Behavior-Based Mapping and Tracking with Multi-Robot Teams Using Probabilistic Techniques

Ashley W. Stroupe[1] and Tucker Balch[2]
[1]Carnegie Mellon University, Pittsburgh, Pennsylvania
[2]Georgia Institute of Technology, Atlanta, Georgia

Abstract—We present a behavior-based approach for directing the movements of robot teams engaged in mapping and tracking target objects in their environment. The resulting paths optimize the vantage points for all the robots on the team – therefore maximizing their information gain. At each step, each robot selects a direction for movement that will maximize the utility of its next observation. The global team behavior serves to optimize the team's knowledge because each robot considers the observational contributions of teammates. The utility of an observation is evaluated as the reduction of target uncertainty resulting from it. The approach is evaluated in simulation with static and moving targets. We evaluate performance by measuring the resulting expected uncertainty about target object locations compared to that of robots acting without regard to teammate locations and to that of global optimization over all robots for each single step. The qualitative behavior of the team is sensible and close to the single-step optimal set of trajectories.

I. Introduction

In this work we address the problem of moving a team of robots so as to discover the locations of identifiable targets. The targets may be stationary or moving. Furthermore, the robots may be homogeneous, or heterogeneous with sensor and movement constraints. The algorithm can take advantage of prior, uncertain knowledge of the locations of the targets, but it does not require this information. Also, we do not assume that the robot team knows how many targets there are in advance.

Our approach to this problem is a behavior-based solution where each robot makes movement decisions based on its own sensor information. Behavior-based robotic systems typically have the advantages of speed and reactivity, which make them desirable for high-speed operation and dynamic environments. However, direct mapping between current state and action often means that full advantage cannot be taken of robot teams.

In this work we use a behavior-based framework to perform landmark mapping and tracking using teams of robots. At each step, robots move in order to most improve their collective information about the targets according to a predefined value function.

We refer to our approach as behavior-based because each robot independently determines in which direction to travel based on the current situation. The approach is not purely reactive because robots maintain knowledge of tracked target position uncertainty in a covariance matrix. This matrix implicitly reflects the prior positions and prior measurements of teammates. Locations and observations by other robots may be communicated, but could also be inferred if teammates are detectable and their sensor models are known. When communication is absent, the full advantage of the team cannot be taken until information is later combined.

For the tasks of mapping and tracking, value must be related to uncertainty in target object locations. High levels of uncertainty will provide lower value than low levels of uncertainty. Different representations of uncertainty may be used, depending on the task. Other types of criteria could be considered in this framework by including them in the value function. For example, priority of individual targets, distance strayed from an otherwise desirable trajectory, and requirements for other tasks could also be included in the evaluation.

In this paper we consider only the issue of movement to optimize observations. However, we envision the algorithm as part of a complete motor schema-based navigational system [1] that simultaneously accomplishes an overall mission, for which target mapping and tracking is but one component. In the context of the motor schema paradigm, the algorithm generates a vector representing the best direction for the robot to move in for an optimal observation, which is in turn integrated with the outputs of other motor schemas in the system.

II. Background and Related Work

There are several areas of previous work related to this research. Many years of investigation into mapping with robots have provided approaches for single robots and robot teams, with a few typical examples given here. Approaches to mapping include building free space/occupancy cell-based maps (primarily for indoor environments) [2][3][11], cell-based traversability maps [6] and landmark mapping (typically Kalman-Bucy filter) [5][8], both for outdoor environments. Most mapping work has focused on using coverage patterns to completely explore a space by one or more robots. In the multi-robot case, areas are typically broken up into sub-areas, which are each covered by a robot; robot sub-maps are later combined into a single map [3][5][8][11]. A second approach is to move some robots while others remain fixed as landmarks [7][10].

The task explored in this research is somewhat different than the coverage task – in particular, we are interested in mapping and tracking targets rather than covering or exploring. Coverage/exploration is appropriate in indoor structured environments when occupancy/free-space maps are required for navigation. Full coverage, however, may not be required for landmark mapping or for measuring specific targets in an otherwise known environment. In this case, the task becomes determining where robots should travel to best measure the landmarks or targets of interest in continuous space as quickly as possible.

Dynamic objects have also been considered in several applications. Several groups have looked at robot placement to keep objects under surveillance by robots [9][12]. Kalman-Bucy filters may be used to locate and track a dynamic target from multiple robots [4]. The instantaneous position of a moving target can be refined by combining simultaneous estimates from multiple platforms [15]. Next Best View predicting where a single robot should to go best improve the quality of a static surface model of an object [13]. Predicting where a team should go to best observe a group of moving objects has also been studied. This approach approximates the distributions from future measurements with particles and attempts to minimize the spread of the particles.

