

Tracking Targets using Multiple Robots: The Effect of Environment Occlusion

Boyoon Jung (boyoon@robotics.usc.edu) and Gaurav S. Sukhatme
(gaurav@robotics.usc.edu)

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

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 which maintains urgency estimates for each region, 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*, which is fully distributed and scalable. We compared the *Region-based Approach* to a 'naive' local-following strategy in three environments with varying degree of occlusion. The experimental results showed that the *Region-based Approach* performs better than the naive strategy when the environment has significant occlusion. Second, we performed experiments (the environment was held constant) in which two techniques for computing urgency estimates were compared. Last, different combinations of mobile sensors and stationary sensors were compared in a given environment.

Keywords: Multi-target tracking, mobile robotics, embedded sensors



1. Introduction

Autonomous target tracking using a network of sensors has many potential applications; for example, surveillance, security, etc. We are particularly interested in the case where the sensor network used for tracking is partially mobile i.e. some nodes of the network have the ability to re-position themselves in response to target motion. Such mobile robot-based trackers are attractive for two reasons: they can potentially reduce the overall number of sensors needed in the tracking network and they can re-configure in response to the movement of the targets.

Formally, we are interested in the problem of tracking multiple anonymous targets in a bounded planar environment using a network of communicating robots and stationary sensors. The key problem is to develop an on-line, 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. In previous work (Dedeoglu et al., 1999; Dedeoglu and Sukhatme, 2000) we have demonstrated an on-line, multi-robot, topological mapping technique which segments the environment into regions automatically. Other research on topological map building techniques includes (Kuipers and Byun, 1991) which describes how to construct a topological map directly from sensor measurements, and Voronoi Graph-based methods (Choset and Burdick, 1995a; Choset and Burdick, 1995b). In this research we focus on the high-level cooperative tracking problem assuming a topological map of the environment is given.

Our approach (Jung and Sukhatme, 2001a; Jung and Sukhatme, 2001b) controls robot deployment 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 or to maximize total coverage. The performance of this *Region-based Approach* depends on the characteristics of environment. Through intensive simulations and real-robot experiments, we show that our approach works better than a simple *Local-following Strategy* when the environment has a high degree of occlusion. Also, our experiments suggest that an optimal ratio of robots to stationary sensors may exist for a given environment with certain occlusion characteristics.

2. Related Work

Previous research treating multi-target tracking as a task allocation problem across multiple robots was reported in (Parker, 1999), where the ALLIANCE architecture was utilized to achieve target-assignment; if a target was not tracked for a while, the robot which was supposed to track the target would give up (*acquiescence*), and another robot would take up the target (*impatience*). *Impatience* and *acquiescence* were the two implicit communication methods between robots. In contrast (Werger and Matarić, 2000) used explicit communication; if a particular robot thinks it is best suited to track a specified target, it stops other robots from tracking the target by broadcasting inhibition signals over the network. (Gerkey and Matarić, 2001) showed that the target tracking problem can be solved using a principled publish/subscribe messaging model; the best capable robot is assigned to each tracking task using a one-round auction.

Considerable effort has been focused on building reliable trackers. In (Schultz et al., 2001), a single robot tracked multiple targets using a particle filter for object tracking and Joint Probabilistic Data Association Filters were used for measurement assignment. In (Chung and Yang, 1995) a real-time visual tracking method was proposed and implemented on an outdoor robot CAIR-2. In (Coue and Bessiere, 2001) a chasing robot using Bayesian Robot Programming and

visual servoing techniques was presented. (Koyasu et al., 2001) presented a stereo-visual tracking robot using two omni-directional camera aligned vertically. A laser rangefinder was used to track multiple moving targets in (Lindstrom and Eklundh, 2001; Kluge et al., 2001). (Bar-Shalom and Blair, 2000) described a Variable Structure Interacting Multiple Model (VS-IMM) estimator combined with an assignment algorithm for tracking multiple ground targets.

The target tracking problem has also been studied in the context of distributed immobile sensor networks. (Horling et al., 2001) introduced a distributed sensor network for moving target tracking. Their approach is to decompose an environment into a number of small sectors in order to reduce the interactions needed between sensors. Each sensor has only limited capability and changes its role dynamically, i.e. managing a sector, tracking a target, producing sensor data, or triangulating the position of a target. (Kumar et al., 2000) also describe a scalable sensor network for locating objects in a dynamic environment when small, networked devices are attached to the objects.

