
Coordinated Search in Cluttered Environments Using Range from Multiple Robots

Geoffrey Hollinger, Joseph Djughash, and Sanjiv Singh

Carnegie Mellon University, Pittsburgh, PA USA
{gholling@andrew.cmu.edu, robojoe@cmu.edu, ssingh@ri.cmu.edu}

Summary. In this paper, we describe real-time methods for incorporating non-line-of-sight range measurements into a framework for finding a non-adversarial target in cluttered environments using multiple robotic searchers. We extend previous coordinated search strategies to utilize information from noisy non-line-of-sight range measurements. Sensors using ultra-wideband radio are becoming available that provide range measurements to targets even when they are occluded. We present two Bayesian methods for updating the expected location of a mobile target and integrating these updates into planning. We present simulated results in a complex museum environment as well as on mobile robots. Our results show the success of our algorithms at utilizing information from measurements in a coordinated search framework.

1 Introduction

The problem of coordinating multiple robots to search for a non-adversarial target is relevant to many applications throughout robotics. Urban search and rescue teams in military and first response scenarios often need to search for both moving and stationary targets in cluttered environments. The major application that has inspired our work is the assistance of first responders in disaster scenarios. Previous work by Liao et al. has examined the potential for range-only measurements to track a mobile first responder in disaster scenarios [8]. As in this previous work, we assume that non-line-of-sight range measurements to a target are available. Ultra-wideband radios are currently in development that can provide range measurements through walls using a synchronized time-of-arrival method [9]. The availability of these sensors motivates the development of algorithms to utilize their potential.

This paper describes real-time methods for incorporating noisy non-line-of-sight sensor measurements of range to a target into a probabilistic multi-robot coordinated search framework. Applications to first response assistance have led us to make the assumption that the target is *non-adversarial*. This assumption is consistent with the task of locating a moving first responder

or survivor in a disaster scenario. We discuss our underlying framework and then outline two Bayesian methods for utilizing information from sensors. This paper is an extension of previous work by Hollinger, Kehagias, and Singh on locating a non-adversarial target in the minimum expected time [6]. The novelties of this paper include the formulation and testing of the measurement update methods, coordinated search results on the museum environment, and the implementation of this work on mobile robots.

This paper is organized as follows. Section 2 discusses related work in the areas of tracking, search, and pursuit-evasion. Section 3 describes the problem of finding a non-adversarial target with multiple searchers. Section 4 presents our algorithms for coordinated search and incorporating measurements. Section 5 gives both simulated results and results on mobile robots. Finally, Section 6 draws conclusions and discusses future directions for this work.

2 Related Work

Research in the areas of tracking, pursuit-evasion, and multi-robot coordination are all relevant to the coordinated search problem. The potential for multi-robot assistance in first response scenarios is examined at a high level by Kumar, Rus, and Singh [7]. They propose a general system for autonomous robotic assistance for first responders. More specific solutions to the first response assistance problem have been formulated by Djugash et al. [2] and later by Liao et al. [8]. These works utilize an ad hoc network of range-only sensors to track first responders, map the locations of new nodes in the network, and coordinate mobile nodes to minimize uncertainty.

Recent work in tracking mobile targets is also relevant to the coordinated search problem. Ferris, Hahnel, and Fox develop an algorithm for tracking a mobile target using wireless signal strength [3]. Their algorithm discretizes the environment in a similar manner to ours and then uses a particle filter to estimate position. We provide an alternative method for position estimation that closely couples with planning and remains computationally tractable. Our method also does not require prior training data, and it is completely online. As a tradeoff, our method is less appropriate than theirs for sensors with very complex or unknown models.

Coordinated search is closely related to work in the field of pursuit-evasion. One of the first algorithms for capturing an adversarial evader in an indoor environment was developed by Guibas et al. [5]. Their algorithm discretizes a polygonal environment based on visibility boundaries and searches the resulting information space for a path that guarantees capture. The problem studied by Guibas and LaValle differs from the one we discuss in this paper in that it assumes an adversarial evader. Sarmiento et al. also directly studied the coordinated search problem by developing strategies for searching indoor environments for a stationary target using a cost heuristic [10]. In previous

work, we developed a heuristic graph search solution to the coordinated search problem for a mobile, non-adversarial evader [6].