III. APPROACH

Our approach to mapping and tracking has each robot team member choose the best move at each step given the current situation. This is a distributed algorithm because each robot makes its own movement decisions. For a robot, the best move is the one that maximizes a given value function. In this case, value is associated with reducing uncertainty in target measurements. The algorithm considers a number of candidate moves and chooses the one with the greatest value (the number of candidate moves and size of move are parameters of the algorithm), given that teammates will be making additional observations near their current locations. Candidate moves that result in an overlap with a target, teammate, or obstacle are excluded, as are moves that violate holonomic constraints of the robot.

State information considered by the algorithm includes the current estimate of target positions, the uncertainty covariance of the target estimates among all teammates (communicated or inferred), and current positions of all robots (communicated or inferred). Target locations and uncertainties are represented probabilistically as two-dimensional Gaussian distributions in covariance matrices. The sensor models enable estimating the extent to which uncertainty can be reduced by a subsequent observation. If necessary to infer teammate contributions, sensor models are also employed to estimate them.

After a robot determines its best direction of movement, it moves a short distance in that direction. New sensor measurements are taken and the covariance matrices are updated. The new covariance is used as the prior in the next step. In the case of static targets, new measurements can be directly combined with previous estimates. For dynamic targets, new measurements can be incorporated by first updating the targets based on a motion model.

If robots are able to communicate, they share probabilistic representations of the observations they have made of the targets. Each robot then uses other robots' observations to update the probabilities of the target locations. Teammate positions are used to estimate the next set of measurements so that they can best be complemented. When communication is not possible but teammates positions are observable and sensor models are available, each robot can individually approximate the observations its teammates would make. It is these estimates of covariance that are used in the next movement step. However, to take full advantage of the team measurements, the individual series of measurements taken by each robot must be later recombined.

The algorithm for each robot is summarized as follows:

- Initialize: Take measurements and initialize positions and covariance with all shared (or estimated) measurements
- Loop until termination condition:
  - Predict approximate results of teammate measurements using sensor models and positions
  - Combine approximated teammate measurements with current covariance to predict contribution
  - Choose move that provides the best value, given predicted covariance
  - Move to new position
  - Take measurements; update covariance with all shared (or estimated) measurements

IV. VALUE MEASURES

For mapping and target tracking in the absence of other requirements, the goal is to obtain the best quality estimates of target locations. High quality estimates have low uncertainty. Four types of value functions were applied. In each, measures of uncertainty are negated so larger uncertainty produces lower value. For all value functions, the 1-σ (1 standard deviation) oval (reflecting location and orientation of the Gaussian p.d.f.) for each target is considered. The first value function, V_1, reflects total uncertainty for all targets by using the total area of 1-σ ovals (total area, units of distance-squared):

$$V_1 = \sum_{i=1:T} \pi\sigma_{i\,maj}\sigma_{i\,min}$$

T is the number of targets. V_2 is the largest uncertainty among all targets (greatest area, distance-squared units):

$$V_2 = \max_{i=1:T}\left(\pi\sigma_{i\,maj}\sigma_{i\,min}\right)$$

The third, V_3, is the total of the largest uncertainty of all targets (total σ_{maj}, units of distance):

$$V_3 = \sum_{i=1:T} \sigma_{i\ maj}$$

The fourth, V_4, is the largest largest uncertainty of all targets' (maximum σ_{maj}, units of distance):

$$V_4 = \max_{i=1:T}\left(\sigma_{i\ maj}\right)$$

The optimal location from which to take an observation depends significantly on the value function. Consider two targets that have been observed, and the covariance represented by the two-dimensional probability density function illustrated in Figure 1. The target at (0,0) has smaller standard deviations than the target at (2,2).

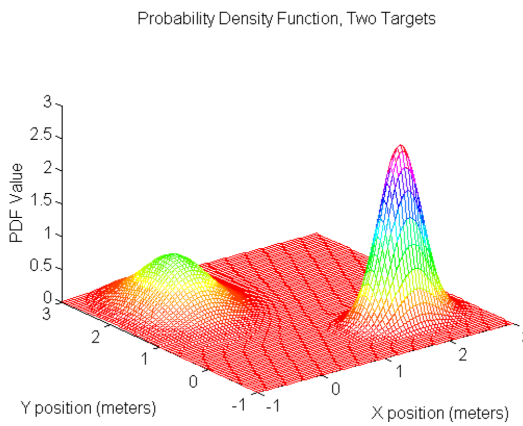

Figure 1. Sample probability density function showing knowledge of two targets, one more certain than the other.