The Pursuit-Evasion problem introduced in (Yamashita et al., 1997) is a formally simplified tracking problem. The goal is to find continuously-moving intruders using a single or multiple searchers with flashlights that emit a single ray of light. (Yamashita et al., 1997) presents upper and lower bounds of the number of necessary searchers in a given environment (a simple polygon) and four measures of shape complexity of the environment (the number of edges, the number of reflex vertices, the bushiness, and the size of a minimum guard set). (Guibas et al., 1997) extend the problem to exploit a visible area instead of a single ray of light. Several bounds on the number of pursuers are defined and the complete algorithm for a single pursuer case is presented.

In this paper, we focus on the high-level cooperation among sensors rather than on tracker accuracy, so our tracker design is relatively simple compared to the trackers reported in (Chung and Yang, 1995; Bar-Shalom and Blair, 2000). However, we consider a more realistic office-like environment composed of corridors and offices instead of a big bounded, empty, open environment used in (Parker, 1999; Werger and Matarić, 2000) or a simple polygon used in (Yamashita et al., 1997). We do not make any assumption about sensors either; our tracker has a limited field-of-view and localization is achieved using odometry and a system of beacons, in contrast to a perfect field-of-view sensor and a known global coordinate system used in (Parker, 1999). Also, our system does not require target identification nor pre-programming of target movements as in (Werger and Matarić, 2000). We also differ from previous research in the utilization of 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.

3. The Region-based Approach

Our approach is based on two fundamental assumptions:

- *the environment can be divided into topologically simple convex regions using certain landmarks as demarcators.* Each region is assumed to be simply connected¹. Figure 1 shows the 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. For example, the link between two corridor regions may be a corner. In our implementation the topology (represented as a planar graph) is made available *a priori* to each robot.

¹ Two regions are connected with a single edge only when they are physically connected such that a robot can traverse from a region to another.




Figure 1. The simulation environments: The black regions are occupied and the rest is free space. The associated topological maps are shown as graphs. (a) *Corridors* consists of long, narrow regions which cause many occlusions. (b) *OpenSpace* consists of several relatively big regions with few occlusions. (c) *EmptySpace* is a single empty space which has no occlusion at all.

- 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 normalized by the sensor coverage area of a single robot.

$$D_r(R) = \frac{\text{the number of robots in region } R}{\text{area of region } R / \text{unit coverage}} \quad (1)$$

The *Target Density* (Equation 2) is an estimate of the number of targets in a region normalized by the sensor coverage area of a single robot.

$$D_t(R) = \begin{cases} -1 & D_t(R) = 0 \ \& \ D_r(R) \neq 0 \\ D_t(R) + \alpha & D_t(R) < 0 \ \& \ D_r(R) = 0 \\ \frac{\text{the number of targets in region } R}{\text{area of region } R / \text{unit coverage}} & \text{otherwise} \end{cases} \quad (2)$$

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 on-board each robot using the third formula of 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.

3.1. 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.

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

Equation 3 was constructed based on the following facts:

- the less targets a region has, the less robots the region requires.
- the more robots a region has, the more robots in that region are free to move to other regions.
- the fewer targets a robot is tracking at a given moment, the less the robot is required to stay in its current region.

θ_t is a parameter that decides how easily a robot would give up tracking the current targets and travel to other regions hoping to find a new, bigger group of targets.

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 $U(R)$ of a region R is calculated using Equation 4. A robot selects the largest $U(R)$ value, and if the value is greater than a threshold θ_u , it navigates using the topological map to region R ; otherwise, it stays within its current region R_c .

$$U(R) = uf \left(D_t(R), D_r(R) \right) \times \frac{d_{avg}}{d} \quad (4.1)$$

$$uf(D_r, D_t) = \begin{cases} \frac{D_t}{D_r} & D_r \neq 0 \\ \frac{D_t}{D_t} \times \alpha & D_r = 0 \ \& \ D_t \neq 0 \\ 1.0 & D_r = D_t = 0 \end{cases} \quad (4.2)$$

$$\alpha = \frac{\text{area of region } R}{\text{unit coverage}} \times \theta_d \quad (4.3)$$

Equation 4 is a formalization of the following ideas:

If there is a region that has targets but does not have a robot, or if there is a region that has more targets than robots, then some robot should move to the region. Otherwise, a robot should move to an unobserved region ($D_r = D_t = 0$) for exploration. θ_d is a parameter that controls how far a robot is willing to travel; 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 farther region with more targets, and vice versa.

3.2. 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 this 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 (Equation 5) from the COG where d_{max} is the distance between the COG and the target furthest from the COG in the group.

$$\text{distance} = \frac{d_{max}}{\sin(FOV/2)} \quad (5)$$




Figure 2. Following targets within a region: A robot maintains a proper distance from the center of gravity (COG) of the current tracked targets.




