In *the 2004 IEEE International Conference on Robotics and Automation (ICRA)*
pp. 2189-2195, New Orleans, LA, April 26 - May 1, 2004.

A Generalized Region-based Approach for Multi-target Tracking in Outdoor Environments

Boyoon Jung* and Gaurav S. Sukhatme†

Robotic Embedded Systems Laboratory
Center for Robotics and Embedded Systems
University of Southern California
Los Angeles, California 90089-0781

Email: *boyoon@robotics.usc.edu**, *gaurav@robotics.usc.edu*†

***Abstract*— We propose a generalized region-based approach to multi-target tracking, which is applicable to structured and unstructured environments. In this approach each robot constructs virtual regions based on the latest tracking information from other robots. Without pre-partitioned region information, each robot independently estimates the most urgent region that needs to be visited. The idea is for robots to coarsely estimate where the targets are present, and to navigate there. A multi-robot system to track moving objects outdoors has been designed using this approach in order to validate the idea. The performance of the individual motion trackers and the cooperative tracking behaviors is evaluated through experiments with different robot bases (a helicopter, a Segway RMP, and a Pioneer) and in simulation. Experimental results indicate that robots are able to distribute themselves appropriately in response to target movement.**

I. Introduction

Automated surveillance or security systems require capability to track the positions of multiple targets autonomously and effectively. Using a group of mobile robots for multi-target tracking is beneficial in the sense that a mobile robot can cover a wide area over time (which means the number of sensors required is smaller) and can adapt to targets' movement patterns for more efficient tracking. Especially, when the number of targets is much bigger than the number of sensors available or when sensors cannot be deployed in advance, the mobility of sensors become indispensable.

The multi-target tracking problem using a group of mobile robots can be treated as a task allocation problem; given a group of agents (*robots*) and a group of tasks (*targets*), assign each task to a proper agent so that the overall performance (*the total number of tracked targets over time*) is maximized. Various distributed algorithms have been proposed. The ALLIANCE architecture [1] achieves target-assignment by the interaction of motivational behaviors. If a target was not tracked for a while, the robot which was supposed to track the target would give up and another robot in a better position would take up the target. In the BLE architecture [2], 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. The Murdoch architecture [3] showed that target-assignment 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.

There have been other approaches that determine robot poses without explicit target assignment, especially when the ratio of the number of robots to the number of targets is close to 1.0. In [4], the configuration of a team of mobile robots was actively controlled by minimizing the expected error in tracking target positions, and a reactive motion planner was reported in [5] that maximizes the shortest distance that a target needs to move in order to escape an observer's visibility region.

We proposed the *Region-based Approach* [6] for multi-target tracking in an *indoor* environment, which converts the target-assignment problem into a simpler region-assignment problem by utilizing the characteristics of indoor environments. After constructing the topological map of a naturally-partitioned environment, robots are assigned to sparse regions by evaluating density estimates of each region.

However, the *Region-based Approach* has a restricted application in the sense that it requires a clear partition and a topological map of an environment. Outdoor environments, for example, do not have clear partition or structure, and it is not trivial to build a topological map for cooperation and navigation. In this paper, we propose a generalized *Region-based Approach*, which is applicable to any environment by constructing virtual regions based on the latest sensor information.

The rest of the paper is organized as follows. Section II provides the brief description of the original *Region-based Approach*, and the extension for general environments is explained in Section III. The control architecture and implementation issues are discussed in Section IV, and the result is analyzed in Section V. The current status and possible improvements are discussed in Section VI.

