

Cooperative Tracking using Mobile Robots and Environment-Embedded, Networked 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. An architecture for robot motion coordination is presented which exploits a shared topological map of the environment. 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.

1 Introduction

Autonomous target tracking has many 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). The robot-based target tracking problem (CMOMMT: Cooperative Multirobot Observation of Multiple Moving Targets), has been formally defined in [1] and has received recent attention in the robotics community[2, 3].

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

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

The goal of the CMOMMT is to maximize the observation.

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

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. However, the environment in [3] was very simple, and movements of targets were pre-programmed. Each target was also identified a priori. In [4], a Variable Structure Interacting Multiple Model (VS-IMM) estimator combined with an assignment algorithm for tracking multiple ground targets was described.

In this paper, we consider a more realistic office-like environment. It consists of corridors; offices will be added in the near future. The major difference from previous research is to utilize environment-embedded, stationary sensors installed at fixed positions in the environment. These sensors are used to track moving targets in their sensor 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 the tracked target location information. We

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

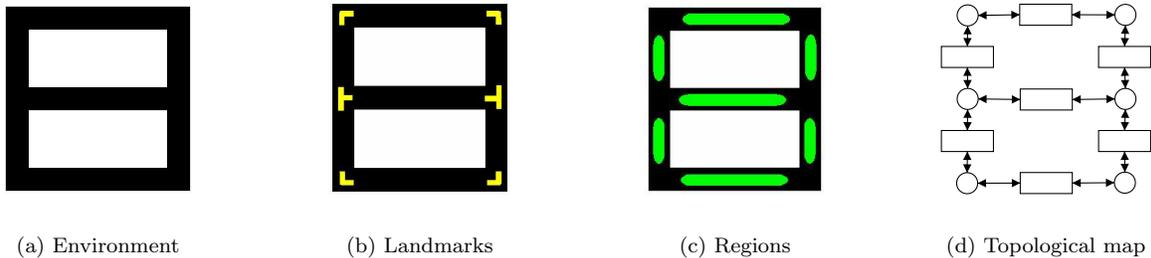


Figure 1: A structured environment segmented into landmarks and regions

present a region-based strategy for robot coordination which uses a topological map, and compare it to a naive ‘target-following’ strategy using the observation metric similar to Equation 1. Our results show that the region-based strategy works better than the naive strategy when the number of targets is large.

In Section 2, the region-based method and system architecture is described. The simulation environment and experimental results are discussed in Section 3. Concluding remarks and future work is discussed in Section 4.

2 Region-based Robot Coordination

When the environment is an empty open space, the main 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. 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. Previous research on map building is extensive [5, 6, 7]. In [5], a simple, modular, and scalable behavior-based technique for incremental on-line mapping is presented, and in [6], a simple, yet robust cooperative mapping method using multiple robots is presented. In [7], a probabilistic approach to building large-scale maps of indoor environments with mobile robots is presented. In this paper, the data structures from [6] 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. Since odometry is used for localization in the experiments reported here, the initial positions of the mobile robots must be known. However, localization information is used only for estimating the positions of moving objects, not for 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 demarcaters. 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.

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. The definition for $D_t(r)$ is not complete as stated in Equation (3); as we explain later, $D_t(r)$ can also assume negative values if no moving objects are detected in region r . 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

¹Embedded stationary sensors are counted as robots when robot density is calculated.

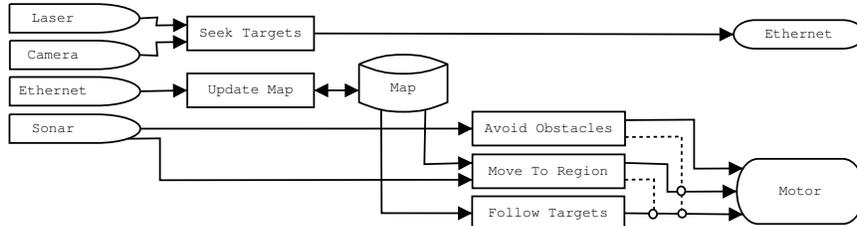


Figure 2: System architecture for mobile robots and embedded sensors

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 . This is signified by the ratio in Equation (5) being less than a prespecified threshold θ . 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 choose one. The following equations show how these properties are used:

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

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

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

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 five modules in the controller: one for detecting moving targets and four 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 [8, 9, 10]. 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. Figure 3 (a) and (b) shows two example actual laser readings. The upper window shows two consecutive laser readings and edges, and the lower window shows the difference between the two readings and a detected moving object.

The idea can be implemented without any limitation for stationary embedded sensors, but several limitations exist for mobile robots carrying laser rangefinders. In the mobile robots' case, simply comparing two consecutive laser readings is not correct because the robot actually moves during the scan process. Figure 4 shows two different positions of a robot when the laser was used. In order to compare these scans correctly, the old reading must be transformed to the new coordinate system. However, during the transformation, there may be several parts of the scan that have no valid data because of rounding errors or two exceptions. The first exception is when the old reading contains the maximum value, which means there is an empty region in front of it. The second exception is a corner occlusion. When there is a corner, the old scan does not have any information behind a corner, but the new scan may have (the fan-shaped region in Figure 4). Therefore, these areas must be ignored. The lower window in Figure 3 (b) shows these ignored regions; only the gray region is compared to