3 Problem Definition

To describe the state of the world in a coordinated search scenario, we must develop representations for the environment and the locations of the searchers and targets. Assume that the state of a searcher at time t is known to be $X^P(t)$, and the state of a target at any time t is known with a certain probability to be $X^E(t)$. In a polygonal environment, this problem is continuous. To collapse the problem into a discrete domain, discretize the environment into a number of cells, $1 \dots N$. The state of the i^{th} searcher is now fully defined by $X_i^P(t) = n$ where n is the cell in which the searcher is currently located. Now, define a capture event at time t as the occurrence of $X_i^P(t) = X^E(t)$ for any searcher i . To define the state of the target, let p be a row vector such that $p = [p_0, \dots, p_N]$ where each value represents the probability that the target is in the corresponding cell. Let the value p_0 represent the probability that the target has already been found by the searchers. Refer to this as the “capture state.” The vector p now defines a probability distribution over the target’s position in the environment (with the addition of a capture state).

The searchers’ goal is to minimize the expected time of reaching a capture event. Thus, the searchers seek to maximize the probability that the target is in the capture state at any given time t . The coordination problem is then defined as the determination of paths for the searchers such that the probability of capture is maximized at any given time. The problem of incorporating measurements becomes one of modifying the probability vector p to take into account new information from sensors.

4 Algorithm Description

4.1 Planning and Coordination

To reduce the task of finding a mobile target in a polygonal environment from a continuous to a discrete planning problem, we discretize a polygonal map into convex cells to form an undirected graph. Figure 1 shows an example discretization for a museum environment. The museum environment is challenging for coordinated search because it contains many cycles, which give a moving target many paths to avoid detection.

To integrate a motion model of the target into our pursuit-evasion framework and better define capture events, we use “capture” and “dispersion” matrices to modify the target’s state vector. We can represent a capture event by defining a matrix that moves probability from the searcher’s current cell to the capture state. Similarly, we can define dispersion matrices to represent the




Fig. 1. Museum map used for simulated coordinated search trials. Gray lines show discretization boundaries. The museum map has many cycles by which moving targets can escape detection.

expected motion of the target in the environment. Concatenating the capture and dispersion matrices yields a forward model for the target’s state vector at each time step.

Since the expected time of capture of the target cannot be known exactly during planning, we use a heuristic cost function to guide coordination of robotic searchers. The entropy of the target’s estimated state vector provides a heuristic that directly relates to expected time of capture. Adding a summation to the cost function over a path allows the searchers to look several cells ahead to determine the optimal search path. When planning using cost functions, locations in the environment must be assigned as replanning points. These locations represent the points in the environment between which the robots move. The most obvious replanning points to use are the centroids of cells because they provide quick access to adjacent cells.

Once a cost function has been defined, coordination strategies for robotic searchers can be found by enumerating possible paths on the discretized floor plan. To develop a decoupled planning algorithm, allow each searcher to enumerate all possible paths on the graph to a given depth while assuming that the states of the other searchers are fixed. This prevents the search space from growing in complexity with the number of searchers. The searchers can passively coordinate with each other by communicating their current paths. Each searcher can then plan for its own path while assuming that the other searchers will adhere to theirs. The coordination algorithm presented in this section is given in much greater detail in previous work [6].

4.2 Incorporation of Measurements

To incorporate noisy, sporadic measurements into our coordinated search framework, we propose two methods for modifying the target’s state probability vector given new sensor data. Both methods can be seen as estimating a posterior distribution given a prior distribution on the target’s state. Thus, they both fit into a Bayesian probabilistic framework. For all cells x_i , we seek

to estimate the posterior distribution $P(x_i^t) = P(x_i^t|z, x_i^{t-1})$ given a measurement z and a prior $P(x_i^{t-1})$. The prior is an estimate of the evader’s position after the application of the dispersion matrix, $P(x_i^{t-1}) = P(x_i^t|x_i^{t-1})P(x_i^{t-1})$. Once we calculate a posterior, we can renormalize and define the target’s state probability vector as $p_t = [P(x_c^t), P(x_1^t), \dots, P(x_n^t)]$, where $P(x_c^t)$ is the probability that evader has already been captured (the capture state). The value of $P(x_c^t)$ remains constant since we assume that noisy measurements cannot capture the target.