II. Region-based Approach

The *Region-based Approach* [6] for target tracking is based on two fundamental assumptions: (a) *the environment can be*

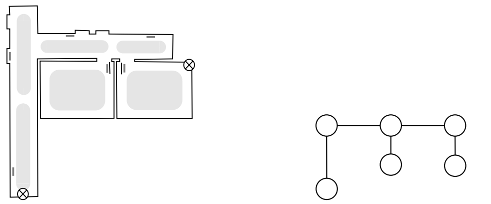

(a) Indoor environment (b) Topological Map

Fig. 1. Typical Office-type Environment

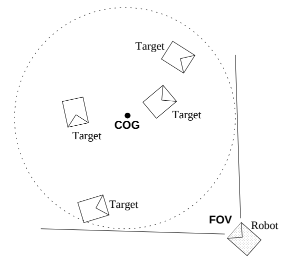

Fig. 2. Following targets within a region

divided into topologically simple convex regions using certain landmarks as demarcaters. (b) *For two comparably sized regions, more robots should be deployed in the one with the higher number of targets.* Figure 1 shows a typical indoor environment and the associated topological map shown as a graph. The nodes in the graph correspond to regions, and links correspond to landmarks. Each region is assumed to be simply connected[1].

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, and 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.

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

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

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.

A. The Coarse Deployment Strategy

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

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

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.

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.

$$\begin{aligned} U(R) &= uf\left(D_t(R), D_r(R) \right) \times \tfrac{d_{avg}}{d} \\ uf(D_r, D_t) &= \begin{cases} \frac{D_t}{D_r} & \text{if } D_r \neq 0 \\ D_t \times \alpha & \text{if } D_r = 0 \;\&\; D_t \neq 0 \\ 1.0 & \text{if } D_r = D_t = 0 \end{cases} \\ \alpha &= \frac{\text{area of region } R}{\text{unit coverage}} \times \theta_d \end{aligned} \tag{4}$$

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.

B. 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

[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.

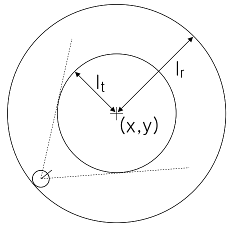

Fig. 3. Parameterized virtual region

attempts to keep at a distance $d_{max}/sin(FOV/2)$ from the COG where d_{max} is the distance between the COG and the target furthest from the COG in the group.

III. Generalized Region-based Approach

For the *Region-based Approach*, the density estimates are maintained for nodes on a topological map, which are discrete and sparse. The advantage of the approach is that the computation for robot distribution is cheap, but there is a restriction such that it is applicable only to structured, indoor (office-type) environments. The *Region-based Approach* can be extended for unstructured, outdoor environments by estimating *virtual* regions based on the latest tracking information from all robots.

A virtual region is represented by four parameters: a center position (x, y), a radius l_r for robot density estimation, and a radius l_t for target density estimation. As shown in Figure 3, the radius l_r and l_t are determined based on the sensor range and the field-of-view of a target tracker.

Since the areas of all regions are identical, for a given position (x, y), two density estimates are computed simply as follows:

$$\begin{aligned} D_r(x, y, l_r) &= \text{the number of robots within a distance } l_r \\ D_t(x, y, l_t) &= \text{the number of targets within a distance } l_t \end{aligned} \tag{5}$$

The formula to check if a robot is available to migrate to a more urgent region is re-defined as Equation 6, when $T(x, y, l_t)$ is the number of targets tracked by a robot in the region $R(x, y, l_r, l_t)$.

$$\max_{\forall x,y} \left[\frac{D_t(x, y, l_t)}{D_r(x, y, l_r)} \times T(x, y, l_t) \right] < \theta_t \tag{6}$$

Similar to the *Region-based Approach*, the urgency function for every virtual region is defined as in Equation 7, where d is the Euclidean distance between the current robot position and the center of a virtual region (x, y).

$$\begin{aligned} U(x,y) &= \begin{cases} \frac{D_t(x,y,l_t)}{D_r(x,y,l_r)} \times f(d) & \text{if } D_r \neq 0 \\ D_t(x,y,l_t) \times \alpha \times f(d) & \text{if } D_r = 0 \,\&\, D_t \neq 0 \\ 1.0 & \text{if } D_r = D_t = 0 \end{cases} \\ f(d) &= \quad e^{-\frac{1}{2} \times \left(\frac{d}{\theta_d}\right)^2} \end{aligned} \tag{7}$$

Each robot re-estimates the urgency function at every cycle, and selects the region with the highest urgency value as a destination. By changing the parameter θ_d of the function $f(d)$, the chance for a robot to select a closer region with less urgency over a farther region with higher urgency can be controlled.

