

Value-Based Observation with Robot Teams (VBORT) Using Probabilistic Techniques

Ashley W. Stroupe
Carnegie Mellon University
Pittsburgh, Pennsylvania
ashley@ri.cmu.edu

Tucker Balch
Georgia Institute of Technology
Atlanta, Georgia
tucker@cc.gatech.edu

Abstract

We present a behavior-based approach for directing the movements of robot teams engaged in mapping target objects in their environment. The resulting paths of the robots optimize the vantage points for all the robots on the team, maximizing information gain. At each step, each robot selects a movement to maximize the utility (in this case, reduction in uncertainty) of its next observation. Trajectories are not guaranteed to be optimal, but team behavior serves to maximize the team's knowledge since each robot considers the observational contributions of teammates. The VBORT approach is evaluated in simulation by measuring the resulting uncertainty about target locations compared to that obtained by 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.

1. Introduction

In this work we address the problem of moving a team of robots so as to discover the locations of identifiable targets. In this paper we consider stationary targets, but the approach can easily be applied to moving targets as well. The robots may be homogeneous, or heterogeneous with regard to their sensors 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 the number of targets in advance.

A behavior-based framework generates trajectories for the robots (first proposed in [12][13]). Trajectories are computed one step at a time by optimizing over a value function (or set of functions). VBORT is applied to perform target mapping using teams of communicating robots. At each step, robots move to most improve the collective information about the targets according to a predefined value function that attempts to minimize target uncertainty.

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

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 be inferred if teammates are detectable. 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, we assume that robots should minimize uncertainty about target locations while also minimizing the length of their trajectories. Accordingly, 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. While this approach can be used to find trajectories to optimize any type of value, this paper focuses on our results for mapping sets of static targets.

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.

2. Background and Related Work

There are several areas of work related to this research. Years of work in 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][9], cell-based traversability maps [5] and landmark mapping (typically Kalman-Bucy filter) [4][7], 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][4][7][9]. A second approach is to move some robots while others remain fixed as landmarks [6][8].

The task explored in this research is somewhat different than the coverage task – in particular, we are interested in mapping and tracking 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. Next Best View does this for one robot [10]. Several approaches optimize over all joint robot team actions, which is computationally expensive [11][15].

3. Approach

In our approach to mapping and tracking, each robot chooses 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 (section 5). The algorithm considers a number of candidate one-step 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 for a robot and its teammates.

After a robot determines the best move, it executes that move. New sensor measurements are taken and the covariance matrices are updated. The new covariance is used as the prior in the next step. For static targets, new measurements can be directly combined with previous estimates. For dynamic targets, measurements can be incorporated by updating targets based on motion models.

To choose a move, each robot independently estimates the measurements each teammate would make given their current locations, their sensor models, and the current belief of target locations. This approximation is reasonable if movements are small, and the next set of measurements will be similar. Once these approximate contributions are incorporated, the move that best complements the probability density function (pdf) estimate is chosen by each robot. If robots can communicate, they share probabilistic representations of the observations they have made of the targets, updating the pdf from the previous step to incorporate all new measurements; the pdf is common to all teammates. When communication is not possible but teammates positions are observable, each robot can use the approximations made for move selection to update the pdf. To take full advantage of the team measurements, the individual series of measurements taken by each robot must be later recombined. This may result in paths that are further from optimal, but still incorporate team contributions in movement decisions.

The algorithm, run independently on each robot, is summarized below and illustrated in Figure 1.

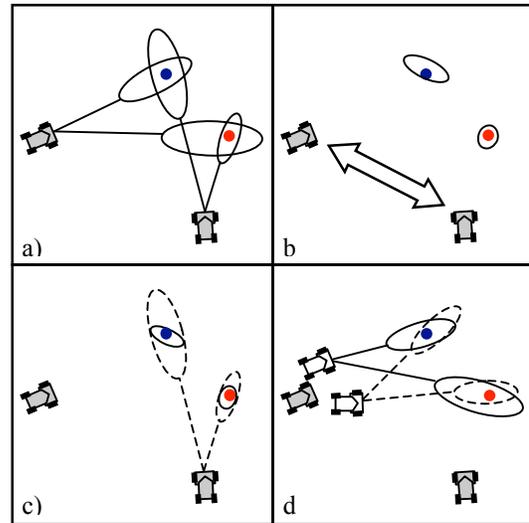


Figure 1. Mapping algorithm. a) Robots measure targets and compute pdfs. b) Robots combine information into single estimates. c) Left robot predicts right robot's observations (dotted) and resulting pdf (solid) in next time step. d) Left robot computes the value of candidate moves (white), given predicted teammate effects.