Figure 3. Coverage Computation: The coverage of a single robot is computed by integrating the distance readings, and the total coverage of all robots is the sum of each coverage. For more accurate computation, overlap is subtracted from the result.

3.3. EXTENSION: VISIBILITY MAXIMIZATION

The region-based approach is able to adopt other indicators as well. The approach described in the previous section might be improved by using visibility (LaValle et al., 1997) instead of the robot density. The region-based approach using the robot and target densities does not always maximize the total coverage in a region; for example, if the sensor ranges of robots in a region overlap, the total coverage within the region is often small even though there are enough robots in the region. In order to obtain a good robot spread within a region, a visibility maximization method is used. The visibility of a region R is defined as:

$$Visibility(R) = \frac{\text{the covered area of region } R}{\text{area of region } R} \quad (6)$$

Calculating an accurate covered area of each region requires a single, unified global coordinate system and geometrical information of all robots in it. This is unrealistic in dynamic environments and does not scale. In addition, every computation needs to be done in real-time. Therefore, the covered area should be estimated with given imperfect information. The system performance is expected to depend on the accuracy of the visibility estimation. In our system, the estimation process consists of two steps: *coverage assignment* and *overlap subtraction*.

Step 1: Coverage Allocation The coverage of a robot r is computed by integrating the distance function $l_r(\theta)$ of its laser rangefinder in the *FOV* of its camera (Figure 3 (a)).

$$coverage(r) = \int_0^{FOV} \frac{1}{2} l_r^2(\theta) d\theta \quad (7)$$




Figure 4. Visibility Maximization Method relying on Local Sensing: A robot decides its behavior based on a local situation; therefore, explicit communication or negotiation is not required.

Once the coverage of each robot has been computed, it is assigned to regions based on the current robot position. When a robot is positioned in the middle of multiple regions, the coverage is divided and assigned to those regions evenly.

Step 2: Overlap Subtraction For real-time computation, overlapping coverage is checked only between robots in the same region, and the coverage of each robot is approximated as a single sector with a radius of the average distance of laser readings (Figure 3 (b)). The overlapping coverage between two sectors is subtracted from the total coverage (Figure 3 (c)) since it has been counted twice.

For the coarse deployment strategy, the robot density $D_r(R)$ is simply replaced with $Visibility(R)$, and the rest of the strategy remains the same as described in Section 3.1. However, the movement within a region is modified in order to pursue two goals in parallel; each robot tries to maximize the number of targets tracked, and to maximize the coverage of regions.

Maximizing coverage implies reducing the overlap of robots' coverage, so robots should keep maximum distance from each other. We use the following method relying on only local sensing and computation.

1. If there are only targets in its sensing range as shown in Figure 4 (a), a robot tracks the group of targets as described in Section 3.2.
2. If there are only robots in its sensing range as shown Figure 4 (b), a robot turns in a direction opposite to the group of robots.
3. If there are both targets and robots in its sensing range as shown in Figure 4 (c), a robot tracks only the group of targets which are closer to itself, and assumes that the rest of them would be tracked by the other robots.

The advantage of this method is that it does not require explicit communication or negotiation among robots.