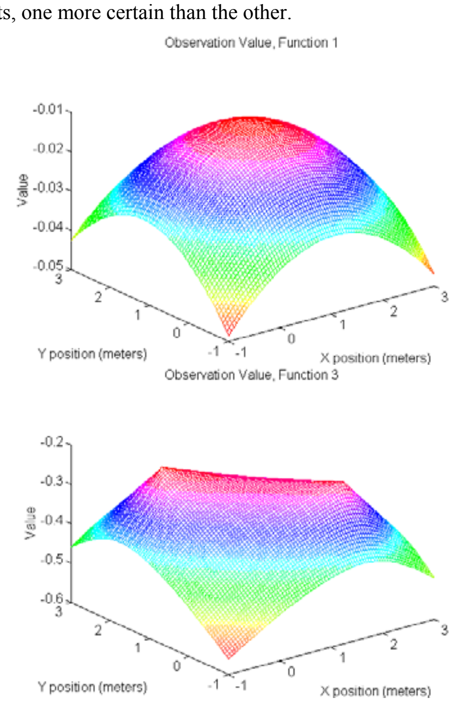

Figure 2. The value of taking observations at each point, given the above p.d.f. Top: V_1 (total area). Bottom: V_3 (total σ_{maj}).

The value of single measurements at points in the environment (V_1 and V_3) is shown in Figure 2. The V_1 maximum of -0.0106 occurs where area reduces equally: point (-1,1), where V_3=-0.019. V_3 is maximized at –0.245 where the observation is nearly perpendicular to both σ_{maj} axes: point (-1,0), where V_1=-0.2742. Both points provide observations aligned roughly perpendicular to the p.d.f.'s major axes and are closer to the less certain target.

V. Experimental Approach

A series of experiments was conducted in a *Matlab* simulation of Minnow robots (Figure 3) [15]. Robot movement decisions were made simultaneously based on the same state information. After each step, the resulting covariance that would result from a measurement by all robots is recorded to use as the covariance for the next step. Termination occurs when all robots prefer current locations to any further moves. An a priori map or an initial observation provides initial uncertainty and location on targets. Targets are assumed unique and identifiable; association is not addressed in this work.

Figure 3. Minnow Robot for multi-robot tasks, including soccer.

Several experiments were conducted using static targets. Numbers of robots and targets were varied, as were initial conditions and robot capabilities. Baseline performance included holonomic motion and omni-directional vision with a range of 30 meters. This range ensures full vision of environments (6x6 and 20x20 meters). Experiments with one and four targets used the small environment; those with ten used the large environment. The vision noise model has standard deviations of 10% in range in meters) and 0.5º in bearing. Simulated non-holonomic motion reduces maximum turning angle to 50º. Simulated limited sensing limits range to 2.5 m and angle to 50º. Robot step size (speed) was also varied. Angular resolution for candidate moves was 2º. Values were chosen based on Minnow performance. For most static target experiments, measurement noise is reduced to zero and maps are initialized with exact target locations but very large uncertainties (1000 m). This allows direct comparison of what is being optimized, the success in optimizing, and the quality of the paths generated by the approach. Comparison of resulting values in the presence of noise may be misleading, as differences may be due to measurement differences. Thus, comparisons are made in the noiseless case. In noisy measurement experiments, the noise added was drawn from Gaussian distributions with the sensor model parameters and the map is initialized with an observation from all robots.

Experiments were also conducted with dynamic targets. Robots again use current covariance and teammate positions to determine independently where to move and predicted covariance is recorded. After this, both robots and targets are moved. Prior to measurement, but after movement selection, target position uncertainty is grown to account for motion. New target positions and covariance matrices are obtained in the measurement step and combined with the previous estimate. Unseen targets increase uncertainty and are assumed static for simplicity; this could be replaced with a KBF prediction. Number of targets and robots were varied, as was speed of targets relative to robot speed. Robot capabilities are baseline (30 meter, omni-directional vision, holonomic motion). Value comparisons are made in the absence of noise.

VI. Experiments with Static Targets

The experiments with static targets were designed primarily to investigate the effects of: value function, team size, initial conditions, measurement noise, and robot capability limitations. Results are compared to robots using individual reasoning and to one-step optimal.

A. *Effects of Value*

Experiments used to evaluate the effects of value were:

- 1 to 4 robots with 1 target
- 1 to 5 robots with 4 targets
- 1 to 7 robots with 10 targets

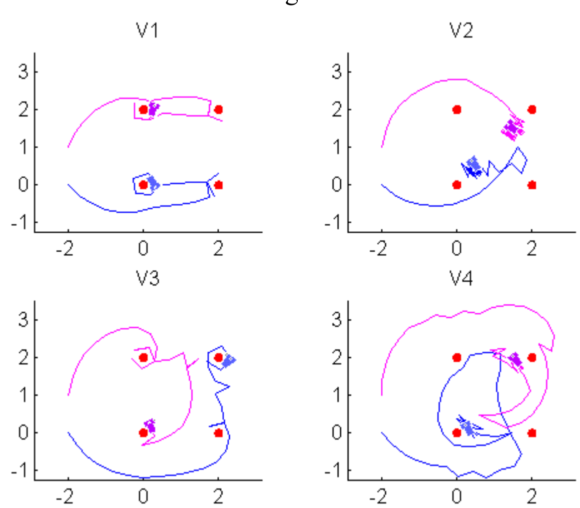

Figure 4. Paths generated by two robots using four value functions. Left plots (minimizing totals) concentrate close to targets while right plots (minimizing maximum) result in paths in the middle.

