
Preliminary Results in Tracking Mobile Targets Using Range Sensors from Multiple Robots

Elizabeth Liao, Geoffrey Hollinger, Joseph Djugash, and Sanjiv Singh

Carnegie Mellon University {eliao@andrew.cmu.edu, ghollinger@andrew.cmu.edu,
robojoe@cmu.edu, ssingh@ri.cmu.edu}

Summary. In urban search and rescue scenarios, human first responders risk their lives as they routinely encounter hazardous environments. A team of robots, equipped with various sensors, deployed in such an environment can be used to track emergency personnel such as firefighters, reducing the risk to human life. This paper explores techniques for tracking a mobile target and coordinating a team of robots, equipped with range-only sensors, through smoke-filled, high-temperature environments. The particular strengths of our tracking and cooperative control algorithms are identified through a set of simulated examples.

1 Introduction

Smoke, darkness, and clutter handicap first responders during time critical emergency situations. Robot teams can help by providing feedback on environmental hazards and tracking human agents in the environment. Creating and maintaining a communications network and tracking the location of emergency response personnel are examples of some of the tasks that the robot teams can perform.

Tracking a moving target using a single robot in a near-ideal, clutter-free environment is a solved problem. Environments with dense obstacle configurations, however, can enforce constraints that limit the robot's tracking capabilities. Since speed and efficiency are critical, the use of multiple robots also becomes desirable. Although the applications for a robot team's assistance are numerous, a necessary functionality of a robot team in this scenario is the ability to reliably and accurately track a human over time. In order to achieve this higher tracking accuracy, the proper coordination of the robot team is important and mandatory. In this paper we explore these two separate yet inter-related problems, tracking a mobile target and coordinating a robot team to further improve the achieved tracking accuracy.

At a high level, the job of a robot team is to track a human while he/she is within range of the robot team. In order to perform such a task, the robot team requires sensors that can provide sufficient information to help localize the human. Cameras, sonar and other commonly used sensors fail to provide

the necessary robustness and accuracy in the presence of smoke, darkness, and debris. Radio then becomes the obvious choice for our particular problem. Radio based sensors provide a unique advantage over the more commonly used sensors by delivering reliable measurements even under the presence of heavy smoke and debris. For these reasons, the results in this paper focus on the use of radio ranging sensors for tracking mobile targets.

The tracking and coordination algorithms presented in this paper are backed by a set of experiments conducted in simulation. Experiments in team coordination are performed with a four member robot team. Given only a floor map and an estimate of the robot positions, the robot team achieves reliable tracking of the mobile target.

2 Related Work

Robotics researchers have examined the potential for multiple robot assistance in urban search and rescue environments [9], and previous research has explored the use of multiple robots in indoor environments for mapping and intruder detection [6].

Recent advances in radio technology allow for range-only measurements between inexpensive transponders and receivers. These sensors can provide information in dynamic, noisy environments without the necessity of line-of-sight. To take advantage of these sensors, algorithms must be developed for range-only localization and tracking. Kantor and Singh present preliminary work in localizing robots using range-only measurements using an extended Kalman filter (EKF) and a particle filter in [7]. Djugash et al. give further research and testing of range-only localization and mapping in [2].

Other researchers have examined the potential for using prior knowledge from the map during tracking and localization. Liao et al. presents a system for limiting particle paths to the voronoi diagram of the map space in [12]. Additionally, Min Oh et al. have described a system for tracking humans in outdoor environments using GPS and map-based priors [14]. These papers use sensors that provide both range and bearing information, and they explore map environments with limited complexity.

Tracking using formation control of a multi-robot team can be adapted for the human tracking scenario. The system described by Saber et al. [15] tracks in formation by generating cost functions based on formation control, collision avoidance, and tracking. Fierro et al. present a framework for the cooperative control of a robot that is broken up into two major parts: control law selection and trajectory generation [3]. Naffin and Sukhatme present a method to dynamically grow and maintain formations through negotiations [13]. While these approaches address maintaining formations in free space with relatively few obstacles, they do not explicitly mention how they handle cluttered environments.