The *Generalized Region-based Approach* is more efficient in the sense that (a) the best region boundary is selected according to the tracker properties, (b) targets are grouped automatically, and (c) motion planning within a region is not required anymore because a virtual region moves adaptively based on the target motions. However, the urgency function $U(x, y)$ is continuous, and the state space is unbounded. These disadvantages can be overcome by limiting the effective search space as described in Section IV.

IV. System Design and Implementation

For most surveillance or security applications, motion may be the most interesting feature to track. Therefore, we designed a multi-robot system that tracks and reports the positions of moving objects in outdoor environments. The *Generalized Region-based Approach* was utilized for the cooperation among robots. The system architecture is shown in Figure 4.

A. Motion Tracker

Our approach to motion tracking using a single camera on a mobile robot is described fully in [7]. We describe it briefly here. For motion detection using a single camera, frame differencing, which compares two consecutive image frames and finds moving objects based on the difference, is perhaps the most intuitive and fast algorithm, especially when the viewing camera is static. However, when the camera moves (*eg.* when it is mounted on a mobile robot), straightforward differencing is not applicable because a big difference is generated by simply moving the camera even if nothing moves in the environment. There are two *independent motions* involved in the moving camera scenario: motions of moving objects and the ego-motion of the camera. Since these two motions are blended into a single image, the ego-motion of the camera should be eliminated so that the remaining motions, which are due to moving objects, can be detected. Figure 5 shows the processing sequence of our motion detection algorithm. Frame differencing is utilized, but the the ego-motion of the camera in the previous image is compensated before comparing it with the current image.

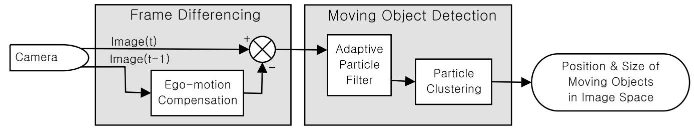

Fig. 5. Processing sequence for motion detection

The ego-motion of the camera can be estimated by tracking features between images [8]. When the camera moves, two consecutive images, I^t (the image at time t) and I^{t-1} (the image at time $t-1$), are in different coordinate systems. Ego-motion compensation is a transformation from the image coordinates of I^{t-1} to that of I^t so that the two images can be

