

Cooperative Tracking with Mobile Robots and Networked Embedded Sensors*

Boyoon Jung and Gaurav S. Sukhatme

boyoon|gaurav@robotics.usc.edu

Robotic Embedded Systems Laboratory
Robotics Research Laboratory
Department of Computer Science
University of Southern California
Los Angeles, CA 90089-0781

Abstract

We study the target tracking problem using multiple, environment-embedded, stationary sensors and mobile robots. The stationary sensors and robots maintain region-based density estimates which are used to guide the robots to parts of the environment where unobserved targets may be present. Experiments in simulation show that the region-based approach works better than a 'naive' target following approach when the number of targets in the environment is high. We present real-robot experiments which support the results from the simulation study.

1 Introduction

Autonomous target tracking has several potential applications; e.g. surveillance, security, etc. Mobile robot-based trackers are attractive for two reasons: they can potentially reduce the overall number of sensors needed and they can adapt to the movement of the targets (e.g. follow targets to occluded areas). We focus on the robot-based target tracking problem (CMOMMT: Cooperative Multirobot Observation of Multiple Moving Targets [1, 2]) which has received recent attention in the robotics community.

The CMOMMT problem is defined as follows. Given a bounded, enclosed region S , a team of m robots R , a set of n targets $O(t)$, and a binary variable $In(o_j(t), S)$ defined to be true when target $o_j(t)$ is located within region S at time t , and $m \times n$ matrix $A(x)$ is defined where

$$a_{ij}(t) = \begin{cases} 1 & \text{if a robot } r_i \text{ is monitoring target } o_j(t) \\ & \text{in } S \text{ at time } t \\ 0 & \text{otherwise} \end{cases}$$

and the logical OR operator is defined as

$$\bigvee_{i=1}^k h_i = \begin{cases} 1 & \text{if there exists an } i \text{ such that } h_i = 1 \\ 0 & \text{otherwise} \end{cases}$$

*This work is sponsored in part by DARPA grant DABT63-99-1-0015 and NSF grants ANI-9979457 and ANI-0082498

The goal is to maximize the observation.

$$\text{Observation} = \frac{\sum_{t=0}^T \sum_{j=1}^m \bigvee_{i=1}^k a_{ij}(t)}{t \times m} \quad (1)$$

In [1, 2], the ALLIANCE architecture was used to coordinate robots in the CMOMMT task; role assignment among mobile robots was achieved implicitly through one-way communication. However, it was assumed that the observation sensors had a perfect field-of-view and a known global coordinate system. Experiments were performed in a bounded, enclosed spatial region, and an indoor global positioning system was utilized as a substitute for vision or range-sensor-based tracking. In [3], an approach to a similar problem using the BLE (Broadcast of Local Eligibility) technique was presented which used a real video camera to track moving objects, and one-way communication for explicit role-assignment. The environment in [3] was simple, and target movements were pre-programmed. Each target was also identified a priori.

In this paper, we consider a more realistic office-like environment composed of corridors. The major difference from previous research is to utilize environment-embedded, stationary sensors installed at fixed positions in the environment. This is particularly apt given the proliferation of small, low-power sensors and communication hardware in today's built environments. These sensors are used to track moving targets within their range, and broadcast target location information over a wireless channel. The mobile robots are used to explore regions which are not covered by the fixed sensors. The robots also broadcast tracked target location information over the wireless channel. We present a region-based strategy for robot coordination and task allocation which uses a topological map, and compare it to a naive 'target-following' strategy using an observation metric similar to Equation 1. Our results in simulation and with real robots

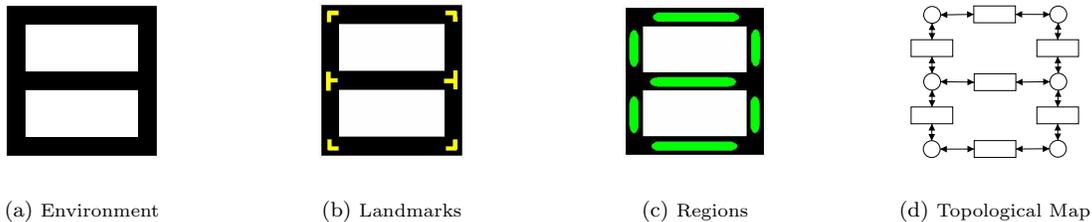


Figure 1: A structured environment segmented into landmarks and regions

and sensors show that the region-based strategy works better than the naive strategy when the number of targets is large.