3 Tracking with Range-Only Data

The task of tracking humans with range-only data introduces severe nonlinearities into any tracking algorithm. In addition, the lack of human odometry and motion constraints further complicate the problem. This section presents two algorithms for tracking humans using range-only data. The first is a simple particle filter, while the second incorporates the knowledge of a map to improve the accuracy of the tracker.

3.1 Particle Filter Tracking

One technique for tracking using range-only sensor data can be derived using Monte Carlo methods. Our algorithm (a modified version of the one presented for localization by Kantor and Singh [7]) uses many particle estimates of the first responder's position, propagates these particles randomly, and then moves them towards the high probability regions provided by the range measurements. This calculates the maximum a posteriori estimate of the first responder's position.

Letting $x_p(k)$, $p \in \{1, 2, \dots, N_p\}$, be a distribution of particles at time k , where N_p is the number of particles in that distribution, the algorithm for updating the particle filter proceeds as follows:

1. Propagate the state of each particle by computing $\tilde{x}_p(k) = x_p(k) + \omega_p(k)$ where ω_p is generated as zero-mean Gaussian noise with covariance R .
2. Assuming that each measurement is independent, compute the weight of each particle from the product of the Gaussian PDFs associated with current range measurements ($m_r(k)$) from each robot

$$w_p = \prod_r p(r_p | m_r(k)) \quad (1)$$

where r_p is the distance between $x_r(k)$ and $\tilde{x}_p(k)$.

3. Normalize the weights so that $\sum_{p=1}^{N_p} w_p = 1$.
4. Resample the particles such that each $x_p(k+1)$ particle is selected from $\tilde{x}_i(k)$, and the probability of selecting each particle is w_i . The resulting distribution moves towards areas of high probability based on measurements from all robots.

The particle filter provides a method for tracking humans without the use of odometry data. Additionally, it avoids any linearization of the range measurement probability distribution function.

3.2 Using Prior Knowledge from the Map

To fully utilize the benefits of operating in an indoor environment, we extend our particle filter implementation to incorporate prior knowledge of the map. Given an estimate of the firefighter's position, the uncertainty of subsequent measurements is bounded by obstacles on the map. For instance, if we know

that a firefighter is traveling in a hallway, subsequent estimates should remain in that hallway. To take advantage of this insight, we make two modifications to the particle filter.

1. Particles are restricted from moving through walls. When calculating particle weights in Step 2, particles that have moved through walls are discounted by ten orders of magnitude.
2. Under the assumption that firefighters generally move in the center of corridors, the weights of particles are adjusted such that they equalize their distance to walls. This is accomplished by viewing the map as an additional 2D Gaussian measurement. The mean of this Gaussian is given by the center of the current corridor, and the standard deviation is determined by the width of the corridor for the minor-axis and infinity for the major-axis. In implementation, a large factor times the length of the corridor can be used as the standard deviation along the major-axis. This approach ensures that the particles stay within the bounds of the wall (the minor-axis) while still allowing the particles to explore the full length of the corridor (the major-axis). The use of a Gaussian provides a principled method for integrating the information from the map into the particle filter while still maintaining the assumption that firefighters tend to remain in the center of hallways. The use of a uniform probability distribution function, for instance, would not take this assumption into account. This changes the weight update step (Step 2) to the following:

$$w_p = p(\tilde{x}_p(k)|M) \prod_r p(r_p|m_r(k)) \quad (2)$$

where $p(\tilde{x}_p(k)|M)$ is the probability of the particle’s position given the Gaussian measurement from the map.

4 Multi-Robot Coordination

One method for increasing the accuracy of firefighter tracking with range-only sensors is through cooperative control of a multi-robot team. The goal of a cooperative control algorithm is to assist in tracking by producing desirable team configurations to reduce the uncertainty of the firefighter’s position.