- 1) Each robot takes measurements on all visible targets and generates a Gaussian covariance of target location (Figure 1a)
- 2) Robots communicate their locations and target position estimates. Target position estimates are combined together and with any previous estimate (Figure 1b).

- 3) Each robot predicts subsequent teammate measurements using current teammate location and sensor model. Predictions are combined with current estimates to predict resulting pdfs (Figure 1c).
- 4) Each robot predicts the measurements and resulting pdf from several candidate moves, combining measurements with the pdf in Step 3 (Figure 1d).
- 5) Each robot selects and executes the candidate move that provides the highest value. Go to Step 1.

4. Implementation

For these experiments the robots can detect range and bearing to observed targets with some sensor noise. The sensor model assumes uncertainty in range scales with range (r) and bearing (θ) uncertainty is a constant angle (Figure 2) as cameras in the Minnow robots [14]. This model may approximate cameras, laser, or sonar. This formulation, with appropriate covariance and Jacobian matrices, may be used for other models.

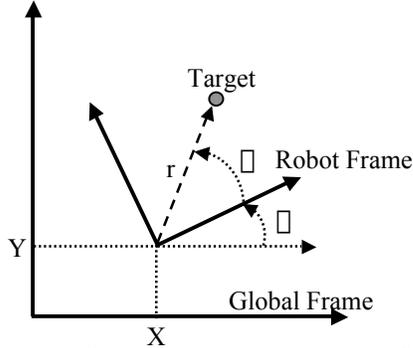


Figure 2. Coordinate frame and parameter definitions.

In Step 1, measurements are taken. The measurement covariance (range-bearing, C_m) is computed as:

$$C_m = \begin{bmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\theta^2 \end{bmatrix} = \begin{bmatrix} (\sigma_r)^2 & 0 \\ 0 & \sigma_\theta^2 \end{bmatrix} \quad (1)$$

where σ_r and σ_θ are sensor parameters described below.

To obtain target covariance in the robot frame, C_t , the Jacobian of the frame transformation, J_m , is applied:

$$J_m^T = \begin{bmatrix} \cos \theta & r \sin \theta \\ \sin \theta & r \cos \theta \end{bmatrix} \quad (2)$$

$$C_T = J_m^T C_m J_m \quad (3)$$

To obtain target covariance in the global frame, C , the transformation Jacobian, J , is used, taking the robot pose uncertainty into account (C_R , if provided):

$$J^T = \begin{bmatrix} \cos \theta & \sin \theta & 1 & 0 & X \sin \theta & Y \cos \theta \\ \sin \theta & \cos \theta & 0 & 1 & X \cos \theta & Y \sin \theta \end{bmatrix} \quad (4)$$

$$C = J^T \begin{bmatrix} C_T \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ C_R \end{bmatrix} J \quad (5)$$

In Step 2, these resulting estimates are communicated and combined (along with any previous estimates) using Equation 6 to produce a single, shared pdf for each target. This is done using a Kalman-Bucy update.

$$C = C \square C [C + C_{new}]^{-1} C \quad (6)$$

To estimate teammate measurements in Step 3, Equations 1-3 are applied again, and Equations 4-5 predict effects on the pdfs.

For Step 4, a set of candidate moves are determined as the set of points reachable at the next time step. For each candidate move, the affects of predicted measurements on the pdfs are determined in the same way as teammate measurements, using Equations 1-5. For each candidate move, the value of the resulting pdfs is determined. Steps 3 and 4 are done simultaneously for all teammates.

The candidate move that maximizes the value function (from step 4) is selected and executed in Step 5.

5. Value Function

VBORT seeks to optimize the value of team observation positions. Clearly the definition of “value” in this context is critical. For mapping and target tracking only, we assume the objective is to obtain the best quality estimates of target locations. We equate “high quality” with low uncertainty. As there is no universally accepted single measure of uncertainty, we must select one.