2 Region-based Robot Coordination

When the environment is an empty open space, the challenge is to assign targets to a fixed number of robots based on the distances between robots and targets. However, when the environment has structure (e.g. an office-type environment), it is important to disperse robots properly. We propose a region-based approach for this purpose.

2.1 ASSUMPTIONS

We make several assumptions about the environment and robot capabilities. First, a topological map of the environment is assumed to be given. This is reasonable given that previous research on map building is extensive [4, 5, 6]. In this paper, the data structures from [5] have been adopted to build a topological map.

The second assumption is that global communication between robots and the fixed sensors is allowed. However, this does not imply two-way communication, like a negotiation. We only use one-way broadcast among sensors and robots; whenever a sensor detects a moving target, the sensor broadcasts the estimated position of the target. Perfect communication is not necessary either; a small rate of packet loss will not degrade the performance of the system.

Third, the initial position of the mobile robots is assumed to be known for localization. However, localization information is used only for inferring the positions of moving objects, not for robot navigation. Navigation is based on a landmark detector, not global positioning.

2.2 THE REGION-BASED METHOD

The basic idea of the region-based approach is that the environment can be divided into several (topologically simple) regions using landmarks as demarcators. In Figure 1 (a), a simple office-type environment that consists of corridors is shown, (b) shows landmarks, (c) shows how the environment can be divided into regions by the landmarks, and (d) is a topological map

of the environment. Given the sensors on our robots, intersections are a natural landmark choice.

Assuming that a topological map is given, we need to decide which region “needs” more robots and which region does not. In order to answer this question, each region is assigned two properties: a *robot density* (D_r) and a *target density* (D_t). They are defined as follows:

$$D_r(r) = \frac{\text{The number of robots in region } r}{\text{Area of region } r} \quad (2)$$

$$D_t(r) = \frac{\text{The number of targets in region } r}{\text{Area of region } r} \quad (3)$$

Robot density indicates how many robots¹ are in a region, and target density indicates how many tracked targets are in a region. Both values are normalized by area. If a region has low robot density and high target density, the region needs more mobile robots, and vice versa.

Sometimes, a robot must stay in its current region even though there is another region that needs more robots; for example, when it is the only robot tracking objects in its region or when there are too many moving objects in the region. Therefore, each robot must check its *availability* on the basis of the following criteria:

$$D_t(r_c) < 0 \quad (4)$$

$$\frac{D_t(r_c)}{D_r(r_c)} < \theta \quad (5)$$

Equation (4) models the situation when the robot has observed the current region r_c , but couldn’t find any target in it, and equation (5) models the situation when there are more than enough robots in the current region r_c . If the situation falls under one of the above criteria, the robot is available and decides to move to another region.

Another problem is how to choose the most urgent region to be observed. The two density properties of each region are used to make this decision. The following equations show how these properties are used:

$$D_r(r_i) = 0 \wedge D_t(r_i) > 0 \quad (6)$$

¹Environment-embedded stationary sensors are counted as robots when robot density is calculated.