Before trying to position robots into configurations, a metric is needed to identify *good* configurations. Preferable configurations will magnify sensor error less, assuming that ranging devices are inherently noisy. We use a metric called the Geometric Dilution of Precision (GDOP), which relates the configuration of ranging devices to location error. Kelly uses Dilution of Precision by evaluating a robot’s position error using landmark configuration for localization purposes [8]. We are in essence using a similar technique in reverse by having robots be the “landmarks”. A two dimensional version of the GDOP formulation [1] was used such that only longitude and latitude were incorporated into the calculations. Since the problem domain is only in two dimensions, we refer to 2D GDOP as GDOP for the remainder of this paper.

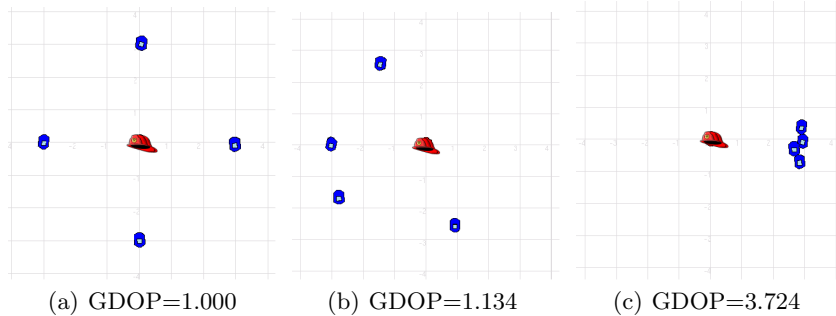


Fig. 1. Example four-robot configurations and their corresponding GDOP values. The values range from the best(left-most) to least desirable(right-most)

Fig. 1 shows example robot team configurations and the corresponding GDOP values. The firefighter is located in the center while the angles between the four team members are varied to form different configurations.

4.1 Cooperative Control Algorithm

In the cooperative control algorithm, each robot considers its teammates as part of the environment (stationary) and creates a plan to position itself with that belief. Here, a *plan* is defined as a path consisting of a series of waypoints ending with a desired goal. The only information available to create the plan are the positions of other team members and the firefighter. Key decisions of the algorithm are summarized below and are detailed by Liao [11].

4.2 Cost Map

A discretized cost map is used to decouple global planning from local navigation. Once a plan is produced, it can be passed off to another algorithm for local obstacle avoidance and waypoint following. Each robot initializes its cost map using an occupancy grid representing the floor plan. Every grid cell represents a possible goal location and is assigned a cost based on configuration of the team and distance from the firefighter assuming that the robot was located in that cell. The most desirable location of the robot corresponds to the cell with the lowest cost, while the best path to the goal is generated using a wavefront propagation path planner [10].

A weighted distance from the robots to the firefighter needs to be incorporated with the configuration cost into an overall cost so that robots are not too close to obstruct the human’s movement but not so far so that they are out of sensor range. This weighted wavefront distance function assigns the lowest value to an ideal distance and then increases in value as robots get farther from the ideal distance.

For the configuration policy, we developed the *Mean Angle* policy, which produces a score according to how much the configuration deviates from an equal angle spacing between robots. Mean angle can be calculated using the following equation:

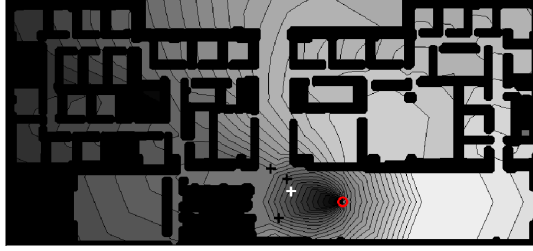


Fig. 2. Cost Contour map using weighted wavefront and mean angle policies. The light regions are the most favorable for the robot represented by the white cross.

$$\sum_{i=1}^n \exp \left| \frac{360^\circ}{n} - \angle i, \text{mod}(i+1, n) \right| - 1$$

This equation calculates a score by looking at the angles between adjacent robots with the firefighter as the vertex. Larger differences between an angle and the ideal angle will contribute to a higher score.