The value function used is the negative areas of the 1- σ ovals of the Gaussian pdfs (units of distance squared):

$$V = \sum_{i=1:T} \sigma_{i \text{ maj}} \sigma_{i \text{ min}} \quad (7)$$

where σ_{maj} and σ_{min} are major and minor axis standard deviations, respectively, and T is the number of targets. Larger areas correspond to greater uncertainty and to lower value. The optimal location from which to take an observation depends significantly on the value function. Many value functions could be used; we make no claims this one is the “best” though it does provide sensible behavior. Several value functions were investigated, but are not presented for space reasons.

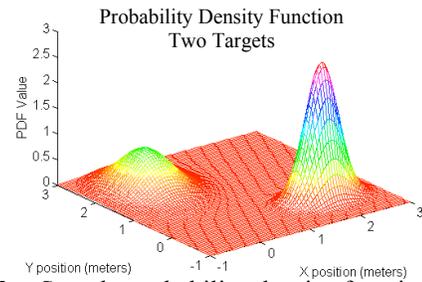


Figure 3. Sample probability density function showing knowledge of two targets, one more certain than the other.

Consider two targets that have been observed, and the

covariance represented by the two-dimensional pdf illustrated in Figure 3. The target at (2,0) has smaller standard deviations than the target at (0,2).

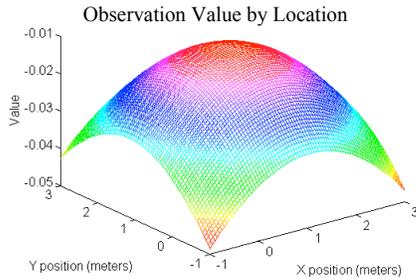


Figure 4. The value of taking one additional observation from each point, given the pdf in Figure 1 and value function in Equation 7.

The value of single measurements from points in the environment (resolution 0.05m) is shown in Figure 4. The maximum value -0.0106 occurs where area reduces equally: (1,1.05). This point provides observations aligned to be as close to perpendicular to the pdf’s major axes as possible and is closer to the less certain target. Thus, this value function serves to encourage observing perpendicular to major axes, essentially “triangulating.”

6. Experimental Approach

A series of experiments was conducted in a *Matlab* simulation of Minnow robots [14]. Movement decisions for all robots on the team were made simultaneously based on the same state information. Termination occurs when all robots prefer their 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; the association problem is not addressed in this work.

Numbers of robots and targets were varied, as were initial conditions and robot capabilities. The candidate moves considered is the set of points on the circle of radius 1-step-size at a resolution of 2°. Two environments were explored: 6x6 meters with 1 or 4 targets (small), and 20x20 meters with 10 targets (large).

For most experiments, measurement noise was reduced to zero and maps were 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 VBORT. Comparing robot trajectories in the presence of noise may be misleading, as differences may be due to measurement differences instead of experimental factors. 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.

6.1. Robot Sensor Model

The vision noise model is Gaussian and has standard deviations of $\sigma_r=10\%$ in range in meters ($\sigma_r=0.1r$) and $\sigma_b=0.5^\circ$ in bearing in radians ($\sigma_b=0.5\pi/180$). Range is 30 meters, 360° to ensure full vision in both experimental environments. Limited sensing reduces range to 2.5 m and bearing to 50°. Parameters are based on Minnow experimental performance, with slightly higher range uncertainty to emphasize asymmetry in measurements.

6.2. Robot Motion Model

Baseline performance includes holonomic motion and speed (step size) of 0.4 meters per 1-second step (roughly one robot length). Simulation of non-holonomic motion reduces maximum turning angle to 50° and does not allow driving backward. Parameters are based on Minnow performance. Robot motion error and uncertainty were ignored in this set of experiments.

7. Experiments

The experiments were designed primarily to investigate the effects of team size, initial conditions, measurement noise, and robot capability limitations. Results are compared to robots using individual reasoning (without regard to teammates) and to one-step optimal.