From the sample trajectories illustrated in Figure 4, one can see that total area and total σ_{maj} value functions (V_1 and V_3) cause robots to closely circle single targets sequentially, while maximum uncertainty value functions (V_2 and V_4) cause the robots to spend more time in the middle, taking equal measurements on all the targets.

B. *Effects of Team Size*

The experiments used to evaluate the effects of increasing team size were the same as those for evaluating value functions (previous section). As team size increases, robots tend to spread out less. This is likely due to several factors. More robots must cluster more in the same space. Also, with more measurements, the required angles between robots drop (three robots form a triangle to observe one target). Lastly, as the team moves, uncertainties on close targets drop quickly, making the farther targets higher priority; robots will concentrate on

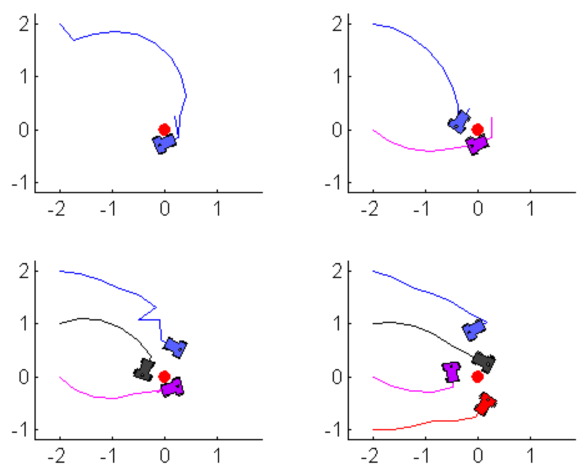

Figure 5. One to four robots observe a target (V_1). Multiple robots distribute around the target. One drives around the target to compensate.

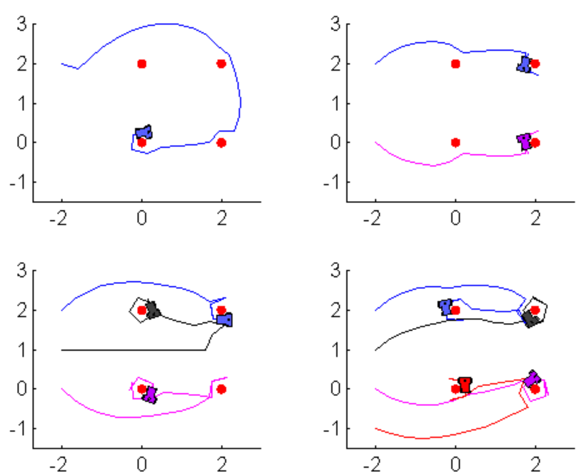

Figure 6. One to four robots observing four targets using V_1. One robot must visit all targets from multiple angles. Two robots arc apart to get perpendicular views. With more robots, they specialize.

these targets, even if multiple robots head toward the same target. Examples in Figures 5-7 optimize with V_1.

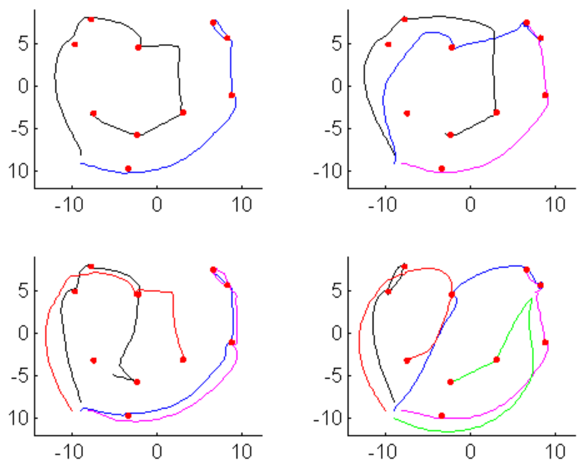

Figure 7. Two to five robots observing ten targets (V_1). With few robots, they must split up to examine different targets. With larger groups, they concentrate on far away (less well known) targets together.

C. *Effects of Initial Conditions*

Experiments run to evaluate initial condition impact were:

- 8 starting locations for 1 robot with 4 targets
- 6 starting locations for 2 robots with 4 targets
- 7 starting locations for 1 robot with 10 targets

Experiments used all value functions. Varying initial conditions changes the points reachable by each robot at each step. It also changes the initial relative points of view of the team members. The further apart robots are initially started, the more different the resulting trajectory will be. Resulting trajectories (Figure 8) are similar in that robots first obtain measurements at varied angles and then take tight measurements on targets as needed.