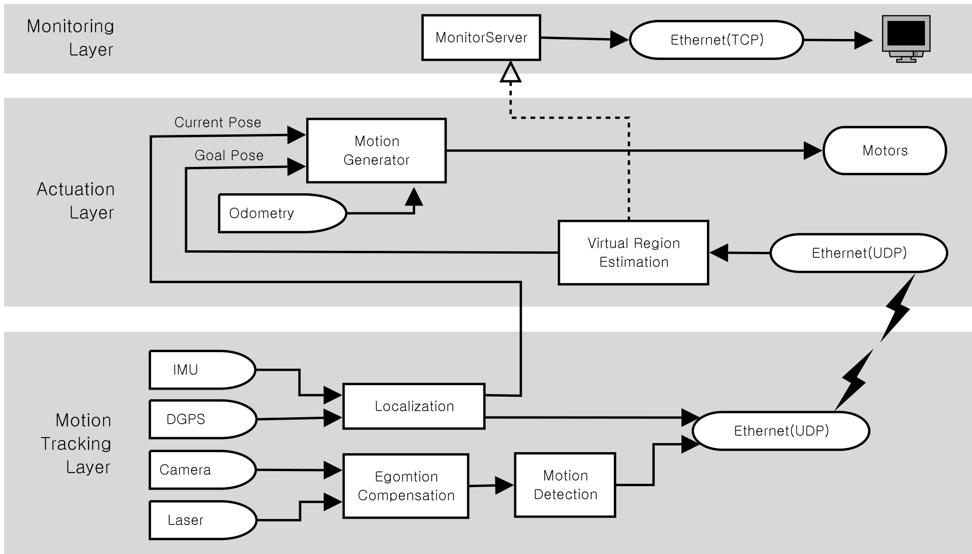

Fig. 4. System architecture for *Generalized Region-based Approach*

compared directly in the same coordinates. The transformation can be estimated using two corresponding feature sets: a set of features in I^t and a set of corresponding features in I^{t-1}. However, since there are independently moving objects in the images, a transform model and an outlier detection algorithm were designed so that the result of ego-motion compensation is not sensitive to those object motions.

Real outdoor images are contaminated by various noise sources, *eg.* poor lighting conditions, camera distortion, unstructured and changing shape of objects, etc. Thus perfect ego-motion compensation is rarely achievable. Even assuming that the ego-motion compensation is perfect, the difference image would still contain structured noise on the boundaries of objects because of the lack of depth information from a monocular image. Some of these noise terms are transient and some of them are constant over time. We use a probabilistic model to filter them out and to perform robust detection and tracking. The probability distribution of moving objects in image space is estimated using an adaptive particle filter [9] as shown in Figure 6 (a), and the final particles are clustered using a mixture of Gaussians [10] for the position estimation (Figure 6 (b)).

The depth information can be retrieved from a laser rangefinder. Given the optical properties of a camera and the geometry between the camera and the laser rangefinder, the distance information from the laser rangefinder can be projected onto the image coordinates as shown in Figure 7. As a result, the image pixels at the same height as the laser rangefinder will have depth information. For ground robots, this partial 3D information can be enough assuming all moving objects are on the the same plane as the robot.

The tracking results (the positions of moving objects in a local coordinate system) and the tracker information (the robot pose in the global coordinate system) are broadcasted for cooperation over the wireless network.

(a) Particle filter outputs

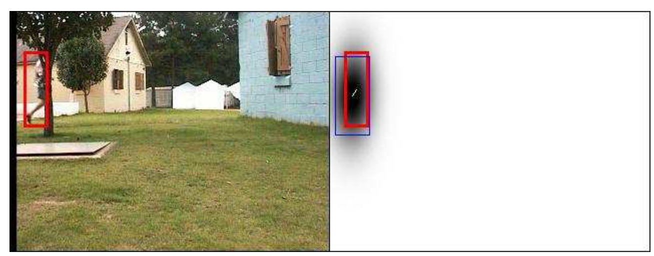

(b) Particle clustering result

Fig. 6. Motion tracking using a single camera

B. Localization

All robots share the global coordinates for multi-robot cooperation, which requires a global localization. Unlike the indoor case, GPS (Global Positioning System) is available outdoors. Combination of a differential GPS and an IMU (Inertial Measurement Unit) provides full pose (x, y, α) correction. An extended Kalman Filter [11] is implemented for stable pose estimation. The state vector is defined as

$$\boldsymbol{x} = [x\ \dot{x}\ y\ \dot{y}\ \alpha\ \dot{\alpha}]^T \tag{8}$$

and the user input vector is defined as

$$\boldsymbol{u} = [v_t\ v_r]^T \tag{9}$$

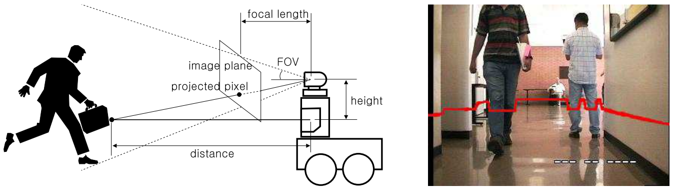

Fig. 7. Projection of laser scans onto the image coordinates

when v_t is the translational velocity and v_r is the rotational velocity. The dynamics equations are