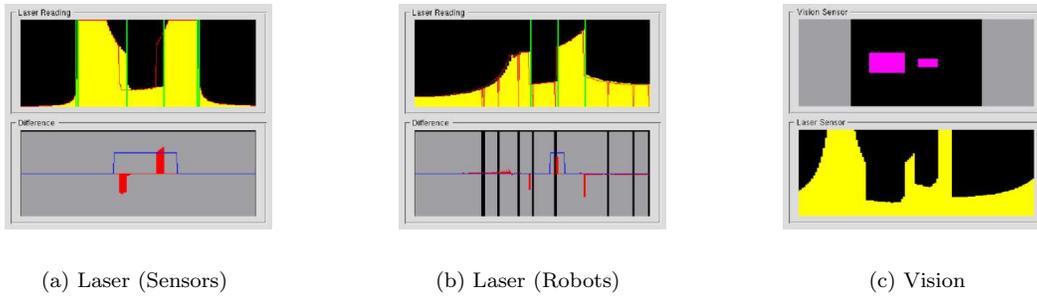


Figure 3: Moving-object tracker: various sensory readings

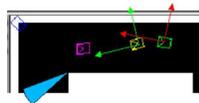


Figure 4: Coordinate transformation

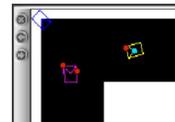


Figure 5: Grouping

calculate a difference. For example, in Figure 3 (a), there is no ignored region because the embedded sensors never move, but there are several ignored regions in Figure 3 (b).

The vision-based tracker uses a camera and a laser rangefinder. A color-blob detector was used to simplify 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. Figure 3 (c) shows corresponding camera and laser readings taken from a single robot.

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. Figure 5 shows a situation that requires grouping; A stationary sensor and a robot both detect a moving target. The mobile robot broadcasts the position of the target, and the embedded sensor does the same, but the position estimates are different. These two estimated positions of the moving object 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. According to Equation (3), target density cannot be negative. *Update-Map* uses the negative range of target density 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 forgotten 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

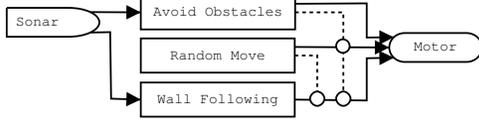


Figure 6: System architecture for target simulation

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.

3 Experimental Results

To test our region-based cooperative target tracking approach, several experiments have been performed using a multiple robot simulator.

3.1 Stage and Player

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/>.

3.2 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 6 shows the control architecture of moving targets.

Wall-Following uses two pairs of side sonars. Target motion is divided into two parts: speed control

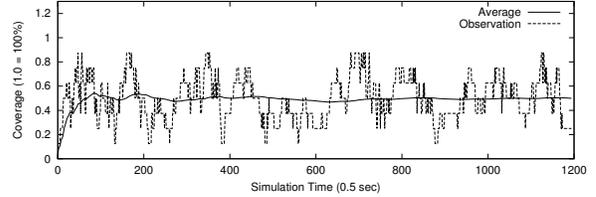


Figure 7: Convergence of the average value

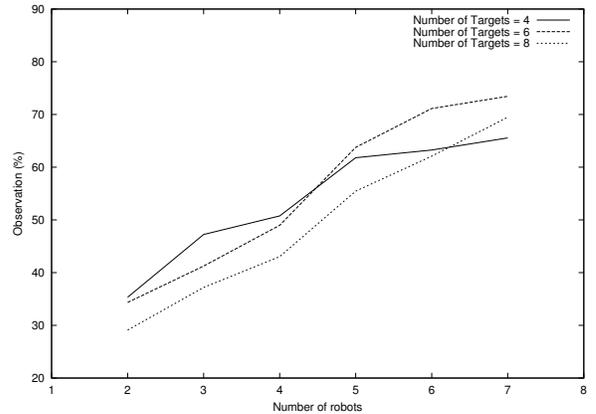


Figure 8: Performance of the region-based method

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. Currently, only one random move is being generated, turning around. However, due to interactions with other targets and robots, each target's moves are quite complicated and 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.

3.3 Experimental Results

The simulation experiments were done with various configurations in order to evaluate the region-based approach. A performance metric for the CMOMMT task was proposed in [1, 2], and the metric (Equation 1) is used to evaluate performance. Each trial ran for 10 minutes. Figure 7 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.

3.3.1 Performance Evaluation

The performance of the system varies according to the number of sensors, and the number of moving objects. In our experiments, a total of 18 different configurations were tested. We changed the number

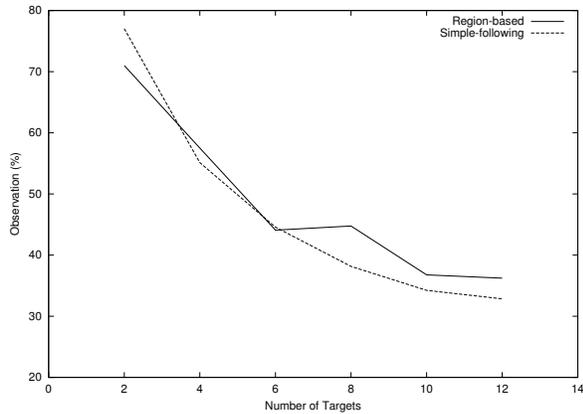


Figure 9: Comparison to a simple following method

of sensors from 2 to 6, and the number of targets from 4 to 8 in steps of 2. Figure 8 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.

3.3.2 Comparison to a simple strategy

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 12 in steps of 2, and four mobile robots were used for all cases. Figure 9 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.

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 when there are many targets to be tracked.

As future work, research on the ratio of mobile robots to embedded sensors is a topic we plan to address. In addition, real robot experiments are planned for the near future. 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.

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] Yaakov Bar-Shalom and William Dale Blair, Eds., *Multitarget-Multisensor Tracking: Applications and Advances*, vol. 3, Artech House, 2000.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [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.