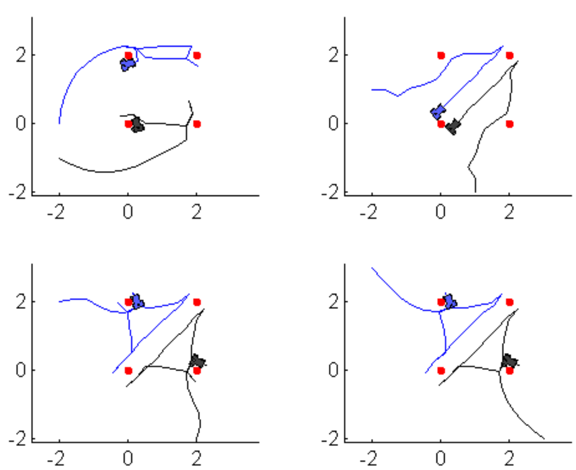

Figure 8. Two robots map four targets from different initial starting conditions optimizing V_1. Despite variety in actual paths, qualitatively they are similar in that robots first move to get complementary angles.

D. *Effects of Measurement Noise*

Experiments run to evaluate the effects of measurement noise, using V_1 and V_3, were:

- 2 to 5 robots with 4 targets
- 2 to 5 robots with 10 targets

Adding measurement noise demonstrates the approach's effectiveness in more realistic conditions (Figures 9, 10). The uncertainty model used to introduce noise into simulated sensor readings was based on Minnow Robots. Resulting trajectories are very similar to those without noise. Resulting values are slightly lower (uncertainty higher), since motions are not perfectly chosen with respect to actual target locations, especially early on before target locations converge to the correct position. Final target positions do converge on the correct values.

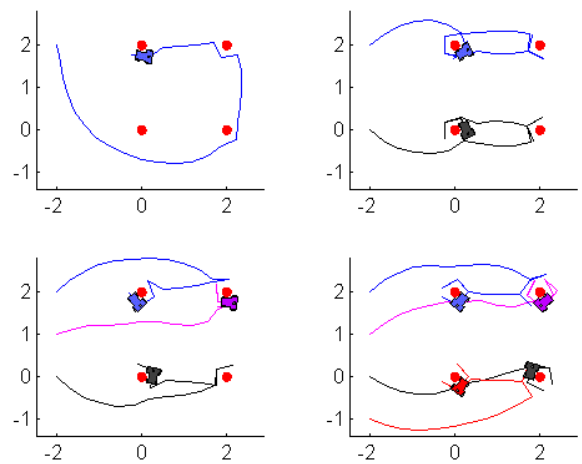

Figure 9. One to four robots observing four targets with added measurement noise using V_1. Trajectories are not quite as smooth, but present the same general shape as in the absence of noise (Figure 6).

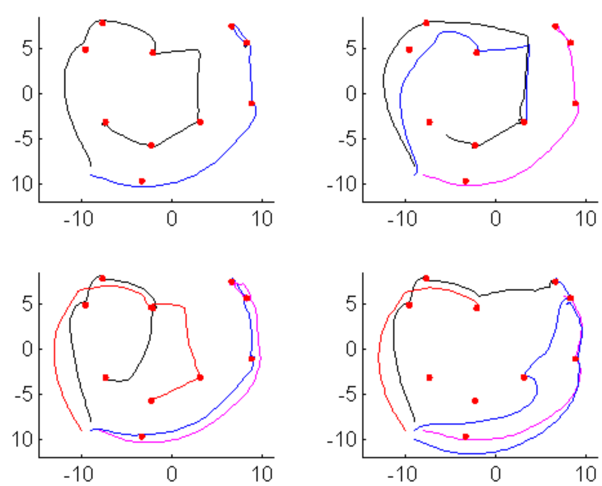

Figure 10. Two to five robots observing ten targets with measurement noise (V_1). Paths are similar to those without measurement noise.

E. Effects of Robot Capability Limitations

Experiments run to evaluate performance of robots with limited capabilities, using V_1 and V_3, were:

- 1 to 4 robots with limited visual range, limited visual angle, non-holonomic motion, and all limitations.
- 7 velocities (step sizes) for 2 robots with 4 targets

Examples of the effects of reducing robot performance are shown in Figure 11. Reducing the robot's turn angle causes robots to circle targets at a larger radius, but paths are qualitatively similar. Limited visual range forces robots to visit more closely to each target. Limiting visual angle forces robots to turn within an area to view all areas. When vision angle is limited, robots may fail to pass by targets; the improvement in targets still in visual range does not always outweigh the loss of observation of the passed target. To prevent early termination, robots are not allowed to consider the current location's value, forcing them out of the local optimum.

