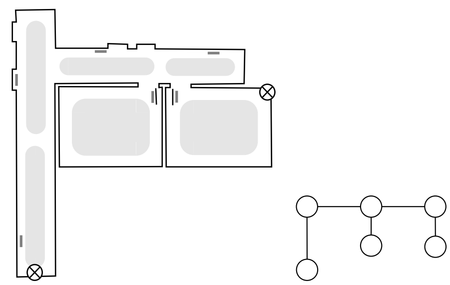

Figure 6: *Environment for real-robot experiments*

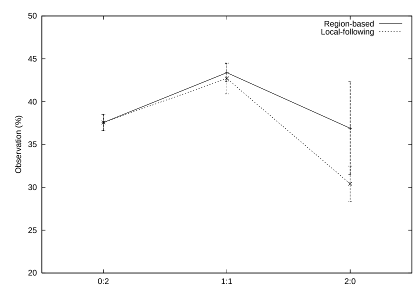

Figure 7: *Performance of the real-robot system*

gies showed almost the same performance in *OpenSpace*, whose shape causes a small amount of occlusion.

We performed a t-test with each data points pair in order to see if the results from the two strategies show significant difference statistically. Table 1 shows the significance values of the t-test; a small number implies that the two distributions are statistically different. As expected, the significance values in *OpenSpace* were big, and those in *Corridors* and *EmptySpace* were small.

5.3 Real-Robot Experiments

Another interesting issue in multi-target tracking using multiple sensors is whether adding a mobile sensor (e.g. a mobile robot) is always more helpful to improve tracking performance compared to adding a stationary environment-embedded sensor (e.g. a security camera). Each sensor has its own advantages. For example, a mobile sensor can cover wider area over time, and can adapt to targets' movement patterns. On the other hand, a fixed sensor can be installed at the best position depending on the environments, and it causes less interference.

Since we wanted a realistic environment for this test, the experiments were performed using real robots on the second floor of our Computer Science Building. Figure 6 shows the map of the floor and its topological map; it consists of two offices and two long corridors. The figure also shows the positions of the laser beacons (dark gray, thick lines) and stationary sensors ('x' marks).

We fixed the number of targets at 4, and used a total of 2 sensors in three different combinations: with only two stationary sensors, with one stationary sensor and one mobile robot, and with two mobile robots. Each configuration was tested over 3 trials, and each trial ran for 10 minutes. The average performance is shown in Figure 7. When the robots were used, we performed the experiment twice, using both the *Region-based Strategy* and the *Local-following Strategy*.

In case of using two stationary sensors, the performance depends on targets' movement. As shown in Figure 7, the standard deviation was small ($\pm 2\%$), which signals that the targets spread out evenly over the environment. When a mobile robot was used together with an embedded sensor, the overall performance was higher than any other cases. As shown in Figure 8 (a), the robot was able to position itself properly to track multiple targets,

(a) 1 robot tracking 2 targets

(b) 2 robots tracking 1 target

Figure 8: *Tracking examples*

and also able to follow the targets continuously. However, when only two robots were used, the performance was worse than the previous cases, especially with the *Local-following strategy*. Since two robots moved independently to follow targets, they often ended with following the same target, as shown in Figure 8 (b). That worst case was observed less often when the system used the *Region-based Strategy*.

Although the case of using a mobile robot and an embedded sensor together showed the best performance, more research on this topic will be necessary to come to a firm conclusion.

6 Conclusion and Future Work

In this paper, we presented a multi-robot behavior-based system that is able to track multiple targets effectively. In contrast to prior research on cooperative tracking, our research focuses on allocating multiple robots to multiple targets in a bounded, structured, complex environment, like an office environment, instead of a large single space. We have developed an online, coordinated motion strategy for robot positioning that leads to a high degree of coverage of the environment. Our approach to the problem is to divide the environment into topologically simple constituents called regions. Robot deployment is controlled at two levels; a coarse level of control causes robots to distribute themselves across regions depending on the estimated number of targets in a region relative to the size of a region. Within a region, robots employ other strategies to look for and follow targets. For the coarse level of control, sensors interchange their tracking information; however, each sensor is independent in the sense that each sensor maintains its topological map and other information independently.