After integrating the Mean angle configuration cost with the weighted wavefront distance, the overall cost map is produced. Fig. 2 shows an overall cost contour map from the perspective of the white robot (far right cross). The light regions symbolize lower (better) cost, and dark regions symbolize higher (worse) cost. The black crosses represent other members of the robot team, and the circle represents the firefighter.

5 Simulation Results

Our cooperative control algorithm was implemented using software from the Player/Stage Project, an open source project that supports research in robots and sensors [4, 5]. Each robot executes the cooperative control algorithm by having an individualized version of the cost map. Initially the cost map is populated with obstacles contained in the provided floor plan. The distance and angle based costs are then filled into the remainder of the grid cells.

At every planning iteration, after obtaining the new positions of team members and the firefighter, the robot updates its cost map. The wavefront planner then generates a new path, which is passed off to the built-in Vector Field Histogram (VFH) [17] driver in Player/Stage for waypoint following and local obstacle avoidance. Since the algorithm implementation relies on frequent replanning, robots almost never reach their current goal and will typically choose different goals at every iteration.

The runtime of this implementation is highly dependent on how fine the floor plan is discretized in the cost map. The size of the cost map is an important factor to consider because every cell in the cost map is updated at the beginning of every iteration. Then, an exhaustive search over the cost map is used to locate the best position for a robot. Finally, the cost map is used again with wavefront planner to find the lowest cost path to the lowest cost cell.

This cooperative control algorithm is similar to Stroupe’s Move Value Estimation for Robot Teams (MVERT) architecture for action selection in multi-robot teams [16]. While a robot running MVERT tries to predict the actions



Fig. 3. Floor plan with initial robot positions and firefighter path. Initial position of the firefighter is the topmost point of the path.

of its teammates to plan for the next step, our cooperative control algorithm treats teammates as stationary objects and makes an entire plan to the goal.

5.1 Experiments in Player/Stage

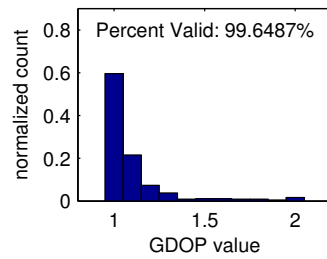
Currently, many real-world applications use sonar based line-of-sight (LOS) range sensors, but we would also like to investigate the added benefits of sensors that can range through walls (non-LOS ability). To quantify the difference between LOS and non-LOS ranging performance, experiments were performed using GDOP as the evaluation metric. Only robots with LOS to the firefighter contributed to GDOP in the appropriate experiments.

For each experiment, a robot representing the firefighter traveled on a pre-determined path while a team of four robots continually positioned themselves according to the commands generated by the cooperative control algorithm. During each iteration, GDOP was calculated using only the firefighter position and robot positions within sensor range of the firefighter. Since the algorithm ideally performs better without obstacles, experiments were run comparing the original floor plan (Fig. 3) to ones run using an “empty” floor plan (the original floor plan without interior walls).

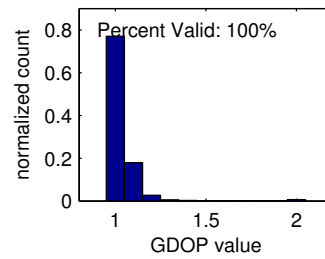
Histograms of the experiments (Fig. 4) show bins of GDOP values recorded during the entire firefighter path. The x-axis shows GDOP bin values while the y-axis shows the normalized histogram counts. Normalization was needed to compare experiments with different length data sets. The percent valid shown on each graph is the amount of time a valid GDOP value was obtained (i.e. when 2 or more robots are in range). Although histograms 4(a) and 4(b) should intuitively be identical, the discrepancy arises by having outside building walls remain in the “empty” floor plan. In the LOS case on an empty floor plan, it is possible that the exterior building walls block LOS to firefighter. This may in turn lead to configurations that evaluate to higher GDOP values.

