Value-Based Observation with Robot Teams (VBORT) for Dynamic Targets

Ashley W. Stroupe[1], Tucker Balch[2]

[1]*Carnegie Mellon University, Pittsburgh, Pennsylvania, ashley@ri.cmu.edu*
[2]*Georgia Institute of Technology, Atlanta, Georgia, tucker@cc.gatech.edu*

Abstract

We present an application of our VBORT algorithm to observation of multiple dynamic targets with robot teams. VBORT is a distributed algorithm that directs the motions of a robot team to maximize effectiveness in observation tasks. The algorithm chooses the most appropriate move for each robot by determining the move with the greatest value. Value is related to reduction in target estimate uncertainty over all simultaneous robot observations. This approach can be applied to any number of tasks if applicable value functions are provided. Multiple objectives may be simultaneously considered through addition of other value functions. In previous work, VBORT has been applied to mapping static targets and shown to provide results close to a one-step optimal [8]. Here VBORT is applied to observing dynamic objects. Simulation experiments are conducted to evaluate the approach. In the dynamic case, robots are able to follow targets and provide more effective complementary observations than without VBORT.

1. Introduction

Much work has been done in applying robot teams to large-scale problems in order to accomplish tasks more efficiently. One particular application is observation: mapping, target tracking, target search, etc. Most approaches have focused on the gain in efficiency provided by dividing a large problem into several small one-robot problems. For the task of observation, additional efficiency can be gained by taking fuller advantage of the multiple platforms to provide multiple points of view (as in stereo) more quickly than possible with individual robots. In order to take the fullest advantage of the multiple points of view possible with robot teams, each robot must consider the contributions of its teammates so that the results of the *collective* observation are optimized. Communication among teammates can make this optimization more accurate.

The VBORT algorithm is an approach in which each robot on a team approximates the future observational contributions of teammates in order to optimize coordination. In this paper we present the application of our VBORT algorithm for observing or tracking dynamic targets. This extends our work with VBORT applied to mapping the locations of static targets [8].

In order to take full advantage of the ability of teams to share information, the observational contributions of each robot at each step must be considered. For a global optimal, *sets of trajectories* for all robots should be compared; this is impractical, particularly in the case of unpredictably moving targets. In the optimal case for a single step, all possible *sets of moves* for all teammates would be considered and the best one chosen. This works for some problems, but is computationally intractable for large numbers of robots or targets. VBORT approximates the one-step optimal to produce near-optimal results with manageable computation time.

VBORT is a behavior-based approach in the sense that movement decisions are based only on the current state of the world. Unlike more traditional behavior-based approaches, this state of the world includes the positions of teammates as well as current goals and a robot's own position. At each step, each robot approximates the contributions that teammates will make in the next time step in order to best complement their contributions relative to a value function. This approximation assumes a measurement similar to that taken from the robot's current location. For observation tasks, value is related to reduction in target location uncertainty.

While current applications (mapping and dynamic target observation) are stand-alone tasks in their own right, we envision VBORT as part of an approach to more complex problems, providing a means for combining multiple competing motor schemas. By incorporating different types of values, multiple tasks can be optimized simultaneously. Value functions are assumed to depend only on current state and time, making computation simple regardless of number of tasks.

2. Related Work

Observation of dynamic objects has been considered in several applications. Several groups have looked at robot placement to keep moving objects under surveillance by robots [3][4]. Kalman-Bucy filters KBFs) may be used to locate and track a dynamic target from multiple robots [2]. The instantaneous position of a moving target can be refined by combining simultaneous estimates from multiple platforms [7]. Rather than observational coverage of space, the focus of this work is to determine where robots should move such that the results of their combined observations will optimize results (relative to some optimization criteria).

Some work on placing robots for optimizing results has been done. Next Best View predicting where a single robot should to go best improve the quality of a static surface model of an object [5]. This, however, does not consider dynamic objects or simultaneous contributions of teammates. 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 [6]. While this work works well for smaller numbers of targets and robots, the exponential increase in number of particles for increased numbers of targets and the simultaneous optimization of all robots increases computational complexity quickly. Additionally, this framework is not as flexible to introducing other criteria (other than observational uncertainty) into the decision process.