7.1. Effects of Team Size

The experiments used to evaluate the effects of team size were:

- A) 1 to 4 robots with 1 target (small environment)
- B) 1 to 5 robots with 4 targets (small environment)
- C) 1 to 7 robots with 10 targets (large environment)

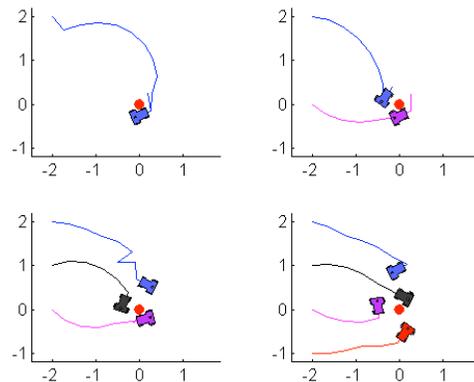


Figure 5. One to four robots observe a target (Exp A). Multiple robots distribute around the target. A single robot must drive around the target to compensate.

As team size increases, robots triangulate more efficiently to reduce uncertainty in all dimensions. Single

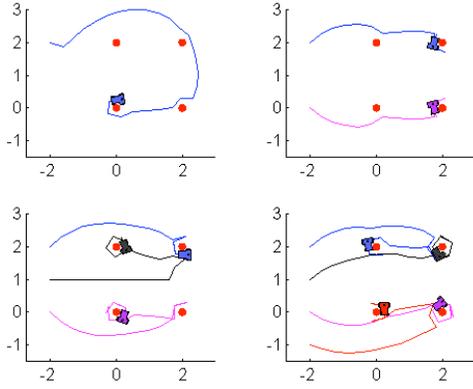


Figure 6. One to four robots observing four targets (Exp B). One robot visits all targets from multiple angles. Two arc apart to get perpendicular views. With more, they examine far targets, then specialize.

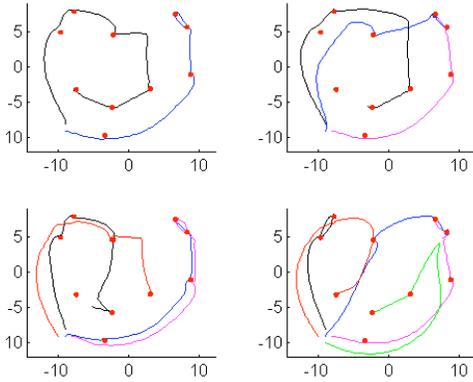


Figure 7. Two to five robots observing ten targets (Exp C). With few robots, they must split up to examine different targets. With larger groups, they concentrate on far away (less known) targets together, then specialize.

robots travel around targets to triangulate with respect to observations over time. Multiple robots directly position themselves to observe targets from complementary axes (orthogonal for two, triangular for three). Generally, more robots allow for more specialization. As the team moves, uncertainties on close targets drop quickly, giving the more distant targets higher priority; robots will concentrate on these far targets, even if multiple robots initially head toward the same target, and later return as the far targets become more certain (Figures 6-7).

7.2. Effects of Initial Conditions

Experiments run to evaluate the effects of varying initial conditions were:

- D) 6 2-robot sets with 4 targets (small environment)

The results of these experiments indicate that varying the initial position relative to the targets leads to markedly different choices of path. At each step, robots move to complement the existing pdf. Different initial conditions

alter the size and orientation of the Gaussians, requiring different complementary observations. Different relative teammate positions must also be complemented with different motions. Example results are shown in Figure 8. When starting close together (upper left) robots must move apart to obtain perpendicular views. When starting with orthogonal views, robots can proceed more directly toward targets. Large changes occur when the robots are moved further apart, but the similar starting locations of the two lower examples result in similar paths.

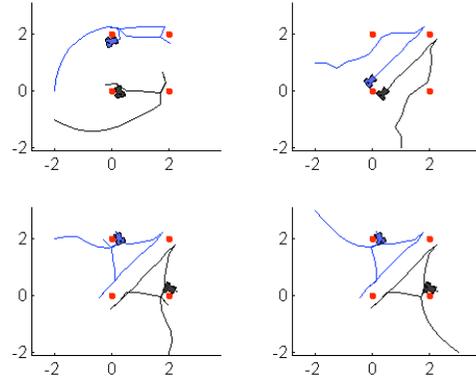


Figure 8. Two robots map four targets from different initial starting conditions (Exp D). Despite variety in actual paths, qualitatively they are similar in that robots first move to get complementary angles.