Sampling Method

Our first method for incorporating sensor data draws samples from a measurement distribution and uses these samples to directly estimate the posterior distribution on the target’s state. Assume that we receive a measurement z with an arbitrary noise model. We can draw n samples from this distribution to represent the “true” measurement.¹ In Euclidean space, each sample represents an (x, y) position of the target, which corresponds to a cell on the map discretization. We then weight each sample according to the prior on the cell in which it falls. This allows for the calculation of the posterior distribution on each cell as in Equation 1.

$$P(x_i^t|z, x_i^{t-1}) = \eta P(x_i^{t-1}) n_i, \quad (1)$$

where n_i is the number of samples that fall in a cell x_i , and η is a normalizing constant.

In this paper, we examine the incorporation of range measurements into the coordinated search framework. A measurement z is a single range value, and we assume that its noise is Gaussian with variance σ^2 . We can draw sample range values from this Gaussian, and these values map to circles in Euclidean space. To map to (x, y) coordinates, we use Bresenham’s circle drawing algorithm [1] on a finely discretized grid and count each hit as a sample. This takes full advantage of the sampled range values and greatly decreases the number of samples necessary to properly estimate the posterior.

Centroid Method

Our second method seeks to improve computational efficiency by calculating the likelihood of a measurement given the target is in each cell. Using Bayes’s rule, the posterior can be rewritten as in Equation 2.

$$P(x_i^t|z, x_i^{t-1}) = \eta P(z|x_i^t) P(x_i^{t-1}), \quad (2)$$

¹ There are other methods for estimating the posterior through sampling, i.e. maintaining particle estimates of the target’s position. Our method is particularly appealing for range measurements because it allows us to sample a 1D space.

where η is a normalizing constant.

This reduces the problem of calculating the posterior to that of calculating a likelihood term $P(z|x_i^t)$. Since each cell is represented as a continuous set of points in the map plane, this calculation is difficult. To reduce the complexity of the problem, we make the assumption that the target will remain close to the centroid of each cell.² Thus, each cell can be reduced to its centroid. Since cells are convex, the centroid is guaranteed to be inside the cell. The calculation of $P(z|x_i^t)$ is now one of calculating a likelihood at a point. This is a common problem in state estimation, and is prevalent throughout Monte Carlo methods and particle filter literature [3].

For range measurements, the problem of calculating $P(z|x_i^t)$ is simply that of determining the expected range value for each cell centroid. Let r_i^c be the range from a ranging sensor to the centroid of cell x_i and let r be the received range measurement with assumed Gaussian noise variance σ^2 . The likelihood is then calculated as in Equation 3.

$$P(z|x_i^t) = N(r; r_i^c, \sigma^2) \tag{3}$$

Comparison of Methods

The sampling and centroid methods each have strengths and weaknesses in different environments. The sampling method biases towards larger cells since those cells will contain more sampled points. In the case of a moving target, the target will spend more time in larger cells, thus warranting this bias. A drawback to the sampling method is that it is computationally expensive if the number of samples grows large. Thus, to maintain real-time measurement incorporation, the number of samples must be kept small. The sampling method is particularly appropriate for small environments (requiring few samples) with highly irregular cells. Such cells are difficult to approximate simply by their centroid.

The major disadvantage of the centroid method is that it makes the strong assumption that the target is near the centroid of the cell. To account for this assumption, a value for σ must be used that is greater than the underlying noise model of the sensor. This degrades the accuracy of the sensor model. Unlike the sampling method, the centroid method does not bias towards larger cells. The major advantage of the centroid method is that it is extremely fast, only requiring one calculation for each cell in the environment. Thus, the centroid method is appropriate for large environments with regularly shaped cells.