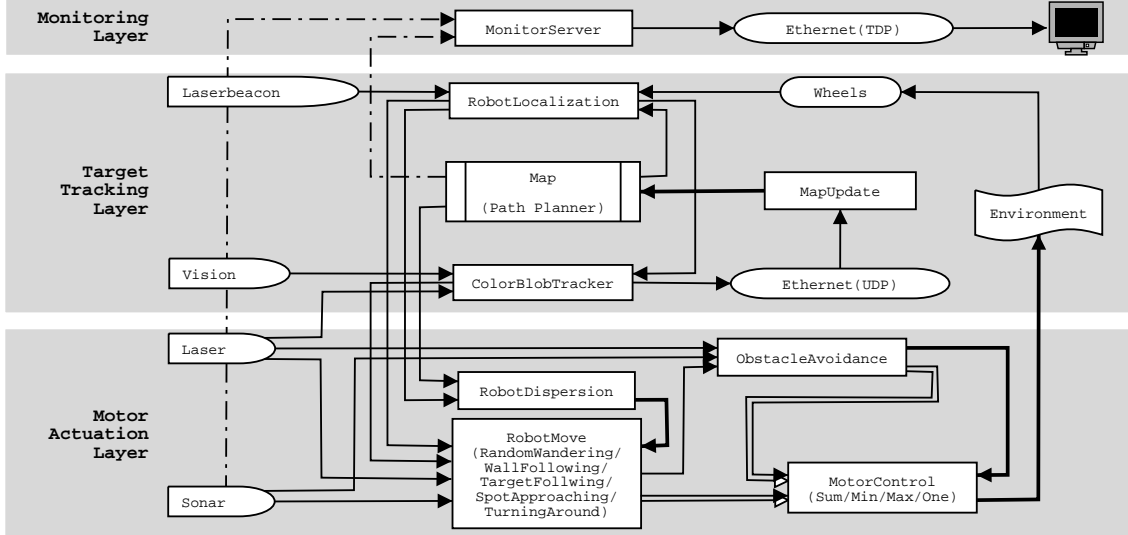


Figure 5. Behavior-based robot control architecture: Behaviors are drawn as rectangles, and sensors are shown as half ellipses (read only) or full ellipses (read/write). Thin arrows indicate information flow (a module can share its internal data with other modules), and thick arrows indicate parameter modification (one module can change other modules' behavior). Dashed lines show user-selected debugging information, and double lines emphasize that a behavior generates decoupled motor control commands (one for speed control and one for heading control).

4. System Architecture

Figure 5 shows a behavior-based control architecture for the mobile robots, which implements the deployment strategies described above. 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 (Matarić, 1997), and always performs one of these behaviors as selected by the *RobotDispersion* module. The basis behaviors are *RandomWandering*, *WallFollowing*, *TargetFollowing*, *SpotApproaching*, and *TurningAround*. Other than *TargetFollowing* which needs some explanation, the others are relatively standard and self-explanatory. *TargetFollowing* acts differently when a robot uses the robot density (Equation 1) and when it uses the visibility (Equation 6). In the former case, *TargetFollowing* causes a robot to approach the center of gravity (COG) of all tracked targets (Figure 2). It also controls the distance between the robot and the COG based on how much the targets are spread. However, in the latter case, *TargetFollowing* makes a robot approach the COG of only selected targets and keep away from other robots (Figure 4).

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*. All inputs are integrated in *Sum* mode, or a single, specified input is used in *One* mode; minimum/maximum inputs are selected in *Min/Max* mode.

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 beacons² which are disseminated in the environment at region boundaries. Given the global positions of the laser beacons and distance/bearing information to them, the *RobotLocalization* module can estimate the position and heading of a robot very accurately, and can correct odometry error.

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 (Jung and Sukhatme, 2001c). 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 at the speed of 10 Hz.

The *MapUpdate* module receives the broadcast packets, and updates density information in the *Map* module. Each target is identified by its global position; if the distance between two target positions is smaller than the twice length of a robot body, the two targets are assumed as a single target. 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. Since the bandwidth of the wireless network is limited, each robot sends its internal information only when there is a request from an external monitoring program. Currently, this layer is used for debugging only, but it can be used for inter-robot communication for future extension.

The *MonitorServer* module is a thread-based server program; it always listens to a specified port, and sends selected internal data of other modules when there is a request. The module can serve more than one client at a time.

² The laser beacon (Howard et al., 2001) is a simple bar-code constructed using strips of retro-reflective paper. Based on the reading from a SICK laser rangefinder, a robot can compute the ID, range, bearing, and orientation for each visible beacon.




Figure 6. Configurations for robots and targets: Pioneer robot is equipped with a SICK laser rangefinder and a Sony PTZ camera, and each target was a Pioneer robot carrying a bright-colored cylinder




Figure 7. System architecture for targets: A target shows four different behaviors (wander, follow walls, turn in place, or stop) in random combination.

5. Experimental Setup

To evaluate our approach, we performed experiments with ActivMedia Pioneer DX-2 robots and the *Player/Stage*³ software platform. *Player* (Gerkey et al., 2001) is a server and protocol that connects robots, sensors and control programs across the network. *Stage* (Vaughan, 2000) simulates a population of Player devices, allowing off-line development of control algorithms. Because *Player* provides the same interface for real Pioneer robots and for virtual robots in the *Stage* simulator, the same configuration was used for real-robot experiments and simulations. Pioneer robots equipped with a SICK laser rangefinder and a Sony PTZ camera were used for tracking, and each target was a Pioneer robot carrying a bright-colored cylinder (see Figure 6).

Tracking performance was evaluated using Equation 8. Following (Parker, 1997) we define the Observation as

$$Observation = \sum_{t=0}^T \frac{m}{M} \times \frac{1}{T} \times 100 \quad (8)$$

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 7, targets have four different behaviors: *Random-Wandering*, *Wall-Following*,

³ *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>.

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. ENVIRONMENT COMPLEXITY

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.