$$\begin{array}{rcll} x^{t+1} & = & x^t + \Delta t \times \dot{x}^t & + w_0^t \\ \dot{x}^{t+1} & = & v_t^t \times \cos(\alpha + \Delta t \times v_r^t) & + w_1^t \\ y^{t+1} & = & y^t + \Delta t \times \dot{y}^t & + w_2^t \\ \dot{y}^{t+1} & = & v_t^t \times \sin(\alpha + \Delta t \times v_r^t) & + w_3^t \\ \alpha^{t+1} & = & \alpha^t + \Delta t \times \dot{\alpha}^t & + w_4^t \\ \dot{\alpha}^{t+1} & = & v_r^t & + w_5^t \end{array} \tag{10}$$

and the measurement vectors are defined as

$$\boldsymbol{z_{gps}} = [x\ y]^T + \boldsymbol{v_{gps}} \tag{11}$$

$$\boldsymbol{z_{imu}} = [\alpha] + \boldsymbol{v_{imu}} \tag{12}$$

when $\boldsymbol{w}$, $\boldsymbol{v_{gps}}$, and $\boldsymbol{v_{imu}}$ are noise terms.

C. Cooperative Motion Planning

The *Generalized Region-based Approach* described in Section III is utilized for cooperative motion planning. The motion trackers on each robot broadcasts its global pose and detected target positions over the wireless network. Each robot *independently* collects those broadcasted messages, and maintains the global position information of robots and targets in the environment. Based on the latest position information, each robot determines if it is available to migrate to other regions using Equation 6. Once the robot decides to move, it searches for the most urgent region in an environment using Equation 7, and sets the center of the region as a goal position.

There are two implementation issues: (a) how to cope with the unbounded state space, and (b) how to represent the continuous urgency function $U(x, y)$ efficiently. Normally outdoor environment is unbounded, and the urgency function $U(x, y)$ is supposed to be evaluated over the entire space theoretically. However, the effective search space can be limited by the fact that the evaluation is performed based on the detected targets only. In other words, the value of $U(x, y)$ will be a constant when the distance from the position (x, y) to the position of any robot in the system is bigger than a certain threshold. In out system, the threshold is $\theta_s = R + l_t$ when R is the effective range of a tracker. Therefore, the size of the effective search space is proportional to the number of robots in the system.

There are two popular methods to represent a continuous function: grid method and non-uniform (adaptive) sampling. The non-uniform sampling is beneficial when the search space is big and the distribution of function values are concentrated on several areas (*eg.* the entropy of the function is small). In our system, the search space is relatively small, and targets can move in and out of the search space quite often. Therefore, a grid method is adopted for the urgency function representation.

D. Navigation

Given the current robot pose (from the *Localizer* module) and the goal position (selected by the *Generalized Region-based Approach*), a robot should be able to perform point-to-point, safe navigation. In a structured indoor environment, wall-following behavior can be a good navigation mechanism since the environment provides basic guidance from a region to another. However, in outdoor environment, other type of local navigation algorithm is required to work without any environmental support.

VFH+ (Vector Field Histogram +) [12] algorithm is implemented for point-to-point navigation. VFH+ algorithm provides a natural way to combine a local occupancy grid map and the potential field method, and the dynamics and kinematics of a mobile robot can be integrated to generate an executable path. In addition, the robot's motion property (*eg.* goal-oriented, energy-efficient, or smooth-path) can be controlled by changing the parameters of a cost function.

Both translational and rotational velocities are controlled by the VFH+ algorithm until the distance between the robot position and the goal position is bigger than l_r. Once a robot is positioned at the distance l_r from a goal position, then only rotational velocity is controlled by the VFH+ algorithm while a translational velocity remains zero. In this manner, a robot would stay close enough not to lose any target, and far enough to keep all targets within its FOV.