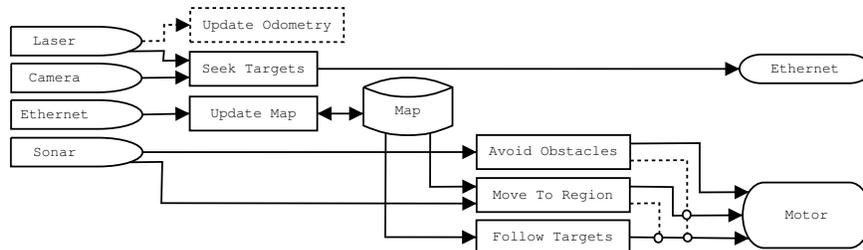


Figure 2: System architecture for mobile robots and embedded sensors

$$\frac{D_t(r_i)}{D_r(r_i)} \geq 1.0 \quad (7)$$

$$D_r(r_i) = 0 \wedge D_t(r_i) = 0 \quad (8)$$

Equation (6) means that a region r_i has moving objects which are not being observed. Equation (7) means that a region has too many objects to be tracked by the current number of robots, and Equation (8) means that a region is not being observed currently. These rules are prioritized; Equation (6) has the highest priority, and Equation (8) has the lowest one. A region for which a higher priority rule is applicable must be observed first. If there are two or more regions with the same score, the region closest to the current robot position is selected to be observed.

2.3 SYSTEM ARCHITECTURE

Figure 2 shows a behavior-based control architecture for the mobile robots which uses the density estimate for role assignment. There are six modules in the controller: one for detecting moving targets and five for dispersing robots according to the criteria discussed in the previous section. The embedded sensors have exactly the same system architecture as the mobile robots, but only one module, *Seek-Targets*, is activated.

2.3.1 *Seek-Targets*

Seek-Targets detects moving objects and broadcasts their estimated positions. As shown in Figure 2, two trackers have been developed: a laser-based tracker and a vision-based tracker. Target tracking is a well studied problem, especially in computer vision [7, 8, 9]. Our trackers are simple by design since our focus is on robot role-assignment.

The laser-based tracker uses the SICK laser rangefinder. It reads the laser rangefinder at 10 Hz and analyzes the data to find moving objects using scan differencing between consecutive laser readings. A big difference is attributed to a moving object. For accurate tracking, a simple edge detection algorithm is used.

The vision-based tracker uses a camera and a laser rangefinder. A color-blob detector was used to sim-

plify the vision problem. It finds the existence and direction of colored objects using a camera, and measures the distance to objects using a laser rangefinder. For details on the implementation of the trackers, the reader is referred to [10].

When moving objects are detected, the *Seek-Targets* behavior broadcasts their estimated positions over the network.

2.3.2 *Update-Map*

Update-Map maintains an internal map. It reads broadcast packets about target locations, and puts them in a queue. By counting the packets in the queue, it can estimate the number of robots and the number of targets in each region. However, before counting them, a proper grouping strategy is required. Consider a situation in which a stationary sensor and a robot both detect a moving target. The mobile robot would broadcast the position of the target, and the embedded sensor would do the same. However, these position estimates would be different due to uncertainty in the system, and they must be grouped as one target. In addition, the embedded sensor would recognize the robot as a moving object because it cannot distinguish a robot from moving objects. This estimated position must be grouped with the robot's position, and removed from the target list.

The robot density and the target density of each region are updated using Equations (2) and (3). The range of robot density is from 0.0 to 1.0, and the range of target density is from -1.0 to 1.0. The expression in Equation (3) for target density, does not use the range -1.0 to 0. This range is used by *Update-Map* in order to mark empty regions. Whenever a robot cannot find any moving objects, it sets the target density of the current region to -1.0, which means that the region does not have any moving objects. By using the negative range, a robot can distinguish a region that does not have any moving object from a region that has not been observed. If target density is negative, *Update-Map* increases it slowly over time to 0.0 because the environment is dynamic. When target density becomes 0.0, it means the system has forgot-

ten that there was no target in the region; the robots may now try to observe the region again if needed.

2.3.3 Avoid-Obstacles

Avoid-Obstacles allows a robot to navigate without collision. It uses the eight front sonars to detect an obstacle. Each sonar uses a different range to detect obstacles, and constructs a virtual oval-shaped region in front of the robot. When any obstacle enters the region, *Avoid-Obstacles* reduces the speed in inverse-proportion to the distance to the obstacle, and turns away from the obstacle. In addition, *Avoid-Obstacles* stops a robot in place when a moving object approaches it, instead of actively avoiding it.