3. Approach

The VBORT algorithm is a behavior-based approach; at each step, robots make a movement decision based on the current state of the world. In order to optimally complement the measurements of teammates, it would normally be necessary to simultaneously consider all potential *sets* of moves. In order to make VBORT manageable for large teams and large numbers of targets, the optimal approach is approximated. If robot motions at each step are small, then all locations reachable in the next time step will produce similar measurements. A robot can thus approximate its teammates' contributions by assuming a set of *average* measurements made from teammates' current locations Measurement results are estimated using the teammates' sensor models. It is relative to this average team contribution that each robot optimizes its next move.

A robot chooses a move in discrete space. Points on the circle of radius one-step-size are potential candidate moves. Typically, candidate moves are selected spaced equally around the circle. More complex sets of candidate moves (such as at variable radii within the reachable circle) may also be considered.

Target positions are represented probabilistically as two-dimensional Gaussian distributions. When targets are observed, positions and covariance are computed using sensor models. All observations from the same time step are combined with the prior covariance or used to initialize for new targets.

To account for position uncertainty caused by target motion, target distributions are "grown'' while leaving the means unchanged. Growing uncertainty increases the Gaussian's area, but does not alter the orientation or ratio of the distribution's major and minor axes (σ_{maj} and σ_{min}) and thus does not affect optimal viewing angles. Growing occurs after movement selection, increasing the uncertainty of the prior distribution before introducing new data. While motion of unseen targets could be predicted using a KBF, we presently assume that unseen targets do not move, allowing the uncertainty to increase unchecked.

For simplicity, robots and targets are assumed to move synchronously. If steps are small and motions relatively slow, this approximation may also produce reasonable results despite slight asynchrony. Largely asynchronous robot and target motions could be handled by providing data timestamps and tracking with the KBF.

The algorithm for each robot, is nearly identical to that for static targets [8] except for the addition of step 6:

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 next 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 predicted measurements with the pdf in Step 3 (Figure 1d).
5. Each robot selects and executes the candidate move that provides the highest value.
6. Uncertainty of targets is grown to account for motion during the previous step. Go to Step 1.

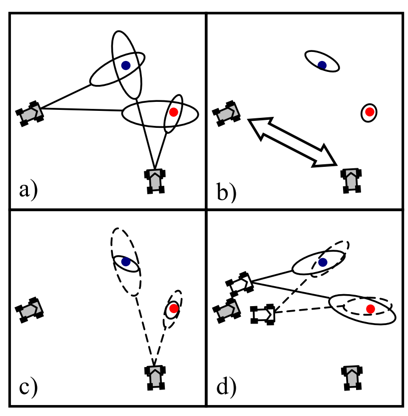

Figure 1. *VBORT a) Robots measure targets and compute pdfs. b) Robots combine information into single estimates. c) Left robot predicts right robot's observations (dotted) and resulting pdf (solid) in next time step. d) Left robot computes the value of candidate moves (white), given predicted teammate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

In step 2, in the absence of communication these actual measurements can be approximated as it is in step 3.

More details of implementation are provided in [8]. In summary, transformations between measurement space and the robot frame and between the robot frame and the global frame are performed using

$$C = J_m^T C_m J_m$$

where C_m is the measurement covariance and J_m is the transformation Jacobian. Combining measurements for prediction (Step 4), evaluation (Step 5), and observation (Step 1) phases use the Kalman-Bucy filter update.

$$C = C - C\left[C + C_{new}\right]^{-1} C$$

4. Value Function

As the object of observing targets is to produce good estimates of targets' locations, certainty of the targets' locations is a reasonable measure of value. The area of the 1-σ oval of the Gaussian pdf representing the target location is larger with greater uncertainty. Thus, we use the negative sum of the 1-σ oval areas for all targets.

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

There may be many ways to measure certainty, and we make no claims that this measure is "best," though it does produce reasonable behavior for both static and dynamic targets [8]. With asymmetrical sensor models (providing more accuracy in one dimension than the other), optimizing this area tends to promote "triangulation" and leads to measurements taken from axes more perpendicular to the major axis of the pdf. As an example in Figure 2, the observation point with the highest value of Equation 3 is (-1.00, -1.05), which is aligned to view partly perpendicular to both axes and is closer to the more uncertain target.