5.2 Simulated Tracking Results

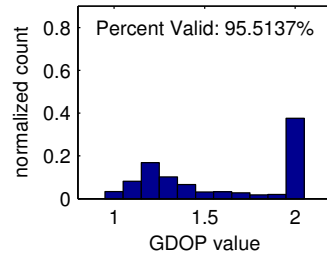
We ran simulated tests at various noise levels to determine the performance of the particle filter tracking with and without the map. Table 1 gives the characteristics of the three noise levels. During the runs, a simulated first responder moves throughout the entire map including both open and cluttered spaces (a total space of about 50x30 meters). Four robots move around the map using the coordination strategy described in Sect. 4 and attempt to localize the firefighter using simulated ranging measurements. Since the measurements are



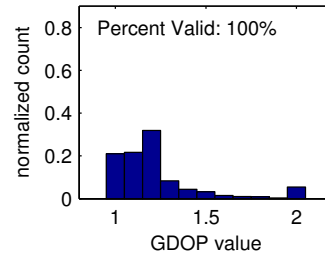
(a) Line-of-sight ranging on an “empty” floorplan



(b) non Line-of-sight ranging on an “empty” floorplan



(c) Line-of-sight ranging on the original floorplan



(d) non Line-of-sight ranging on the original floorplan

Fig. 4. Histogram representing GDOP values of a single experiment run. The x-axis shows the GDOP bin values and y-axis shows the normalized bin counts.

meant to model ranging radio, line-of-sight is not taken into account. The positions of the robots are assumed to be known for these tests. We set the particle filter to use 100 particles, which provides a reasonable loop speed.

Table 1. Characteristics of three simulated noise levels and the average Euclidean errors for various simulated tracking algorithms

	Simulated sensor variance (m^2)	Dropped meas. (%)	Avg PF error (m)	Avg map PF error (m)
Low noise	0.5	25	0.2429	0.2295
Moderate noise	1.0	50	0.4163	0.3982
High noise	1.5	75	0.7461	0.6720

In addition, Table 1 also gives the average Euclidean error over the entire run (approximately 10,000 estimates) for all runs. Using prior knowledge from the map improved the results, and this improvement was more pronounced at higher noise levels. Fig. 5 gives sample histograms of the Euclidean error for the tracking runs.

Additional tests (not presented) were performed using a Pioneer robot and sonar sensors to verify the range-only tracking algorithms on actual hardware. These experiments confirmed the simulated results.

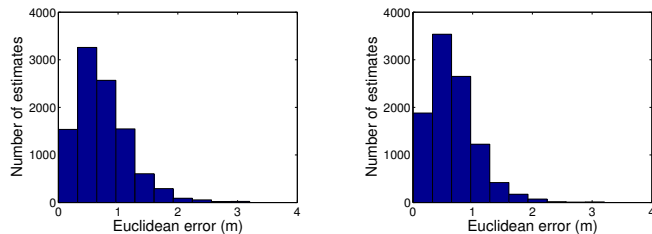


Fig. 5. Sample histograms of Euclidean error at high noise level for simulated tracking using particle filter without map (left), and particle filter augmented with the map (right)

6 Conclusion/Future Work

Our results show that a particle filter provides feasible tracking using noisy, range-only sensor data. The filter is robust enough to avoid divergence due to high sensor variances, dropped measurements, and sensor silences. It has been revealed that using prior knowledge from the map has the potential to further increase the tracking accuracy of the particle filter. However, if robots are operating in scenarios in which the map is not known (or only partially known), the particle filter without the map provides acceptable tracking performance.

To improve tracking accuracy, a cooperative control algorithm was developed for a robot team to minimize the position uncertainty of the target. Using a decentralized cost map approach, the cooperative control algorithm locates the lowest cost position for the robot and then finds the lowest cost path to reach that position. Configurations were evaluated using the GDOP metric, a measure of how team configuration will magnify sensor error. Using experiments running the cooperative control algorithm in the Player/Stage simulation environment, the effect of LOS vs. non-LOS sensors on team performance were compared. In a cluttered environment, the use of non-LOS sensors improved performance by giving the robot team the ability to form preferred configurations.