7.3. Effects of Measurement Noise

Experiments run to evaluate the effects of measurement noise were:

- E) 2 to 5 robots with 4 targets (small environment)
- F) 2 to 5 robots with 10 targets (large environment)

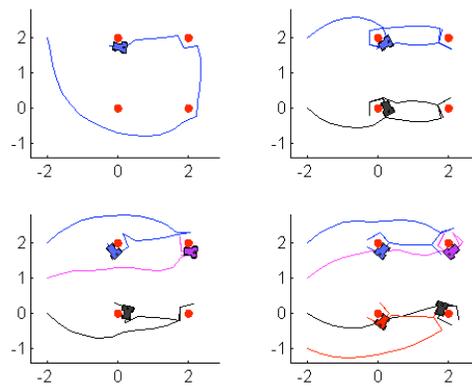


Figure 9. One to four robots observing four targets with added measurement noise (Exp E). Trajectories are not quite as smooth, but present the same general shape as in the absence of noise (Figure 6).

This experiment serves to demonstrate VBORT's robustness in more realistic conditions (Figures 9-10). The model used to introduce noise into simulated sensor

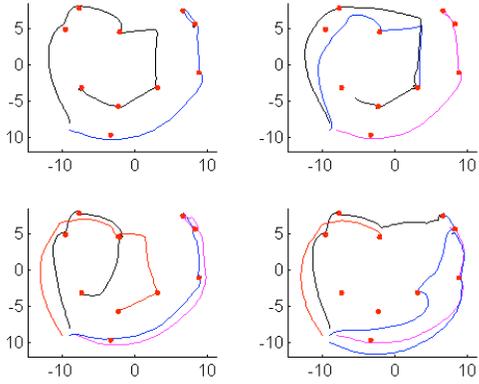


Figure 10. Two to five robots observing ten targets with measurement noise (Exp F). Paths are similar to those without measurement noise.

readings is the same as the sensor model; noise was not introduced into robot motion. Resulting trajectories are very similar to those without noise. Final values are slightly lower (uncertainty higher) since motions are chosen with respect to incorrect target locations and pdfs, especially early on before target locations converge to the correct position. Final target positions do converge to correct positions within one standard deviation.

7.4. Effects of Robot Capability Limitations

Experiments to evaluate performance with limited capabilities (as in section 5, small environment) were:

- G) 1 to 4 robots with limited visual range
- H) 1 to 4 robots with limited visual angle
- I) 1 to 4 robots with non-holonomic motion,
- J) 1 to 4 robots with all above limitations
- K) 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

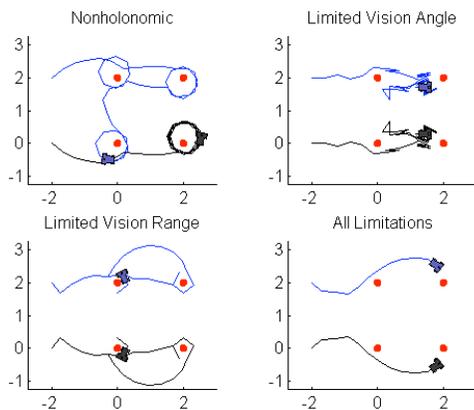


Figure 11. Two robots with various limitations observing four targets (Exp G-J). Despite limitations in turn angle, vision angle, and vision range, robots find paths to achieve good measurements on all targets.

causes robots to circle targets at a larger radius, but paths are qualitatively similar to the case without measurement noise. 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.

Changing step size also affects the results to some degree (Exp K). 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.

7.5. Effects of Target Priority (Preliminary)

One preliminary experimental run was conducted with two robots and four targets to test the effect of varying the priority of targets. Target priority is implemented as a multiplication factor on the value of an individual target; targets with high priority have a high priority coefficient, causing them to contribute more weight to the overall choice of direction. Figure 12 shows how early attention is given to high priority targets; once these are sufficiently explored, one robot moves on to lower priority targets, while one remains at the top priority target.