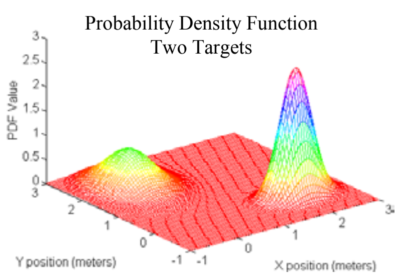

Figure 2. *Probability density function showing beliefs of two targets, one more certain than the other.*

5. Experimental Approach

A series of experiments was conducted in *Matlab* using a simulation of the Minnow robots developed in our lab [7]. All robot movement decisions were simultaneous based on the same state information. After each step, simulated measurements are taken and shared. The covariance resulting from combining all measurements and priors becomes the prior for the next step. In simulation, robots and targets are simultaneously moved after each computational step. 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.

Two types of experiments were conducted: *idealistic* and *realistic*. Idealistic experiments were conducted without simulated measurement noise and were initialized with an a priori target map that provided exactly accurate means but high uncertainty. Experiments without measurement noise allow direct comparison between approaches, with differences in results being due only to differences in approach. The qualitative appropriateness of trajectories is also more evident. Realistic experiments include measurement noise and initialize targets as they are seen with noisy measurements. Experiments with measurement noise demonstrate applicability and robustness to more realistic operating conditions. In both cases, noise in robot motion and position estimation is ignored.

5.1 Sensor Model

We assume a range-bearing sensor (such as a camera) as shown in Figure 3. The vision noise model is Gaussian and has standard deviations of 10% in range in meters ($\sigma_r = 0.1r$) and 0.5º in bearing in radians ($\sigma_\phi = 0.5\pi/180$). Parameters are based on Minnow experimental performance, with slightly higher range uncertainty to emphasize asymmetry in measurements. Vision range varies between 30 meters (for full observation of the environment) and 2.5 meters (for experiments with limited sensing).

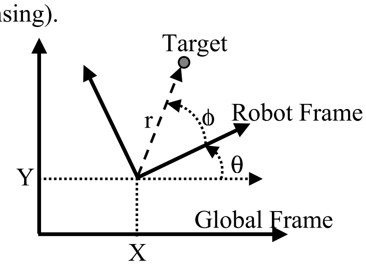

Figure 3. *Coordinate frame and parameter definitions.*

5.2 Robot Motion Model

Baseline performance includes holonomic motion and speed (step size) of 0.4 meters per 1-second step (roughly one robot length). Parameters are based on Minnow performance. Candidate moves were spaced at increments of 2 degrees around the reachable circle. Robots may also remain in place.

5.3 Target Motion Model

The trajectories of targets for the robots to track were designed by hand. Targets, except as noted, move more slowly than robots, 0.2 to 0.3 meters per 1-second step. As noted above, target positions are represented by two-dimensional Gaussian distributions. To allow for target movement, variance, σ^2_{maj} and σ^2_{min}, is grown by 0.5 m^2 in each step, an amount larger than the target step size.

6. Experiments and Results

A set of experiments were conducted using the idealistic approach. Several experiments were repeated using the realistic approach. These experiments are:

A. Three robots (one stationary) with one target moving linearly at a constant velocity.
B. One to three robots with one target moving linearly at a constant velocity the same as robots.
C. Three and four robots with four targets moving linearly away at a constant velocity.
D. Three and four robots with four targets moving piece-wise linearly at nearly constant velocities.
E. One robot with one target moving linearly at three constant velocities (slower than, same as, and faster than robot).

6.1 Tracking Behavior

Examples of the tracking behavior of the robots as demonstrated by Experiments B, C, and D are shown in Figures 4-6. With single targets, or multiple robots per target, robots tend to spread out evenly around targets to produce complementary views.