2.3.4 Move-To-Region

Move-To-Region disperses robots all over the environment. The algorithm for it is divided into three steps: checking robot availability, finding the most urgent region, and moving to the region. First, this behavior checks if a robot itself is free to move to another region. Equations (4) and (5) are the criteria to decide if a robot is available for observing other regions. If available, the behavior finds a region to be observed urgently on the basis of the internal map. It simply examines the internal map, and finds one using the prioritized scoring policy (Equations (6), (7), and (8)). If there are two or more regions that have the same score, the closer region is selected as the most urgent region. Once a starting region and a goal region are decided, a simple graph search is performed to find the shortest path. (The internal maps consist of nodes (landmarks) and regions as shown in Figure 1(d).) A robot follows the shortest path to move to the goal region.

2.3.5 Follow-Targets

The *Follow-Targets* behavior causes robots to follow detected targets. In order to make robots follow more than one target at the same time, *Follow-Targets* calculates the center of mass of detected targets and follows this point, not the targets themselves. The worst case is when two targets move in opposite directions. This does not happen often in our narrow corridor environment.

2.3.6 Update-Odometry

This behavior has been added for real-robot experiments. In the real-robot experiments, odometry information is used to estimate robots' position in a global coordinate system, but odometry always drift. A laser beacon is positioned in each corner, and the *Update-Odometry* compensates the drift in real time whenever it sees the laser beacon.

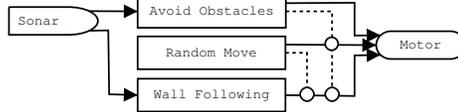


Figure 3: System architecture for target simulation

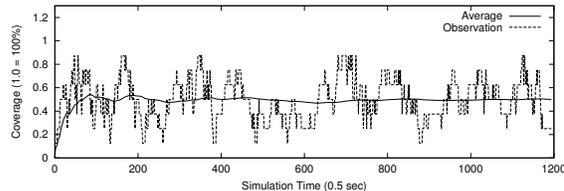


Figure 4: Convergence of the average value

3 Experimental Results

3.1 SIMULATION

To test our region-based cooperative target tracking approach, several experiments have been performed using a multiple robot simulator (Player/Stage) and the structured environment shown in Figure 1 (a). Player [11] is a server and protocol that connects robots, sensors and control programs across the network. Stage [12] simulates a population of Player devices, allowing off-line development of control algorithms. Player and Stage were developed at the USC Robotics Research Labs and are freely available under the GNU Public License from <http://robotics.usc.edu/player/>.

3.1.1 Target Simulation

Because Stage supports only mobile robots, moving targets in the environment were simulated using robots. The target movements are intended to crudely simulate human movements in an office environment, especially in corridors, like wall-following, turning, staying in place with other targets, etc. Figure 3 shows the control architecture of moving targets.

Wall-Following uses two pairs of side sonars. Target motion is divided into two parts: speed control and direction control. *Wall-Following* sets the speed to a maximum value, and uses a proportional controller to align the target parallel to a wall using the front and rear sonars. *Random* was added to make targets' movements somewhat unpredictable.

Avoid-Obstacles is the same module used in the robot controller. The only difference is that a target never stops in place; it always actively avoids obstacles.

3.1.2 Performance Evaluation

The simulation experiments were done with various configurations in order to evaluate the region-