² Another technique for generating likelihoods is to finely discretize the environment and calculate a likelihood at each point. This is equivalent to the histogram filter in the literature. The centroid method can be seen as a special case of the histogram filter.

5 Results

5.1 Simulated Results

We implemented our algorithms on a multi-agent coordinated search simulation in C++ on a 3.2 GHz Pentium 4 processor [6]. This simulation allows for multiple searchers, realistic modeling of ranging measurements, and both stationary and moving targets. For moving trials, we assumed that the average speed of the target is 1 m/s , and that it moves holonomically between cell boundaries. We tested our algorithms with varying numbers of searchers on a museum map (shown in Fig. 1). The searchers start in the same location for all trials, and the location of the target is initialized at random on the map. Each searcher holds a mobile ranging sensor, and a stationary ranging sensor is placed at the top left corner of the map. The stationary sensor corresponds to one placed outside of a building and helps to disambiguate target position uncertainty caused by symmetry. The simulated ranging sensors were given a variance of 2 m^2 and were given a 10% chance of providing a measurement each second.

Fig. 2 shows average capture times for moving and stationary targets with varying numbers of searchers. The results show significant improvement using measurements with both the centroid and the sampling methods. Additionally, the centroid method outperformed the sampling method. This is likely due to the limitations on the number of samples necessary to maintain real-time operation (500 samples were used for each measurement). Furthermore, the sampling method's bias towards large cells is detrimental when the target is stationary. Conversely, the centroid method's assumption that the target remains close to the centroid of the cell is valid for moving targets in small cells, which is most often the case on this map.




Fig. 2. Simulated coordinated search results in a museum comparing two methods for incorporating range measurements. Graphs show average capture times versus number of searchers for a stationary target (left) and a moving target (right). Error bars are one SEM, and averages are over 200 trials. Graphs are zoomed in for comparison.

5.2 Robot Results

To further confirm our algorithms, we implemented them on a mobile robot. We used a single Pioneer One robot from ActivMedia, equipped with a laser range scanner, as a searcher and a Nomad Scout as a mobile target. We utilized the Player/Stage software to control the robots [4]. The Pioneer searcher used the Adaptive Monte Carlo Localization (AMCL) algorithm packaged with Player/Stage to localize within the known environment. The Nomad Scout utilized a pseudo-random walk algorithm with sonar sensors.

For experiments that required the incorporation of measurements to the target, a *Parrot* ranging sensor was used to acquire the measurements. The *Parrot* sensors form a ranging sensor network that uses radio to communicate data and ultrasonic signals to determine range between nodes [11]. A total of three Parrots were used in our experiments. The first was mounted on top of the searcher, the second was placed on top of the target, and the third was placed at a surveyed location near a corner of the map.

We conducted experiments within a $20\text{ ft} \times 17\text{ ft}$ room, which we partitioned as in Fig. 3 using 1 ft cardboard walls. Since the test environment contains a cycle, it is impossible to guarantee capture of a quickly moving target with a single searcher [5]. We placed both mobile and stationary Parrots at a height of approximately 1.5 ft above the floor, making them taller than the cardboard walls. By placing the Parrots above the walls, we were able to convert the otherwise line-of-sight restricted Parrot sensors to “non-line-of-sight” ranging sensors. This simulates the range measurements that one could receive with ultra-wideband radio ranging sensors.




Fig. 3. Simple map used for coordinated search trials on Pioneer robot platforms (left) and corresponding photograph of robots in test environment (right). Gray lines show discretization boundaries, ‘R’ shows position of robotic searcher, ‘T’ shows position of target, and ‘M’ shows position of stationary measurement beacon. Blue and red traces show paths of searcher and target respectively (reconstructed from AMCL output). The robots are in the same locations on the map as they are in the photograph.

In ultrasonic-based ranging systems, there are two types of noises that affect range estimation, multipath reflections and random noise. For the Parrots, the latter is easily modeled as a zero-mean gaussian with a variance

of 0.5 m^2 . In order to detect and remove the multipath measurements from processing, we added a chi-squared test to our algorithm. This helps to filter out erroneous measurements caused by multipath.