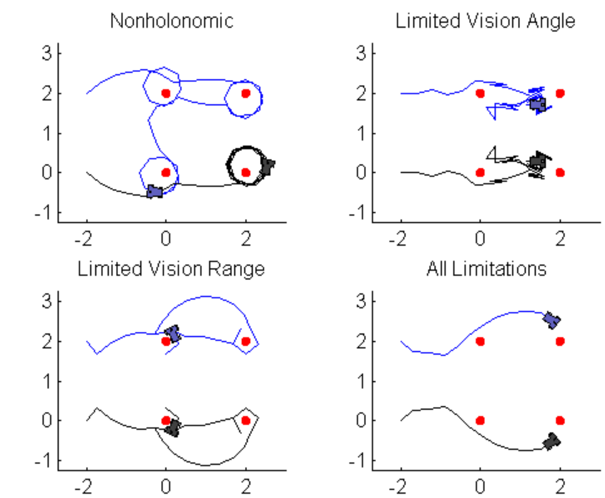

Figure 11. Two robots with various limitations observing four targets (V_1). Despite limitations in turn angle, vision angle, and vision range, robots find paths to achieve good measurements on all targets.

Changing step size also affects the results to some degree. Larger steps may force robots to not be able to reach optima but may allow robots to escape local minima. Typically, paths are similar across step sizes.

F. Comparative Performance

Performance was compared to that where robots use single-robot reasoning to choose moves (varied one-robot initial condition experiments in *C*). In the single-robot approach, at each step, robots use only their own observations to determine the next step. For quantitative comparison with this approach, after all robots have completed independent paths, measurements are combined and the resulting observational value is determined. Comparison can also be made with one-step optimal, in which all combinations of moves are investigated in order to find the most valuable *set* of actions. Due to the complexity of the optimal approach, only tests with small state spaces were conducted.

In comparison to the single-robot approach, this approach consistently results in better values (Figure 12). Results are compared after the same number of steps (20 for 4 targets, 50 for 10). Resulting values are shown for V_1, and results with V_3 show the same improvement. Qualitatively, robots spread out more, despite similar initial positions relative to the single-robot approach; the single-robot approach causes robots starting with similar positions to follow similar paths, reducing coverage. Single-robot approach performs closer to this approach when robots are intentionally initially distributed for coverage and point of view and when targets are clustered so many measurements can be taken. When robots are more naturally started near each other, resulting paths tend to be similar and do not take full advantage of the multi-robot aspects of the team by varying point of view.

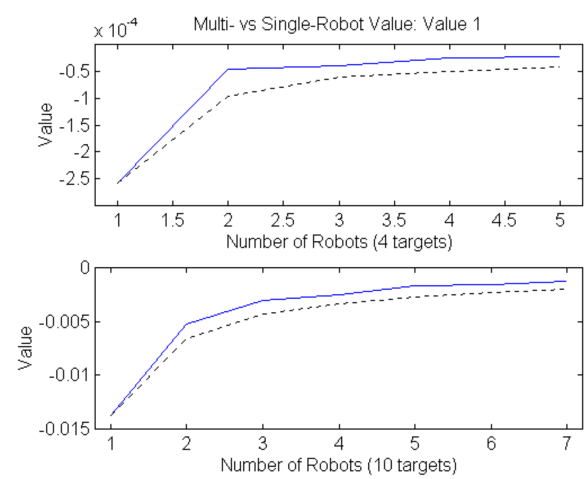

Figure 12. Comparison of value for this approach and the individual robot approach. Values improve using the multi-robot approach.

After many steps, giving the individual robots time to visit many targets, the resulting value of the multi-robot approach is still greater. The difference is much wider early, providing good values much more quickly.

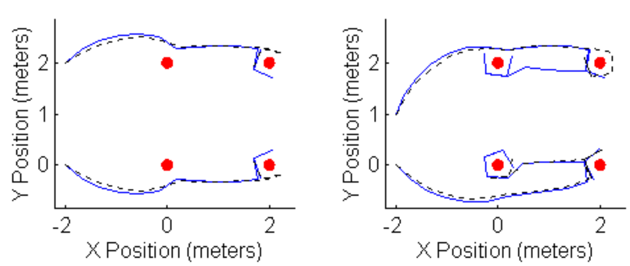

Figure 13. Comparison of one-step optimal trajectories (dotted) to this approach (solid). Trajectories are very similar, with this approach spreading robots slightly more since they don't account for simultaneous teammate movement.

In comparison to the one-step optimal, this approach causes robots spread slightly more, since simultaneous teammate movement is not considered, but trajectories and resulting values are quite similar (Figure 13). For example, after 10 steps in two experiments, this approach produces V_1 values of -2.87×10^{-4} and -3.77×10^{-4}, while the optimal produces values of -2.56×10^{-4} and -3.52×10^{-4}, respectively. The single-robot approach produces values of -4.71×10^{-4} and -0.001 for the same initial conditions.