V. Result and Discussion

A. Motion Tracker

It is obvious that the performance of the cooperative tracking is affected by the performance of the motion tracker. The performance of the motion tracker was evaluated [**?**] on three different robot platforms: robotic helicopter, Segway RMP, and Pioneer2 AT. Each platform has unique characteristics in terms of its ego-motion. The *Robotic Helicopter* in Figure 8 (a) is an autonomous flying vehicle carrying a monocular camera facing downward. Once it takes off and hovers, planar movements become the main motion, and moving objects on the ground stay at a roughly constant distance from the camera most of the time; however, pitch and roll motions still generate complicated video sequences. The *Segway RMP* in Figure 8 (b) is a two-wheeled, dynamically stable robot with self-balancing capability. It works like an inverted pendulum; when the robot accelerates/decelerates, the pitch angle increases seriously. Therefore, the ego-motion compensation step should be able to cope with not only planar movements but also pitch motions. The *Pioneer2 AT* in Figure 8 (c) is a typical four-wheeled, statically stable robot. Since the Pioneer2 robot is the most statically stable platform among these three platforms, we drove it on the most severe (rocky and occlusive) environment.

The tracking performance statistics are summarized in Table I. The *Robotic helicopter* result shows that the tracking algorithm missed six objects, but five of them were the

(a) Robotic Helicopter (b) Segway RMP (c) Pioneer2 AT

Fig. 8. Robot platforms for motion tracker evaluation

Platform	Detected	True +	Rate	Error
Helicopter	29 / 35	29	82.86 %	13.26
Segway RMP	208 / 220	206	93.63 %	20.29
Pioneer2 AT	126 / 172	114	66.28 %	13.54

TABLE I
PERFORMANCE OF THE MOTION TRACKER

cases when a moving object was introduced and showed only partially on the boundary of the image plane. For the *Segway RMP* result, the detection rate was satisfactory, but the average distance error[2] was larger than the others. The reason was that the walking person was closer to the robot and the tracking algorithm often detected the upper body only, which caused a constant distance error. The *Pioneer2 AT* result was the worst; however, the terrain for the experiment was more challenging and the input images were more blurred and unstable.

B. Cooperative Tracking

The performance of the *Generalized Region-based Approach* was studied by inspecting the changes of urgency distribution caused by the position changes of robots and targets in a simulated environment. The environment was an empty, unbounded space, and there were three robots and ten moving targets in it. The target motions were random, and the grid size was fixed to 1 meter for the urgency function representation.

Figure 9 shows how a robot (on the bottom-left) determines a goal position according to an estimated urgency distribution. The gray-scale of each grid indicates its urgency value estimated by the robot; the darker the color is, the higher the urgency value is. It is obvious in Figure 9 (a) that the robot selected the center of the targets on the right as a goal position because there is a peak of the urgency distribution. In Figure 9 (b), it is clearly shown that the urgency distribution becomes balanced after the robot participated in tracking the group of targets on the right.

A switching-region behavior is shown in Figure 10. The center robot was tracking the group of targets on the left initially as shown in Figure 10 (a). However, when the robot on the right discovers more targets, the urgency distribution estimated by the center robot skews to the right (Figure 10 (b)). As a result, the center robot migrates to help tracking the group of targets on the right, and the final urgency distribution is

[2] The average Euclidean distance in pixels; it should not be considered as actual error measurement since the ground truth was the center of objects, not the center of motions.