A formal approach to analyze the complexity of environment models has been introduced in (Cazals and Sbert, 1997). The basic idea of the approach is to generate uniformly distributed random rays, and to compute statistics such as how often rays intersect with models, how long truncated segments are, etc. Various characteristics about 3D models can be inferred from these statistics. We utilize this approach in order to measure the complexities of our planar environments.

As explained in (Cazals and Sbert, 1997), two end-points of rays must be selected from the *circumcircle* of environment images for uniform distribution, not from the bounding box. In order to meet this restriction, two angles (θ_1 , θ_2) were randomly generated, and the two end-points (x_1, y_1) and (x_2, y_2) were computed as follows when r is the radius of the circumcircle:

$$x_1 = r \times \cos(\theta_1), \quad y_1 = r \times \sin(\theta_1)$$

$$x_2 = r \times \cos(\theta_2), \quad y_2 = r \times \sin(\theta_2)$$

If a ray did not intersect with the bounding box of environment images at all, the ray did not count.

Table I shows the statistics of the three environments. The environments with several targets present were also tested in order to investigate how the complexities change when there are robots/targets in the environments. The targets were evenly distributed in regions. In Table I, p_0 is the probability that a ray does not intersect with any objects, i.e. walls, robots, and targets. en_{int}^G is the average of the number of intersection points, and σen_{int}^G is its standard deviation. μCL is the average of the length of free lines⁴, and σCL is its standard deviation. Especially, p_0 is a nice indicator showing how much occlusion an environment has, and μCL can be an indicator showing how big (wide) each region is.

As expected, *Corridors* showed small p_0 and μCL , which means it has significant occlusion and its regions are narrow. On the other hand, p_0 of *Empty Space* was big since it has little occlusion. The μCL values of both *Open Space* and *Empty Space* were big because their regions are wide. Needless to mention, the more targets, the higher complexity.

⁴ Free lines are the segments of rays that divided by intersection points. The lengths of the free lines are normalized with the longest free line.

Table I. Complexity of the environments as a function of number of targets: Only targets are included for the complexity analysis because robots are part of the tracking system, not part of the environment. However, since a robot has the same size with a target, the result can be interpreted such as the number of targets indicates the number of objects including robots.

Environment	# targets	# rays	p_0	en_{int}^G	σen_{int}^G	μCL	σCL
Corridors	0	10,000	0.3003	1.9552	1.5201	0.2714	0.2438
Corridors	2	10,000	0.2963	2.0394	1.6076	0.2657	0.2445
Corridors	4	10,000	0.2901	2.1212	1.6646	0.2596	0.2422
Corridors	6	10,000	0.2876	2.1998	1.7273	0.2555	0.2432
Corridors	8	10,000	0.2849	2.2784	1.8201	0.2509	0.2413
Open Space	0	10,000	0.7600	0.4800	0.8542	0.4678	0.2798
Open Space	2	10,000	0.7260	0.5738	0.9600	0.4478	0.2826
Open Space	4	10,000	0.6784	0.7140	1.0983	0.4265	0.2852
Open Space	6	10,000	0.6331	0.8188	1.1447	0.4131	0.2853
Open Space	8	10,000	0.6015	0.9082	1.1995	0.3992	0.2829
Empty Space	0	10,000	1.0000	0.0000	0.0000	0.5639	0.2462
Empty Space	2	10,000	0.9406	0.1196	0.4776	0.5301	0.2601
Empty Space	4	10,000	0.8800	0.2510	0.6950	0.5057	0.2687
Empty Space	6	10,000	0.8338	0.3678	0.8613	0.4785	0.2724
Empty Space	8	10,000	0.8003	0.4688	0.9985	0.4594	0.2730

5.3. CONTROL STRATEGIES

Three experiments were performed with various configurations. Two of them were performed using the *Stage* simulator, and the third experiment was performed using *Pioneer DX-2* robots.

5.3.1. *Region-based versus Local-following Strategy*

The *Region-based Strategy* was compared with 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. The *Local-following Strategy* still tries to maximize the number of tracked targets in the current region by maintaining a proper distance from a group of targets as explained in Figure 2. The best scenario for the *Local-following Strategy* would be that many targets gather in an open space and stay in a large, circle-shaped group; a single robot would be able to keep track of most of the targets in the group. The worst scenario for the *Local-following Strategy* would be that all robots track the same single target and the rest of targets stay away from the target being tracked; since robots never give up the current tracked target, they would never leave the current region to move to other regions which contain other targets. On the other hand, the *Region-based Strategy* causes a robot to relinquish tracking the current tracked target to explore other regions when there are more than enough robots in the current region; therefore, the worst scenario would never happen. The purpose of this comparison is to measure if the map-based coarse control causes an improvement in tracking performance.