VII. Dynamic Target Experiments

Several experiments were run using dynamic targets initialized with an a priori map and perfect measurements (for comparison) and most were run with measurement initialization and noisy measurements (for realism):

- 3 robots (one stationary) with 1 target moving linearly at a constant velocity slower than robots.
- 1 to 3 robots with 1 target moving linearly at a constant velocity slower than robots.
    - 4 starting locations for 1 robot
- 3 and 4 robots with 4 targets moving linearly away at a constant velocity slower than robots

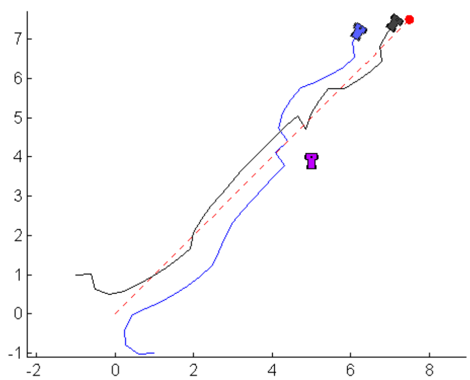

Figure 14. Two robots track a target using V_1. Robots stay on opposite sides of the target to get complementary views. As a third robot is approached, they reconfigure to a triangle.

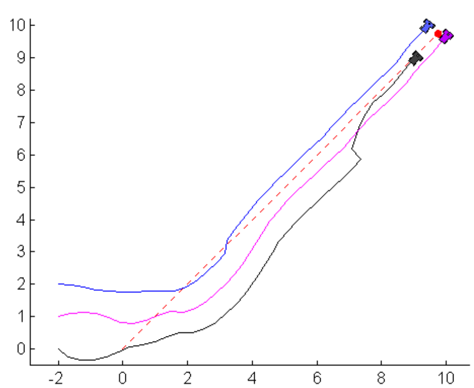

Figure 15. Three robots track a single target. First, they stay on either side, then they form a triangle to obtain complementary views

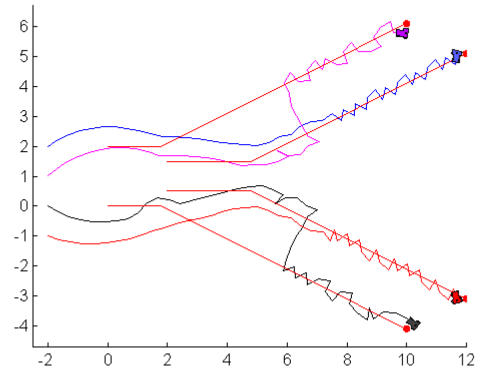

Figure 16. Four robots follow four targets. At first, while targets travel together, robots arc away to obtain complementary views. As the targets diverge, the robots specialize, each following one.

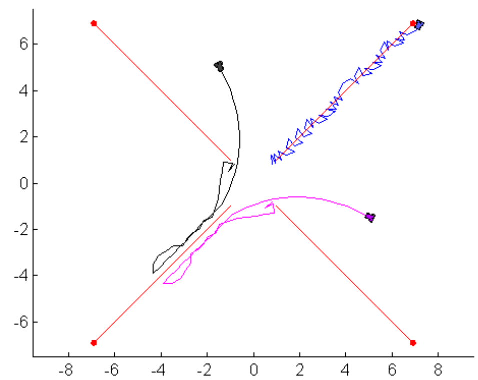

Figure 17. Four robots with limited vision range follow four divergent targets. When two lone targets move out of range, robots switch focus.

- 3 and 4 robots with 4 targets moving piece-wise linearly at velocities slower than robots.
- 1 robot with 1 target moving linearly at constant velocity (slower than, same as, and faster than robot).

Examples of the tracking behavior of the robots are shown in Figures 14-17. Targets moving together allow robots to arc around to vary the observation angle. As targets diverge, robots tend to specialize based on proximity. If targets move out of visual range, robots switch attention as needed. The oscillatory motion around target paths occurs when robots can move more quickly than the target. These paths allow single robots to obtain measurements with different axes to improve results. This behavior reduces and disappears when targets are moving faster than robots, as robots must use all forward velocity to keep the target in visual range. Trajectories resulting from noisy measurements, as in the static case, are very similar to those without noise.

Comparison shows improvement as in the static case. The value improves 14 times for two robots following a single target, and 11 times for three robots (Figure 14). Qualitative differences are evident: three collaborating robots form a triangle for view diversity; individuals follow similar path obtaining similar measurements.

VIII. Limitations And Future Work

Several additions to this work will be made in the near future. The approach will be implemented with unknown environments, allowing addition of new targets as they are observed. Next, the algorithm will be implemented in small teams of real robots. Lastly, the effects of more complex value functions will be investigated. Such value functions will include variable target priority (example in Figure 18) and deviation from pre-planned paths.