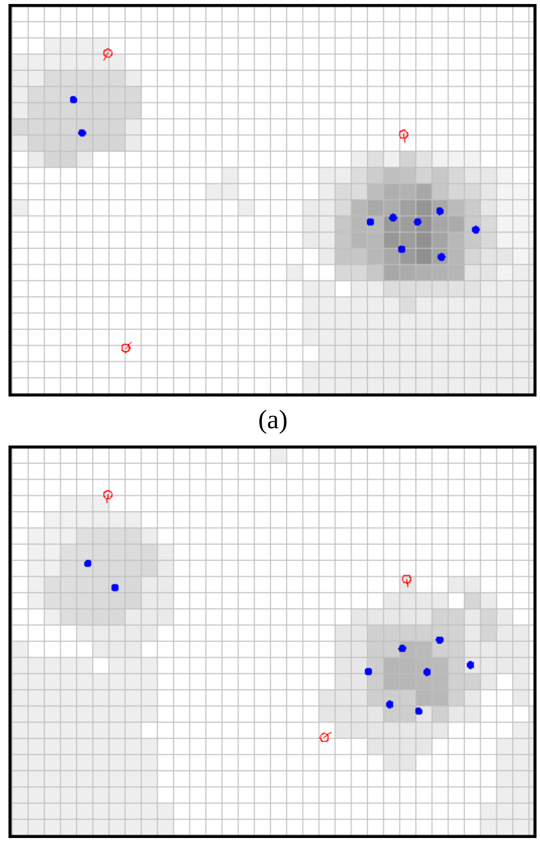

(a)

(b)

Fig. 9. Virtual region selection behavior

shown in Figure 10 (c). It is clear that the *Generalized Region-based Approach* controls robots in the same manner of the original *Region-based Approach*.

It should be clarified that the urgency distributions shown in Figure 10 were the estimations performed by the robot in the center. Because of the function $f(d)$ in Equation 7, the urgency estimation of each robot is different from others, which is the reason why the left robot did not migrate to the right while the central one did.

VI. CONCLUSION AND FUTURE WORK

The *Region-based Approach* for cooperative multi-target tracking was extended to unstructured, outdoor environments by constructing virtual regions based on the latest sensor information. This *Generalized Region-based Approach* was applied to the multiple moving objects tracking problem. In order to validate the approach, a motion tracker using a single camera and a laser rangefinder was developed and integrated with other modules for localization and safe navigation. The cooperation among robots was controlled by the *Generalized Region-based Approach*, and the study of urgency distribution changes confirmed that robots were properly distributed according to the target distribution. The generalized approach does not depend on a topological map of the environment, however it suffers from some limitations. These are discussed

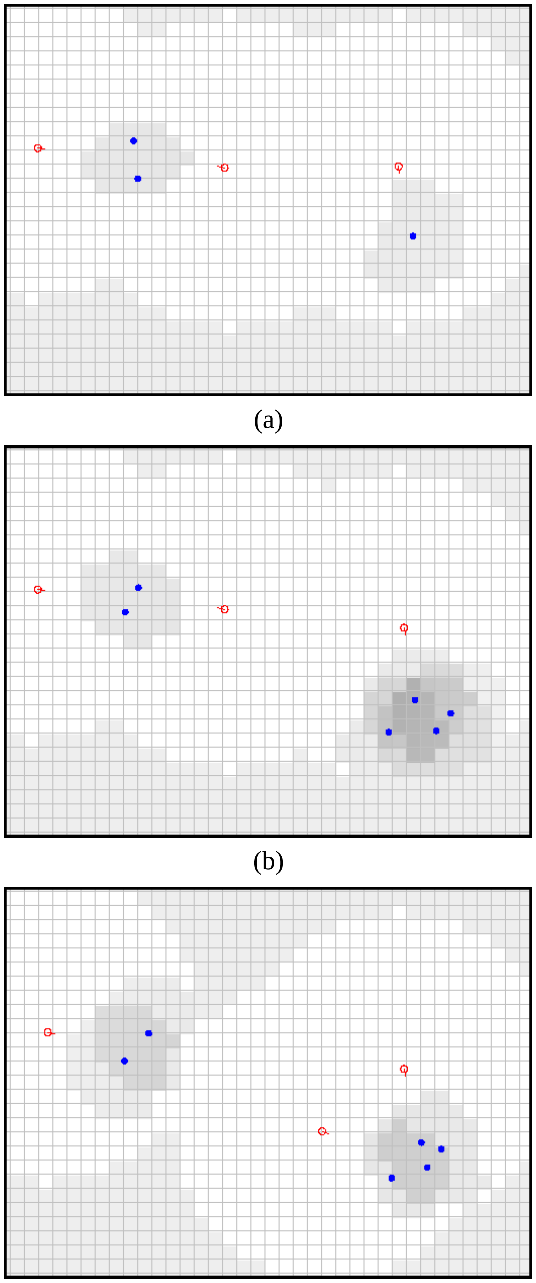

(a)

(b)

(c)

Fig. 10. Region-switching behavior

briefly below, and form the subject of ongoing and future work.