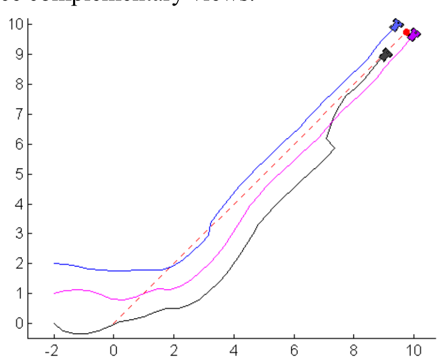

Figure 4. *Three robots track a target (Exp B). Robots settle into a triangle formation to obtain complementary views.*

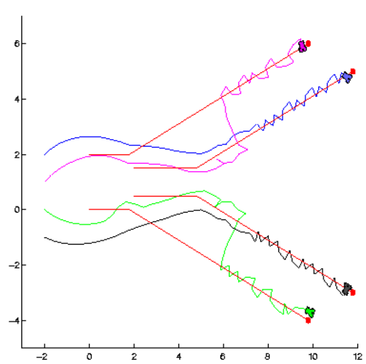

Figure 5. *Four robots observe four targets (Exp D). While targets are together, robots arc to obtain perpendicular views. As targets separate, robots specialize and follow single targets based on proximity.*

This can most clearly be seen in Figure 5, where the three robots form a triangle around the single target. Multiple targets produce more complex behavior. Targets moving together allow robots to arc around to vary the observation angle. As targets diverge, robots tend to focus on a subset of the targets based on their proximity to them. If targets move out of visual range, robots switch attention as needed. In Figure 5, four robots are able to arc around four targets while they are close; once they diverge, the team splits to follow. In Figure 6, divergence requires instant specialization.

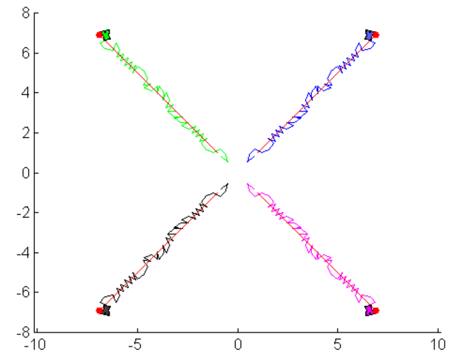

Figure 6. *Four robots track four targets (Exp C). Robot immediately specialize and follow targets based on proximity.*

6.2 Effects of Team Size

Team size effects were investigated with Experiments B, C, and D. When compared with the combined observations of multiple robots moving independently, value improves 14 times for two robots following a single target, and 11 times for three robots (Figure 4). Qualitative differences are evident: three collaborating robots form a triangle while two take perpendicular views (Figure 5); multiple individuals follow similar paths obtaining similar (redundant) measurements. This is also apparent in Figure 7, as the robot configuration changes to accommodate an additional (fixed) teammate's entry and departure in the formation. The formation is not perfectly reformed as one robot falls behind during reconfiguration.

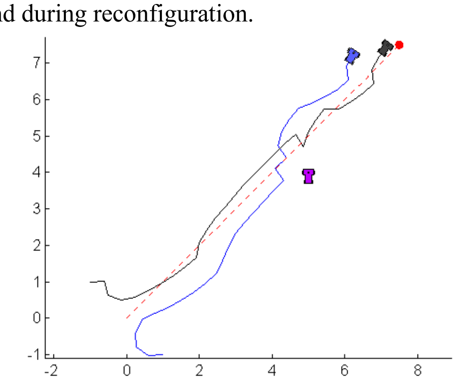

Figure 7. *Two robots track a target (Exp A). Robots stay on opposite sides of the target to get complementary views. As a third robot is approached, they reconfigure to a triangle.*

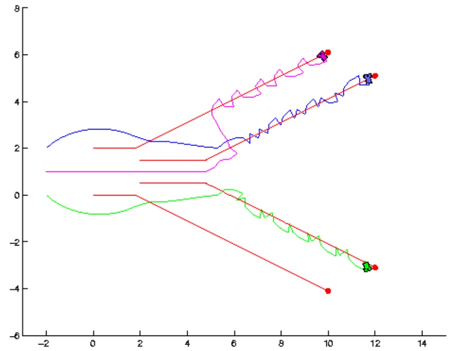

Figure 8. *Three roots observe four targets (Exp D). While targets cluster, robots arc to obtain different views. Once specialization is required, it occurs; the bottom robot oscillates less to remain closer to the fourth target.*