Figure 8. Environment for real-robot experiments: It consists of two offices and two long corridors. The laser beacons are shown as dark gray, thick lines on the walls, and the position of stationary sensors are marked by \otimes symbols.

5.3.2. Robot Density versus Visibility

Two variants of the *Region-based Strategy* are compared. With the *Robot Density*, a robot tracks all targets in its sensing range. However, when robots get too close to each other, the overall performance might reduce. On the other hand, with *Visibility*, a robot covers a wider area within a region by keeping away from other robots, but this may result in the loss of some targets in the process of receding.

In order to perform experiments with various configurations and minimize the effect of unnecessary factors, the *Stage* simulator is used with the *Corridors* environment shown in Figure 1 (a). Thus this experiment does not vary the environment.

5.3.3. Mobile Robots versus Embedded Sensors

Another interesting issue in multi-target tracking using multiple sensors is whether adding a mobile sensor (i.e. a mobile robot) is always more helpful to improve tracking performance compared to adding a stationary environment-embedded sensor (i.e. 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.⁵


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 8 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 (\otimes marks). In this study the environment and control strategy were held constant and the ratio of the mobile to stationary trackers was varied.

6. Experimental Results


6.1. REGION-BASED VERSUS LOCAL-FOLLOWING STRATEGY

The simulations were performed with various configurations in three different environments. 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. The experimental results are shown in Figure 9.


⁵ Actual cost of mobile robots and fixed sensors is not considered for this comparison since we only focus on the performance of cooperative tracking strategies using heterogeneous sensors.



(a) Corridors



(b) Open Space



(c) Empty Space

Figure 9. Simulation results comparing the performance of the two strategies: The number of robots was fixed at 2, and the number of targets varied from 4 to 20 in steps of 2.




Figure 10. Performance with visibility maximization: Three strategies were compared (Region-based strategy with visibility maximization, Region-based strategy with robot density, and Local-following strategy) in the *Corridors* environment. The number of robots was fixed at 2, and the number of targets varied from 4 to 20 in steps of 2.




Figure 11. Performance of the real-robot system: The number of targets was fixed at 4, and a total of 2 sensors were used in three different combinations (with only two stationary sensors, with one stationary sensor and one mobile robot, and with two mobile robots). The experiment was performed on the second floor of our Computer Science Building shown in Figure 8.

6.2. ROBOT DENSITY VERSUS VISIBILITY

The simulations were performed with various configurations using the *Corridors* environment. 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. The result is shown in Figure 10.

6.3. MOBILE ROBOTS VERSUS EMBEDDED SENSORS

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 11. When the robots were used, we performed the experiment twice, using both the *Region-based Strategy* and the *Local-following Strategy*.

Table II. Significance values from T-test as a function of number of targets and environment characteristic

# 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

7. Discussion

7.1. REGION-BASED VERSUS LOCAL-FOLLOWING STRATEGY

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 strategies 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 II shows the significance values of the t-test; a small number implies that the two distributions are statistically different. As expected, the significance values of *OpenSpace* were big, and those of *Corridors* and *EmptySpace* were small.

The simulation results show a strong correlation between robot coordination strategies and environment shape. The data in Figure 9 suggest that in environments with a low value of p_0 , the *Region-based Approach* out performs the *Local-following Strategy*, and in environments with a high value of p_0 , the reverse is true. However, there are unobserved factors which might be critical to the overall performance, i.e. region decomposition. The *Region-based Approach* may perform differently when an environment is decomposed into few big regions and when the same environment is divided into many fine regions. Also, the region decomposition method may have to be varied depending on environment shape; for instance, region size should be small when the complexity of an environment is high, and vice versa. We plan to investigate these issues in future work.

Table III. Significance values from T-test as a function of number of targets and different strategies

# targets	vs Robot Density	vs Local-following
4	0.007200	0.067867
6	0.000737	0.000077
8	0.114616	0.217457
10	0.008551	0.551933
12	0.117852	0.676290
14	0.013156	0.743729
16	0.020076	0.300305
18	0.049254	0.433392
20	0.003620	0.223343

7.2. ROBOT DENSITY VERSUS VISIBILITY

As shown in the t-test result (Table III), the *Region-based Approach* using the *Visibility* did not show significant difference from the *Local-following Strategy*. There are two possible explanations for this.

First, the coverage estimation method might have been inappropriate. Because of lack of global, geometrical information and real-time requirement, simple approximation methods were utilized as much as possible. For example, each sensor’s coverage was computed by integrating real laser readings, but a single simplified sector was used for overlap computation. Also, small error was ignored. For instance, overlap was computed for all possible pairs of robots within a region; when the coverages of three robots (A , B , and C) are overlapped, the coverage intersection $A \cap B \cap C$ is subtracted once more unnecessarily and never compensated.