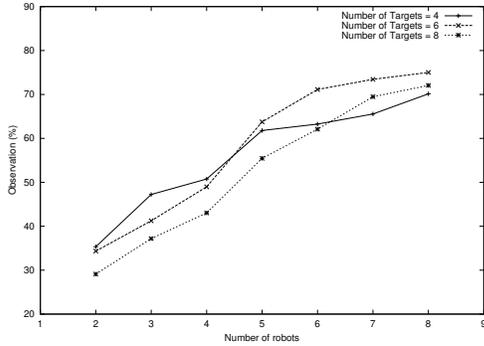


Figure 5: Performance of the region-based method

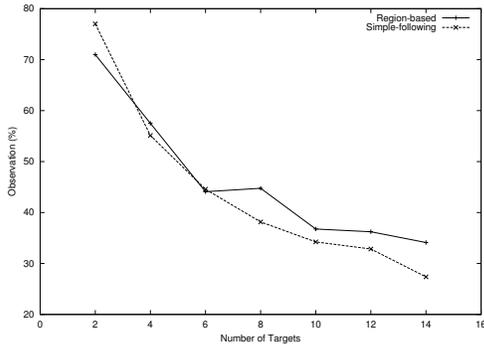


Figure 6: Comparison to a simple following method. Four mobile robots are being used for tracking.

based approach. The metric in Equation 1 is used to evaluate performance. Each trial ran for 10 minutes. Figure 4 shows the average observation rate over time which stabilizes after 6–7 minutes. The difference between the actual position of a target and its position as reported by the sensors was small. The average error was approximately 4 cm.

The performance of the system varies according to the number of sensors, and the number of moving objects. In our experiments, a total of 21 different configurations were tested. We changed the number of sensors from 2 to 8, and the number of targets from 4 to 8 in steps of 2. Figure 5 shows the tracking results. As expected, the more the sensors, the better the tracking performance. One interesting fact observed through the experiments is that the performance improves whenever sensors are added, but this improvement tails off when the number of sensors is greater than the number of objects.

The region-based method was compared to a simple target-following method. In order to implement the simple method, we inhibited the *Move-To-Region* module. The robots follow walls, but after finding moving targets, the robots follow their center of mass. We changed the number of moving targets from 2 to 14 in steps of 2, and four mobile robots were used

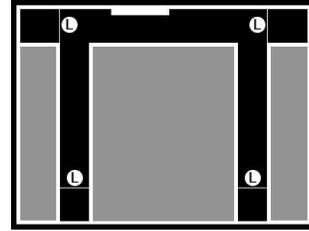


Figure 7: Computer Science Building

for all cases. Figure 6 shows the results. When the number of objects is small, the simple method occasionally showed better performance because robots do not give up following objects to explore other regions that may be more urgent. However, as the number of targets is increased, the region-based method showed better performance because *Move-To-Region* causes robots to move to regions that have more objects.

3.2 REAL-ROBOT EXPERIMENTS

In order to verify the results from simulation, we have performed real-robot experiments. Because Player provides exactly the same interface for a real Pioneer robot as a virtual robot in Stage, the control programs written for simulation can be used for real robot experiments without major modification.

3.2.1 Environment

The region-based approach was implemented on ActivMedia Pioneer DX-2 robots with SICK laser rangefinders and Sony PTZ cameras. They have exactly the same configuration as the virtual robots in Stage. The experiment was performed on the second floor of the Computer Science Building. As shown in Figure 7, there are four landmarks and three regions. The horizontal corridor is approximately 11 m long, and the vertical ones are approximately 9 m long. The width of the corridors is 1.5 m.

One serious difference from simulation was imperfect odometry information. In order to compensate for odometry drift, laser beacons were installed in each corner (marked ‘L’ in Figure 7). The positions of the laser beacons in a global coordinate system are known, and robots update their internal odometry information whenever they see a laser beacon.

Bright colored balls were used as moving targets. In the experiment, each trial ran for 9 minutes, and each target ball was rolled in a region for 3 minutes in total. Each ball was allowed to stay within a region up to 1 minute, after which it was moved to another region which was selected randomly. This 9-minute long experiment was done three times for each configuration, and the average of the three results was taken as the final result for the configuration.