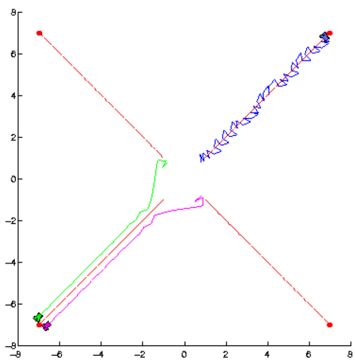

Figure 9. *Three roots observe four divergent targets (Exp C). Robots immediately specialize by proximity. Measurements by all robots allow reasonable observations of the fourth target.*

In figures 8 and 9, compared to 5 and 6, we see that smaller team sizes do not allow perfect specialization once targets separate. While targets remain close, good measurements can be taken on all targets with fewer robots, with especially good measurements on the closest targets. When targets separate, robots typically follow a single target closely; in Figure 9 we observe that the choice of which targets to follow allows perpendicular views on the distant targets. In these experiments, the increased value on the closely observed targets overwhelms the small potential improvement by moving toward the far target, preventing robots from switching targets. Note oscillation may be reduced, keeping robots more often closer to the far target (Figure 8). This is a limitation of a behavior-based approach: while multiple steps could take robots close enough to the far target to produce an overall better result, single steps cannot provide enough improvement on the far target to draw robots away. Faster robots (larger steps), robots may jump over local optima and choose to observe far targets. If, at far distances, the growing uncertainty becomes larger than the reduction from observation, this may also induce robots to move closer to the far target.

6.3 Effects of Measurement Noise

Experiments A-D were repeated using the realistic approach, with the addition of measurement noise at the observation step (Step 1) and initializing new targets with noisy measurements. The uncertainty model used to introduce noise into simulated sensor readings is the same as the sensor model. Resulting trajectories are very similar to those without noise. Trajectories are less smooth as errors from noise slightly divert robots or are corrected. The similarity and the slightly jagged nature of the paths are evident in Figure 10.

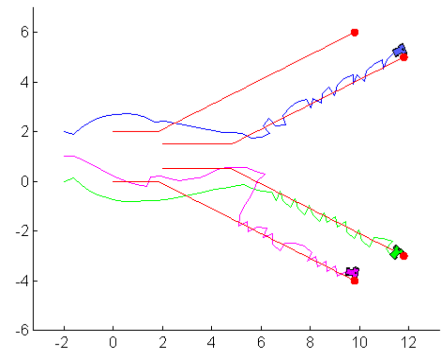

Figure 10. *Three roots observe four targets with noisy measurements (Exp D). Note that behavior is similar to that without noise.*

6.4 Effects of Limited Sensing

Experiments C and D were repeated with a limitation on sensing range in order to investigate resulting changes in behavior. In these experiments, the vision range was limited to 2.5 meters, making it impossible for robots to see all targets at all times. As shown in Figure 11, robots arc to coordinate then specialize in targets. As targets move out of range of one or more robots, remaining low-quality observations on these targets do not significantly decrease uncertainty to counteract the uncertainty growing in Step 6. As uncertainty grows, these far targets begin to dominate the value function,

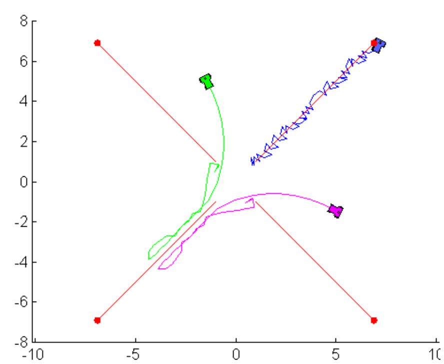

Figure 11. *Three robots observe four divergent targets with limited sensing range (Exp C). With unseen targets, robots switch to observe them. Here, robots switch to the lower right target and then back to individual targets. While following lower right target, perpendicular views of others are gained.*

causing robots near enough to respond to move to observe the uncertain targets. Due to the behavior-based approach looking ahead only one step, if a target moves out of range to where no single step will allow any robot observation, the target becomes lost and no robots will be induced to try to move toward it.