Second, the visibility maximization method described in Section 3.3 assumes that tracking sensor is omni-directional, which is not valid in both simulation and real robots. In the current visibility maximization method, when a robot detects other robots, it abandons tracking targets which are positioned closer to the other robots. By doing this, robots can keep distance from each other without explicit communication. In Figure 4 (c) this method works properly since the upper robot is facing down. However, if the upper robot were facing up, the lower robot would still move same and the other two targets would not be tracked anymore. This problem is caused because the current tracking sensor cannot detect other robots’ heading.

We expect that better implementation of visibility maximization method would improve its overall performance.

7.3. MOBILE ROBOTS VERSUS EMBEDDED SENSORS

The general tracking performance in the real-robot experiments was not as good as that in the simulation. There are several explanations. First, the effective range of the real vision system (around 4.5m) was much shorter than the range of the simulated one (8m). Also, the field of view (FOV) of the real camera (90 degrees) was not as wide as the FOV of the virtual camera (120 degrees). Second, the environment for the real-robot experiments was much more complex; for example, the connection between offices and corridors (doors) was narrow enough to cause the robots difficulties. Lastly, the odometry error of the real robot was larger, and caused unfavorable effects on navigation; for example, it took longer to traverse from a region to another region.



(a) 1 robot tracking 2 targets



(b) 2 robots tracking 1 target

Figure 12. Tracking examples: (a) shows the favorable case which observed more often when the *Region-based Strategy* is used. (b) shows the worst case which observed more often when the *Local-following Strategy* is used.

In case of using two stationary sensors, the performance depends on targets' movement. As shown in Figure 11, 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 12 (a), the robot was able to position itself properly to track multiple targets, 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 12 (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, i.e. larger-scale experiments with more combinations. We plan to address this in future work.

8. Conclusion and Future Work

In this paper, we presented a mobile sensor network that is able to track multiple targets effectively. The proposed *Region-based Approach* attempts to solve the sensor coordination problem without explicit negotiation among sensors by exploiting a topological map. All mobile sensors maintain urgency information independently, and re-position themselves based on that information. The *Region-based Approach* has been implemented with a behavior-based system which is fully distributed and scalable. Through intensive simulations and real-robot experiments, the *Region-based Approach* was shown to perform better than a naive approach (used as a baseline for comparison) when the environment contains significant occlusion. We used a simple metric for measuring the degree of occlusion, based on the average mean free path of a random line segment drawn in the environment.

Our experiments open up two new lines of research, which suggest a. an optimal ratio of robots to stationary sensors may exist for a given environment with certain occlusion characteristics, and b. an optimal ratio of mobile to stationary sensors may exist for a particular environment.

There are several open problems to be investigated in the future. First, the relationship between our *Region-based Approach* and the region decomposition method (responsible for the structure of the topological map) should be clarified. We assumed, in this paper, that the

topological map for an environment is given *a priori*, and focused on the relationship between the robot coordination method and the structure of the environment. However, the topology of the decomposed regions may be a critical factor in the overall tracking performance. We plan to integrate an on-line map construction module with the current system in order to generate the topological map for our tracking strategy.

Second, further research on the optimal combination of mobile robots and stationary sensors to maximize overall tracking performance is required. We performed experiments with real robots in three different combinations in this paper, but larger scale results are needed to be conclusive.

Acknowledgements

This work is supported in part by DARPA grant DABT63-99-1-0015 and NSF grants IIS-013394, ANI-0082498, and EIA-0121141.