Fig. 4 gives average times to capture for trials using the Pioneer and Nomad over a total of 144 runs. The robot results confirm the simulated results by showing that the centroid and sampling methods outperform the case without sensor data. They also show an improvement of the sampling method over the centroid method for moving targets and the reverse for stationary targets. This is likely due to the tendency for moving targets to remain in large cells and the sampling method’s bias towards larger cells. Qualitatively, the searcher moved more quickly towards the target’s cell while incorporating measurements. This demonstrates the success of our methods on a mobile robot platform.




Fig. 4. Coordinated search results on a Pioneer robot comparing two methods for incorporating range measurements. Graphs show average capture times for a single searcher finding a stationary target (left) and a moving target (right). Error bars are one SEM, and averages are over 24 trials. Graphs are zoomed in for comparison.

6 Conclusions and Future Work

We have presented two techniques for incorporating non-line-of-sight range measurements into a coordinated search framework. Both techniques utilize a Bayesian framework for estimating the posterior distribution of a target’s position given a prior estimate and a new measurement. Our techniques generalize to most sensors with a known noise model, and they scale well with increasing searchers. We make the assumption that a target is non-adversarial and present a heuristic graph search approach to minimizing the expected time to capture. We have presented simulated results showing that our methods remain tractable in a large environment with cycles. Our simulated results also show that both of our techniques greatly outperform the case without measurements. Additionally, we have implemented our algorithm on a Pioneer

robot utilizing sonar ranging sensors. The results on the Pioneer confirm the simulated results and demonstrate the real-time performance of our methods.

A major extension of this work is the implementation of our coordinated search algorithms on multiple robots in large-scale environments. We plan to move out of the controlled laboratory environment and into an office building similar to map used in simulation. We will also use ultra-wideband ranging radios to provide non-line-of-sight sensing in this environment. Our algorithms provide an important step towards fully utilizing the new technology of ultra-wideband ranging radios.

Acknowledgements

We thank Athanasios Kehagias for his insightful comments. This work is supported by the National Science Foundation under Grant No. IIS-0426945.

References

1. J.E. Bresenham. Algorithm for computer control of a digital plotter. *IBM Systems Journal*, 4(1), 1965.
2. J. Djugash, S. Singh, and P. Corke. Further results with localization and mapping using range from radio. In *Proc. 5th Int'l Conf. on Field and Service Robotics*, 2005.
3. B. Ferris, D. Hahnel, and D. Fox. Gaussian processes for signal strength-based location estimation. In *Proc. of Robotics Science and Systems*, 2006.
4. B. Gerkey, R. Vaughan, and A. Howard. The Player/Stage project: Tools for multi-robot and distributed sensor systems. In *Proc. Int'l Conf. on Advanced Robotics*, 2003.
5. L. Guibas, J. Latombe, S. LaValle, D. Lin, and R. Motwani. Visibility-based pursuit-evasion in a polygonal environment. *Int'l Journal of Computational Geometry and Applications*, 9(5):471–494, 1999.
6. G. Hollinger, A. Kehagias, and S. Singh. Probabilistic strategies for pursuit in cluttered environments with multiple robots. In *Proc. Int'l Conf. on Robotics and Automation*, 2007.
7. V Kumar, D. Rus, and S. Singh. Robot and sensor networks for first responders. *Pervasive Computing*, pages 24–33, 2004.
8. E. Liao, G. Hollinger, J. Djugash, and S. Singh. Preliminary results in tracking mobile targets using range sensors from multiple robots. In *Distributed Autonomous Robotic Systems*, 2006.
9. R. Saeed, S. Khatun, B.M. Ali, and M. Khazani. Performance of ultra-wideband time-of-arrival estimation enhanced with synchronization scheme. *ECTI Trans. on Electrical Eng., Electronics and Communication*, 4(1), 2006.
10. A. Sarmiento, R. Murrieta-Cid, and S. Hutchinson. A multi-robot strategy for rapidly searching a polygonal environment. In *Proc. 9th Ibero-American Conf. on AI*, 2004.
11. W. Zhang, J. Djugash, and S. Singh. Parrots: A range measuring sensor network. Technical Report CMU-RI-TR-06-05, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.