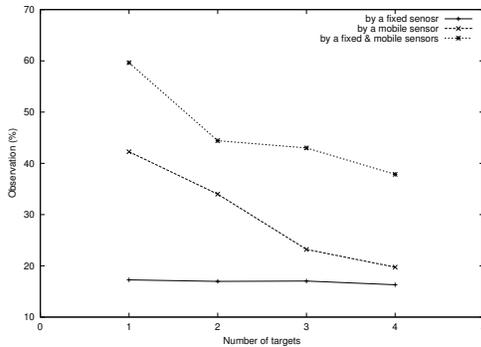


Figure 8: Result of the real-robot experiment

3.2.2 Performance Evaluation

Three different configurations were tested in the real-robot experiment: one fixed sensor, one mobile robot, and one fixed sensor and one mobile sensor. The results are shown in Figure 8. Because the environment is dynamic, i.e. the targets move, one mobile robot showed better performance than one fixed sensor. The overall performance (as measured by the observation) of the real-robot system is worse than that of the system in simulation. There are two possible explanations. First, the global positioning of each robot was worse. In simulation, the *Following-Targets* behavior uses global position of targets to follow them. Since this is not possible in the physical experiments, a robot loses targets more often in reality than in simulation. Second, the effective range of the real vision system was shorter than the simulated vision system. In spite of these differences, the basic trend in the data in Figure 8 is similar to the data observed in simulation.

4 Conclusion and Future Work

Autonomous target tracking systems have many real-world applications. We have presented a region-based tracking system, which is especially well suited to structured environments. The system utilizes embedded sensors, like surveillance cameras already installed in buildings. Initial experiments indicate that our approach shows better performance compared to a naive strategy when there are many targets to be tracked.

Our initial experiments reported here are encouraging. Some related questions we plan to investigate in the future include:

1. What is the optimal ratio of mobile robots to stationary embedded sensors?
2. Can the assumption about the shared global coordinate system be relaxed?
3. Where should the fixed sensors be placed?

4. How can robots modify their exploration strategy to improve system performance?

References

- [1] Lynne E. Parker, “Cooperative motion control for multi-target observation,” in *Proceedings of the 1997 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1997, pp. 1591–1598.
- [2] Lynne 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.
- [3] Barry B. Werger and Maja J. Mataric, “Broadcast of local eligibility for multi-target observation,” in *Proceedings of Distributed Autonomous Robotic Systems*, 2000.
- [4] Goksel Dedeoglu, Maja J. Mataric, and Gaurav S. Sukhatme, “Incremental, on-line topological map building with a mobile robot,” in *Proceedings of Mobile Robots*, Boston, MA, 1999, vol. XIV, pp. 129–139.
- [5] Goksel Dedeoglu and Gaurav S. Sukhatme, “Landmark-based matching algorithm for cooperative mapping by autonomous robots,” in *Distributed Autonomous Robotic Systems (DARS)*, Knoxville, Tennessee, 2000.
- [6] Sebastian Thrun, Wolfram Burgard, and Dieter Fox, “A probabilistic approach to concurrent mapping and localization for mobile robots,” *Machine Learning and Autonomous Robots (joint issue)*, vol. 31 & 5, pp. 29–53 & 253–271, 1998.
- [7] Isaac Cohen and Gerard Medioni, “Detecting and tracking objects in video surveillance,” in *Proceeding of the IEEE Computer Vision and Pattern Recognition 99*, Fort Collins, June 1999.
- [8] Stephen S. Intille, James W. Davis, and Aaron F. Bobick, “Real-time closed-world tracking,” in *Proceeding of the IEEE Conference on Computer Vision and Pattern Recognition*, June 1997, pp. 928–934.
- [9] Alan J. Lipton, Hironobu Fujiyoshi, and Raju S. Patil, “Moving target classification and tracking from real-time video,” in *Proceeding of the IEEE Workshop on Applications of Computer Vision*, 1998.
- [10] Boyoon Jung and Gaurav 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.
- [11] Brian Gerkey, Kasper Stoy, and Richard T. Vaughan, “Player robot server,” Institute for Robotics and Intelligent Systems Technical Report IRIS-00-391, University of Southern California, 2000.
- [12] Richard T. Vaughan, “Stage: A multiple robot simulator,” Institute for Robotics and Intelligent Systems Technical Report IRIS-00-393, University of Southern California, 2000.