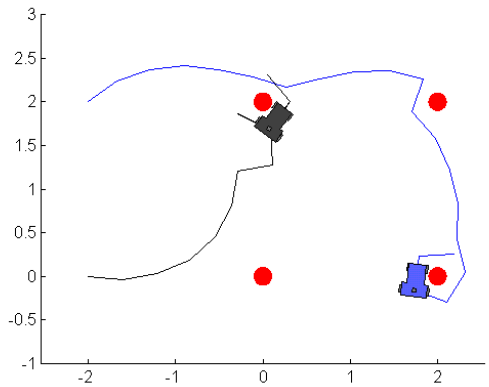

Figure 18. Two robots observe targets with varying priority. Top priority is top left, followed by top right, lower right, and lower left. Both robots immediately proceed to top priority target; after, one robot explores lower priority targets, while one remains at the top priority one.

IX. Discussion And Conclusions

This approach uses greedy search to determine the best action for each robot using current teammate positions and previous teammate measurements (estimated or known). The approach allows robot teams to improve use of a team without structured planning. The resulting combined measurements are of better value (defined by a value function) than are combined measurements of individually minded robots.

As a behavior-based approach with limited mathematical and state-size complexity, robots can quickly react to dynamic situations. Robots may be trapped in local optima rather than reaching a global optimum, though the resulting observations provide results with good value. This approach can be used to determine where robots should move to observe both static and dynamic targets best, given a definition of value. This can be done in noisy environments and with varying robot capabilities.

For N robots, T targets, and m candidate moves per robot, this algorithm requires (N-2)T Gaussian multiplications (using method in [15]) to predict teammate observations, (m-1)T multiplications to evaluate candidate moves, and (N-1)T multiplications to combine measurement results (actual or estimated). Without communication, estimated next measurements use the previous step's predictive measurement; thus computation is not increased. The one-step optimal requires $m^N T$ multiplications. With 4 targets, one step takes 1.65 seconds for each robot for a team of 2 and 2.42 seconds for a team of 3. In comparison, the single-robot approach requires 0.44 seconds and optimal requires 230 seconds for a team of 2 and several hours for a team of 3 (computed by one robot).

Acknowledgement: This research was supported by DAPRA's Mobile Autonomous Robot Software program.

Note: The first author, Ashley W. Stroupe, is a full-time doctoral student at Carnegie Mellon University.

X. References

[1] R. Arkin. "Motor Schema-Based Mobile Robot Navigation." International Journal of Robotics Research, 8(4), 1989.

[2] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun. "Collaborative Multi-Robot Exploration." ***Proceedings ICRA, 2000***, 2000.

[3] A. Cai, T. Fukuda, and F. Arai. "Information Sharing among Multiple Robots for Cooperation in Cellular Robotic System." ***Proceedings IROS***, 1997.

[4] M. Dietl, J.-S. Gutmann, and B. Nebel. "Cooperative Sensing in Dynamic Environments." ***Proc IROS***, 2001.

[5] G. Dissanayake, H. Durant-Whyte, and T. Bailey. "A Computationally Efficient Solution to the Simultaneous Localisation and Map Building (SLAM) Problem." ***Proceedings ICRA, 2000***, 2000.

[6] D. Gennery. "Traversability Analysis and Path Planning for a Planetary Rover." *Autonomous Robots*, 1999.

[7] R. Grabowski, L E. Navarro-Serment, C. J. J. Pareidis, P. K. Khosla. "Heterogeneous Teams of Modular Robots for Mapping and Exploration." *Autonomous Robots Special Issue on Heterogeneous Multi-Robot Systems*, 2000.

[8] J. Guivant and E. Nebot. "Optimization of the Simultaneous Localization and Map-Building Algorithm for Real-Time Implementation." *IEEE Transactions on Robotics and Automation*, vol. 17, no. 3, 2001.

[9] J. Hespanha, H. Kim, and S. Sastry, "Multiple-agent probabilistic pursuit-evasion games." *Proceedings 38th Conf. On Decision and Control*, 1999.

[10] A. Howard and L. Kitchen. "Cooperative Localisation and Mapping." *Proc 1999 International Conference on Field and Service Robots*, 1999.

[11] H. Moravec. "Sensor Fusion in Certainty Grids for Mobile Robots." *AI Magazine*, Summer, 1988.

[12] L. Parker. "Distributed Algorithms for Multi-Robot Observation of Multiple Moving Targets." *Autonomous Robots*, 12(3), 2002.

[13] R. Pito, "A solution to the next best view problem for automated surface acquisition." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21, 1999.

[14] J. Spletzer and C. Taylor. "Sensor Planning and Control in a Dynamic Environment." *Proceedings ICRA 2002*, 2002.

[15] A. Stroupe, M. Martin, and T. Balch. "Distributed Sensor Fusion for Object Position Estimation by Multi-Robot Systems." ***Proceedings ICRA, 2001***, 2000.