For future research, we plan to examine alternative methods for utilizing prior knowledge from the map. Algorithms like Hidden Markov Models have the potential to more accurately model the tendencies of human first responders in environments with maps. Such an approach would also allow robots to learn motion patterns from human movement data.

To deal with real scenarios, robot failure detection could improve the cooperative control algorithm’s robustness and flexibility. Robot failures are important to detect because in hazardous environments, malfunctions are common. For robustness, a team should be able to detect if a robot has stopped or is out of range and compensate for the loss in its behavior.

Acknowledgements

We thank Bradley Hamner for sharing his hardware and software optimization expertise.

References

1. P. Dana. Global Positioning System Overview, The Geographer's Craft Project, Department of Geography, The University of Colorado at Boulder, <http://www.colorado.edu/geography/gcraft/notes/gps/gps.html>.
2. J. Djugash, S. Singh, and G. Kantor. Range-only slam for robots operating cooperatively with sensor networks. In *Proceedings, IEEE International Conference on Robotics and Automation*, 2006.
3. R. Fierro, P. Song, A. Das, and V. Kumar. *Cooperative Control and Optimization*, volume 66, chapter Cooperative control of robot formations, pages 73–93. Kluwer Academic Press, 2002.
4. B. Gerkey, R. Vaughan, and A. Howard. The Player/Stage project: Tools for multi-robot and distributed sensor systems. In *Proceedings of the International Conference on Advanced Robotics (ICAR 2003), Coimbra, Portugal, June 30 - July 3, 2003*, pages 317–323, 2003.
5. B. Gerkey, R. Vaughan, A. Howard, and N.Koenig. The Player/Stage Project, <http://playerstage.sourceforge.net/>.
6. A. Howard, L. Parker, and G. Sukatme. Experiments with a large heterogeneous mobile robot team. *Intl. Journal of Robotics Research.*, 2005.
7. G. Kantor and S. Singh. Preliminary results in range-only localization and mapping. In *Proc. of IEEE Conf. on Robotics and Automation.*, 2002.
8. A. Kelly. Precision dilution in triangulation based mobile robot position estimation. In *Intelligent Autonomous Systems*, 2003.
9. V Kumar, D. Rus, and S. Singh. Robot and sensor networks for first responders. *Pervasive Computing*, pages 24–33, 2004.
10. S. M. LaValle. *Planning Algorithms*. Cambridge University Press (also available at <http://msl.cs.uiuc.edu/planning/>). To be published in 2006.
11. E. Liao. Cooperative control of a multi-robot team in a firefighting scenario. Technical Report CMU-RI-TR-05-50, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, in progress.
12. L. Liao, D. Fox, J. Hightower, H. Kautz, and D. Schulz. Voronoi tracking: Location estimation using sparse and noisy sensor data. In *Proc. Intl. Conf. on Intelligent Robots and Systems*, 2003.
13. D. Naffin and G. Sukhatme. Negotiated formations. In *Proceedings of the Eighth Conference on Intelligent Autonomous Systems*, pages 181–190, 2004.
14. S Min Oh, S. Tariq, S. Walker, and F. Dellaert. Map-based priors for localization. In *Proc. Intl. Conf. on Intelligent Robots and Systems.*, 2004.
15. R. Saber, W. Dunbar, and R. Murray. Cooperative control of multi-vehicle systems using cost graphs and optimization. In *Proc of the American Control Conference, June 2003.*, 2003.
16. A. Stroupe. *Collaborative Execution of Exploration and Tracking Using Move Value Estimation for Robot Teams (MVERT)*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, September 2003.
17. I. Ulrich and J.Borenstein. VFH+: Reliable obstacle avoidance for fast mobile robots. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation.*, pages 1572–1577, 1998.