Through intensive simulations and real-robot experiments, our region-based approach was shown to perform better than a naive approach when the environment contains significant occlusion. This opens up a significant line of inquiry which could further elucidate the exact relationship between robot coordination strategies and environment shape. We also demonstrated the modularity of our algorithm by integrating stationary sensors into the tracking system without any modification. This experiment opens up a new line of research, which suggests that an optimal ratio of mobile robots to stationary sensors may exist for a given environment with certain occlusion characteristics.

References

[1] D. Schultz, W. Burgard, D. Fox, and A. B. Cremers, “Tracking multiple moving targets with a mobile robot using particle filters and statistical data association,” in *Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, 2001, pp. 1165–1170.

[2] S. M. LaValle, H. H. Conzalez-Banos, C. Becker, and J. Latombe, “Motion strategies for maintaining visibility of a moving target,” in *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, 1997.

[3] J. Chung and H. S. Yang, “Fast and effective multiple moving targets tracking method for mobile robots,” in *Proceedings of the 1995 IEEE International Conference on Robotics and Automation*, 1995, pp. 2645–2650.

[4] B. Kluge, C. Kohler, and E. Prassler, “Fast and robust tracking of multiple moving objects with a laser range finder,” in *Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, 2001, pp. 1683–1688.

[5] M. Lindstrom and J. O. Eklundh, “Detecting and tracking moving objects from a mobile platform using a laser range scanner,” in *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, Hawaii, October 2001, pp. 1364–1369.

[6] H. Koyasu, J. Miura, and Y. Shirai, “Realtime omnidirectional stereo for obstacle detection and tracking in dynamic environments,” in *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, Hawaii, October 2001, pp. 31–36.

[7] L. E. Parker, “Cooperative robotics for multi-target observation,” *Intelligent Automation and Soft Computing, special issue on Robotics Research at Oak Ridge National Laboratory*, vol. 5, no. 1, pp. 5–19, 1999.

[8] B. B. Werger and M. J. Mataric, “Broadcast of local eligibility for multi-target observation,” in *Proceedings of Distributed Autonomous Robotic Systems*, 2000.

[9] B. Jung and G. S. Sukhatme, “Cooperative tracking using mobile robots and environment-embedded, networked sensors,” in *Proceedings of the 2001 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 2001, pp. 206–211.

[10] Y. Bar-Shalom and W. D. Blair, Eds., *Multitarget-Multisensor Tracking: Applications and Advances*, vol. 3, Artech House, 2000.

[11] B. Horling, R. Vincent, R. Mailler, J. Shen, R. Becker, K. Rawlins, and V. Lesser, “Distributed sensor network for real time tracking,” in *Proceedings of the 5th International Conference on Autonomous Agents*, 2001, pp. 417–424.

[12] L. J. Guibas, “Sensing, tracking, and reasoning with relations,” *IEEE Signal Processing Magazine*, March 2002.

[13] F. Zhao, J. Shin, and J. Reich, “Information-driven dynamic sensor collaboration for tracking applications,” *IEEE Signal Processing Magazine*, March 2002.

[14] M. Yamashita, H. Umemoto, I. Suzuki, and T. Kameda, “Searching for mobile intruders in a polygonal region by a group of mobile searchers,” in *Symposium on Computational Geometry*, 1997, pp. 448–450.

[15] L. J. Guibas, J. Latombe, S. M. LaValle, D. Lin, and R. Motwani, “A visibility-based pursuit-evasion problem,” *International Journal of computational eometry and Applications*, vol. 9, no. 5, pp. 471–494, October 1997.

[16] M. J. Mataric, “Behavior-based control: Examples from navigation, learning, and group behavior,” *Journal of Experimental and Theoretical Artificial Intelligence, special issue on Software Architectures for Physical Agents*, vol. 9, no. 2-3, pp. 67–83, 1997.

[17] A. Howard, M. J. Mataric, and G. S. Sukhatme, “Relaxation on a mesh: a formalism for generalized localization,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Wailea, Hawaii, October 2001.

[18] B. Jung and G. S. Sukhatme, “Tracking multiple moving targets using a camera and laser rangefinder,” Institute for Robotics and Intelligent Systems Technical Report IRIS-01-397, University of Southern California, 2001.

[19] B. Gerkey, R. Vaughan, K. Stoy, A. Howard, G. S. Sukhatme, and Maja J Mataric, “Most valuable player: A robot device server for distributed control,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Wailea, Hawaii, October 2001, pp. 1226–1231.