Since the algorithms does not 'remember' that a region was empty after a robot leaves the region, it is possible that a robot may oscillate between two empty regions; however, by setting the cost function parameters of the VFH+ navigation algorithm properly, this worst case can be avoided. We also plan to investigate explicit exploration strategies to control robot motion when it is available to move, and no urgent region needs servicing. The current system select a region randomly. By adopting a more sophisticated exploration strategy, performance enhancements can be expected. The individual motion tracker using a camera showed stable motion-tracking capability in 2D image space by adopting a probabilistic filter at the end of the processing sequence. However, the tracker sometimes generates transient position errors in 3D space when the rotational velocity of a robot is high. The conversion from 2D to 3D using a laser rangefinder assumes perfect synchronization between a camera and a laser rangefinder, which is unachievable in real-robot experiments. Since there is always a delay between the image capture time and the range retrieval time, laser range scans are projected at wrong places when a robot rotates at high speed. We believe that those errors can be reduced by adding another filter for 3D-position estimation.

Acknowledgments

This work is supported in part by DARPA grants DABT63-99-1-0015, and 5-39509-A (via UPenn) under the Mobile Autonomous Robot Software (MARS) program.

References

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

[2] B. B. Werger and M. J. Matarić, "Broadcast of local eligibility for multi-target observation," in *Proceedings of Distributed Autonomous Robotic Systems*, 2000, pp. 347–356.

[3] B. Gerkey and M. J. Matarić, "Principled communication for dynamic multi-robot task allocation," in *Experimental Robotics*, ser. VII, D. Rus and S. Singh, Eds., vol. LNCIS 271, Springer-Verlag Berlin Heidelberg, 2001, pp. 353–362.

[4] J. R. Spletzer and C. J. Taylor, "Dynamic sensor planning and control for optimally tracking targets," *International Journal of Robotics Research (IJRR)*, vol. 22, no. 1, pp. 7–20, January 2003.

[5] R. Murrieta-Cid, H. González-Baños, and B. Tovar, "A reactive motion planner to maintain visibility of unpredictable targets," in *the Proceeding of IEEE International Conference on Robotics and Automation*, Washington DC, May 2002, pp. 4242–4247.

[6] B. Jung and G. S. Sukhatme, "Tracking targets using multiple robots: The effect of environment occlusion," *Autonomous Robots*, vol. 13, no. 3, pp. 191–205, 2002.

[7] ——, "Detecting moving objects using a single camera on a mobile robot in an outdoor environment," in *International Conference on Intelligent Autonomous Systems*, The Netherlands, March 2004, pp. 980–987.

[8] A. Censi, A. Fusiello, and V. Roberto, "Image stabilization by features tracking," in *Proceedings of the 10th International Conference on Image Analysis and Processing*, Venice, Italy, September 1999, pp. 665–667.

[9] D. Fox, "KLD-sampling: Adaptive particle filter," in *Advances in Neural Information Processing Systems 14*. MIT Press, 2001.

[10] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2001.

[11] G. Welch and G. Bishop, "An introduction to the kalman filter," Department of Computer Science, University of North Carolina at Chapel Hill, Tech. Rep. 95-041.

[12] I. Ulrich and J. Borenstein, "VFH+: Reliable obstacle avoidance for fast mobile robots," in *Proceeding of the IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May 16–21 1998, pp. 1572–1577.