References

- Bar-Shalom, Y. and W. D. Blair (eds.): 2000, *Multitarget-Multisensor Tracking: Applications and Advances*, Vol. 3. Artech House.
- Cazals, F. and M. Sbert: 1997, 'Some Integral Geometry Tools to Estimate the Complexity of 3D Scenes'. Technical Report RR-3204, Institut National de Recherche en Informatique et en Automatiue (INRIA).
- Choset, H. and J. Burdick: 1995a, 'Sensor Based Planning, Part I: The Generalized Voronoi Graph'. In: *Proceedings of the 1995 IEEE International Conference on Robotics and Automation*, Vol. 2. Nagoya, Japan, pp. 1649–1655.
- Choset, H. and J. Burdick: 1995b, 'Sensor Based Planning, Part II: Incremental Contruction of the Generalized Voronoi Graph'. In: *Proceedings of the 1995 IEEE International Conference on Robotics and Automation*, Vol. 2. Nagoya, Japan, pp. 1643–1648.
- Chung, J. and H. S. Yang: 1995, 'Fast and Effective Multiple Moving Targets Tracking Method for Mobile Robots'. In: *Proceedings of the 1995 IEEE International Conference on Robotics and Automation*. pp. 2645–2650.
- Coue, C. and P. Bessiere: 2001, 'Chasing an Elusive Target With a Mobile Robot'. In: *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Maui, Hawaii, pp. 1370–1375.
- Dedeoglu, G., M. J. Matarić, and G. S. Sukhatme: 1999, 'Incremental, on-line topological map building with a mobile robot'. In: *Proceedings of Mobile Robots*, Vol. XIV. Boston, MA, pp. 129–139.
- Dedeoglu, G. and G. S. Sukhatme: 2000, 'Landmark-based Matching Algorithm for Cooperative Mapping by Autonomous Robots'. In: *Distributed Autonomous Robotic Systems (DARS)*. Knoxville, Tennessee.
- Gerkey, B. and M. J. Matarić: 2001, 'Principled Communication for Dynamic Multi-Robot Task Allocation'. In: D. Rus and S. Singh (eds.): *Experimental Robotics*, Vol. LNCIS 271 of VII. Springer-Verlag Berlin Heidelberg, pp. 353–362.
- Gerkey, B., R. Vaughan, K. Stoy, A. Howard, G. S. Sukhatme, and M. J. Matarić: 2001, '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, pp. 1226–1231.
- Guibas, L. J., J.-C. Latombe, S. M. LaValle, D. Lin, and R. Motwani: 1997, 'A Visibility-Based Pursuit-Evasion Problem'. *International Journal of computational eometry and Applications* **9**(5), 471–494.
- Horling, B., R. Vincent, R. Mailler, J. Shen, R. Becker, K. Rawlins, and V. Lesser: 2001, 'Distributed Sensor Network for Real Time Tracking'. In: *Proceedings of the 5th International Conference on Autonomous Agents*. pp. 417–424.
- Howard, A., M. Matarić, and G. Sukhatme: 2001, 'Relaxation on a Mesh: a formalism for Generalized Localization'. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. Wailea, Hawaii.
- Jung, B. and G. S. Sukhatme: 2001a, '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*. pp. 206–211.
- Jung, B. and G. S. Sukhatme: 2001b, 'Cooperative Tracking with Mobile Robots and Networked Embedded Sensors'. Institute for Robotics and Intelligent Systems Technical Report IRIS-01-404, University of Southern California.

- Jung, B. and G. S. Sukhatme: 2001c, '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.
- Kluge, B., C. Kohler, and E. Prassler: 2001, '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*. pp. 1683–1688.
- Koyasu, H., J. Miura, and Y. Shirai: 2001, '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, pp. 31–36.
- Kuipers, B. and Y.-T. Byun: 1991, 'A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations'. *Journal of Robotics and Autonomous Systems* **8**, 47–63.
- Kumar, S., C. Alaettinoglu, and D. Estrin: 2000, 'SCalable Object-tracking through Unattended Techniques (SCOUT)'. In: *Proceedings of the 8th International Conference on Network Protocols*. Osaka, Japan.
- LaValle, S. M., H. H. Conzalez-Banos, C. Becker, and J.-C. Latombe: 1997, 'Motion Strategies for Maintaining Visibility of a Moving Target'. In: *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*.
- Lindstrom, M. and J. O. Eklundh: 2001, '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, pp. 1364–1369.
- Matarić, M. J.: 1997, '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* **9**(2-3), 67–83.
- Parker, L. E.: 1997, 'Cooperative Motion Control for Multi-Target Observation'. In: *Proceedings of the 1997 IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 1591–1598.
- Parker, L. E.: 1999, 'Cooperative Robotics for Multi-Target Observation'. *Intelligent Automation and Soft Computing, special issue on Robotics Research at Oak Ridge National Laboratory* **5**(1), 5–19.
- Schultz, D., W. Burgard, D. Fox, and A. B. Cremers: 2001, '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*. pp. 1165–1170.
- Vaughan, R. T.: 2000, 'Stage: A Multiple Robot Simulator'. Institute for Robotics and Intelligent Systems Technical Report IRIS-00-393, University of Southern California.
- Wenger, B. B. and M. J. Matarić: 2000, 'Broadcast of Local Eligibility for Multi-Target Observation'. In: *Proceedings of Distributed Autonomous Robotic Systems*.
- Yamashita, M., H. Umemoto, I. Suzuki, and T. Kameda: 1997, 'Searching for Mobile Intruders in a Polygonal Region by a Group of Mobile Searchers'. In: *Symposium on Computational Geometry*. pp. 448–450.