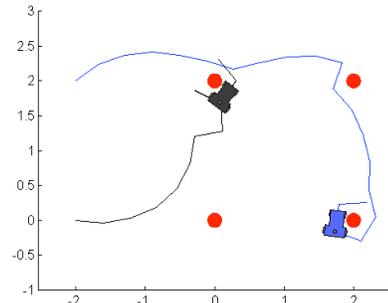


Figure 12. Two robots observe targets with varying priority. Highest to lowest: top left, top right, lower right, lower left. Both robots immediately proceed to top priority target; after that, one robot explores lower priority targets, while one remains at the top priority target.

7.6. Comparative Performance

Performance was compared to that where robots use single-robot reasoning to choose moves and to the one-step optimal.

In the single-robot approach, at each step, robots use only their own observations to determine the next step. For quantitative comparison with VBORT, after all robots have completed independent paths, all collected measurements are combined (resulting in the same total number of measurements) and the resulting observational value is determined.

In comparison to the single-robot approach, VBORT consistently results in better values (Figure 13) when compared after the same number of steps (20 for 4 targets, 50 for 10). The data for single-robot reasoning were gathered using robots run individually from the same starting locations as the team members for the experiments in section 7.2. Resulting values are shown, and other investigated value functions show the same improvement (not reported here). Qualitatively, robots tend to spread out more in the cooperative case, despite similar initial positions relative to the single-robot approach. For individual robots, each robot is essentially mapping the whole space. Robots starting near each other, as would be common in real-robot missions, are drawn to similar locations, taking overlapping measurements and reducing coverage, rather than to complementary locations to take advantage of varying points of view. Individuals perform 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.

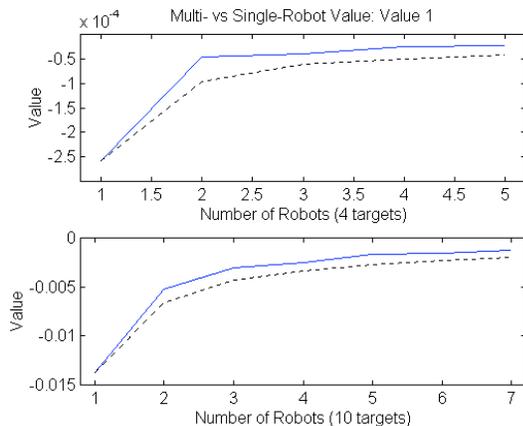


Figure 13. Comparison of value of observations for this approach (solid) and the individual robot approach (dotted). Values improve using the multi-robot approach.

After many steps, giving individual robots time to visit many targets, the resulting observation value provided by VBORT is still greater. The difference is much wider early, providing good values much more quickly. Maps of a desired quality are made in fewer steps using VBORT. Table 1 shows time required to obtain a given maximum uncertainty for the 10-target, 4-robot case.

Table 1. Time (Steps) to Specified Map Quality

| | All Seen | 0.20m | 0.10m | 0.05m | 0.01m |
|--------------|----------|-------|-------|-------|-------|
| Single-Robot | 39 | 41 | 43 | 45 | 47 |
| VBORT | 33 | 36 | 37 | 40 | 49 |

Comparison is also 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. The one-step optimal approach is that in which all combinations of moves are considered in order to find the most valuable *set* of actions for the team. 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 14). Comparative resulting values for the experiment in Figure 14 after 10 steps are provided in Table 2. Our approach and the one-step optimal produce values nearly identical (less than 10% different) and that are almost twice those of the single-robot approach.

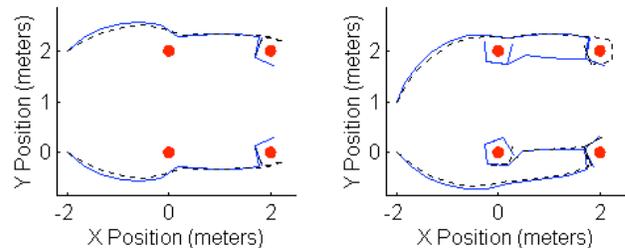


Figure 14. Comparison of one-step optimal trajectories (dotted) to this approach (solid). Trajectories are similar, with this approach spreading robots slightly more, since simultaneous teammate movements are not predicted.