6.5 Effects of Varying Velocity

The effects of varying target velocity relative to robot velocity were investigated in Experiment E. Oscillatory motion around target paths occurs when robots can move more quickly than the target (Figures 4-11). At first one might conclude that the oscillations are due to robots preventing overtaking and moving farther from targets. However, robots have the option of moving as close as possible to a target and then remaining in place until another move allows it to get closer. In other words, oscillation is not the only means available for remaining close to slow targets. Oscillations allow single robots to obtain measurements from more orthogonal vantage points, thus improving results. This is borne out by looking at values: a slightly longer-range measurement along an orthogonal axis reduces area more than a closer-range measurement along the same axis. This oscillation behavior reduces and disappears when targets are moving faster than robots, as robots must use all forward velocity to keep the target in close visual range as long as possible (Figure 12). This behavior also disappears when multiple robots follow single targets (Figures 4 and 9), eliminating need for robots to perform their own triangulation.

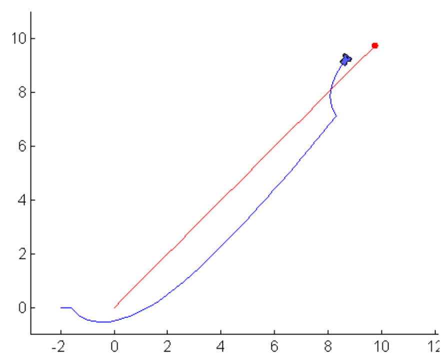

Figure 12. *One robot tracks a target moving at the same velocity (Exp E). The robot cannot often triangulate by moving axes in order to keep up with the target, but does arc at first, and then switches sides later.*

7. Discussion

VBORT can be used to determine where robots should move to observe both static and dynamic targets best, given a definition of "best" is available. In this case, we rely on a provided measure of the value of an observation. VBORT improves the accuracy and speed of overall measurements by considering teammate locations. Using limited mathematical complexity, an approximation to optimizing over teammates, and low state-space complexity, VBORT makes decisions quickly for large teams and target numbers. This allows robots to react quickly to dynamic and changing situations. As target directions change, robots automatically adjust, taking each other into account, without having to plan. This can be done in noisy environments and with varying robot capabilities [8]. The complexity of this algorithm is far less than the one-step optimal.

Comparisons with the one-step optimal using static targets demonstrated that little value is lost due to the approximation [8] and the results are summarized here. In Table 1 performance is compared for two robots and four targets by comparing final map values after a fixed number of steps, indicating improvement over single-robot and similarity to one-step optimal. In Table 2, improvement in computational complexity for a single robot in a single step is demonstrated for N robots, T targets, and m candidate moves. Computation time for 2 and 3 robots, 4 targets, and 2 moves is also improved.

Table 1. Value Comparison of Approaches

Single-Robot	VBORT	Optimal
$-4.71\text{x}10^{-4}$	$-2.87\text{x}10^{-4}$	$-2.56\text{x}10^{-4}$

Table 2. 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

A limitation common to behavior-based approaches is that robots may produce a trajectory that makes use of local optima rather than global optima; the one-step optimal may also have this problem due to the limited look-ahead. While paths may not be globally optimal, paths do produce reasonable results, improving the accuracy and speed of target location and improving reactivity to dynamic targets.

8. Concurrent and Future Work

In addition to the work with dynamic targets here, previous work has been done using VBORT to locate multiple static targets with robot teams. Using VBORT, robot teams are able to produce high quality maps in shorter time than in approaches which do not consider teammate contributions directly [8].

Several additions to this work will be made in the near future. The approach will be implemented with completely unknown environments, requiring mapping of landmarks for localization. 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 and deviation from pre-planned paths.

9. Conclusions

VBORT 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 by incorporating teammate contributions without structured planning. The resulting combined measurements are of better value (defined by a value function) than are combined measurements of individually-minded robots. Resulting trajectories of robots are not globally optimal with respect to target trajectories and values, but are reasonable and produce improvements in results. Despite very simple target motion models, robots keep targets under observation and produce high-quality estimates of target position. Due to a behavior-based approach looking ahead one step, robots may be trapped in local optima (following close targets rather than moving toward farther targets) and may not deliberately recover lost targets.

10. Acknowledgements

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

11. References

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

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

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

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

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

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

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

[8] A. Stroupe and T. Balch. "Value-Based Observation with Robot Teams (VBORT) Using Probabilistic Techniques." Submitted to ICAR 2003.