Table 2. Value Comparison of Approaches

| Single-Robot | VBORT | Optimal |
|------------------------|------------------------|------------------------|
| -4.71×10^{-4} | -2.87×10^{-4} | -2.56×10^{-4} |

8. Discussion

The resulting trajectories are highly dependent on the sensor model. Sensors that provide much higher accuracy in one dimension (highly oval covariances) will require more triangulation and produce trajectories that arc strongly around targets; sensors that produce circular covariances will not require this. Value functions must depend on the task; for mapping, any value function that can optimize uncertainty of estimates is useful.

Table 3 demonstrates comparison in complexity and computation time for the single-robot approach, our approach, and the one-step optimal. Complexity of the single-robot approach, our approach, and the one-step optimal can be compared by considering calculations required by a single robot to choose a single move, given N robots, T targets, and m candidate moves. The calculations required by a single robot are considered. Complexity is related to the number of Gaussian multiplications made. In our approach, the first term is teammate contribution prediction and the second term is

move contribution prediction. Without communication, estimated next measurements use the previous step's predictive measurement; computation is not increased. Comparative results are shown for an experiment with 4 targets, 2 and 3 robots and 180 candidate moves. Times are shown for computation time of a single robot.

Table 3. Complexity Comparison of Approaches

| | Single-Robot | VBORT | Optimal |
|------------|--------------|------------|------------------|
| Complexity | mT | (N-2)T+ mT | m ^N T |
| Time (2) | 0.44 s | 1.65 s | 230 s |
| Time (3) | 0.44 s | 2.42 s | >4 hrs |

9. Recent and Future Work

As discussed previously, this approach has been applied to observing multiple dynamic targets. This work is presented separately due to space limitations. Robot movement decisions are based on the same way as in the static target case; in order to account for target motion, uncertainties are grown between steps.

Several additions to this work are in progress. VBORT will be implemented in unknown environments, allowing addition of new targets as they are observed. The approach will be integrated with mapping of landmarks rather than assuming correct robot localization. VBORT will be implemented in small teams of real robots. Lastly, the effects of more complex value functions will be investigated, including adding additional tasks such as exploration, target-related tasks, and target priority.

10. Conclusion

VBORT uses greedy search to determine the best action for each robot using current teammate positions to approximate teammate contributions in the next step. 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 and closely approach the performance of the one-step optimal. The trajectories resulting from our approach are substantially easier to compute yet produce results still very close to optimal. Improvement versus single-robot thinking is most evident with larger environments, with more robots and targets.

Using this behavior-based approach, with limited state-space size and mathematical 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. VBORT can determine robot trajectories 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.

Acknowledgement

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

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] 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.
- [5] D. Gennery. "Traversability Analysis and Path Planning for a Planetary Rover." *Autonomous Robots*, 1999.
- [6] 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.
- [7] 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.
- [8] A. Howard and L. Kitchen. "Cooperative Localisation and Mapping." *Proc 1999 International Conference on Field and Service Robots*, 1999.
- [9] H. Moravec. "Sensor Fusion in Certainty Grids for Mobile Robots." *AI Magazine*, Summer, 1988.
- [10] R. Pito, "A solution to the next best view problem for automated surface acquisition." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21, 1999.
- [11] J. Spletzer and C. Taylor. "Sensor Planning and Control in a Dynamic Environment." *Proceedings ICRA 2002*, 2002.
- [12] A. Stroupe. "Mission-Driven Collaborative Observation and Localization." *Thesis Proposal, Carnegie Mellon University*, 2001.
- [13] A. Stroupe and T. Balch. "Mission Relevant Collaborative Observation and Localization." In *Multi-Robot Systems: From Swarms to Intelligent Automata*, Vol. I. 2002.
- [14] A. Stroupe, M. Martin, and T. Balch. "Distributed Sensor Fusion for Object Position Estimation by Multi-Robot Systems." *Proceedings ICRA, 2001*, 2000.
- [15] S. Sukkariéh, E. Nettleton, B. Grocholsky, and H. Durrant-Whyte. "Information Fusion and Control for Multiple UAVS." In *Multi-Robot Systems: From Swarms to Intelligent Automata*, Vol. II. 2003.