

Dynamic Sensor Planning and Control for Optimally Tracking Targets

John R. Spletzer Camillo J. Taylor
GRASP Laboratory – University of Pennsylvania
Philadelphia, PA 19104, USA
{spletzer, cjtaylor}@grasp.cis.upenn.edu

Abstract

This paper presents an approach to the problem of actively controlling the configuration of a team of mobile agents equipped with cameras so as to optimize the quality of the estimates derived from their measurements. The issue of optimizing the robots' configuration is particularly important in the context of teams equipped with vision sensors since most estimation schemes of interest will involve some form of triangulation.

We provide a theoretical framework for tackling the sensor planning problem, and a practical computational strategy inspired by work on particle filtering for implementing the approach. We then extend our framework by showing how modeled system dynamics and configuration space obstacles can be handled. These ideas have been applied to a target tracking task, and demonstrated both in simulation and with actual robot platforms. The results indicate that the framework is able to solve fairly difficult sensor planning problems online without requiring excessive amounts of computational resources.

1 Introduction

The idea of using teams of small, inexpensive robotic agents to accomplish various tasks is one that has gained increasing currency in the field of robotics research. Figure 1 shows a picture of a Clodbuster robot which is based on a standard remote controlled motion platform and outfitted with an omnidirectional video camera – its only sensor. Using teams of these modest robots, fairly sophisticated applications such as distributed mapping, formation control and distributed manipulation have been successfully demonstrated [1, 2].

One of the more interesting aspects of these platforms is that estimates for relevant quantities in the world are formed by combining information from multiple distributed sensors. For example, the robots in the team shown in Figure 1 obtain an estimate for their relative configuration by combining the angular measurements obtained from all of the omnidirectional images and performing a simple triangulation operation. Similar techniques can be used to estimate the locations of other features in the environment, such as the box they are manipulating. In fact, one could choose to view the team in




Figure 1: A single Clodbuster robot (left) and the team performing a distributed manipulation task.

Figure 1 as a three-eyed stereo rig where the individual eyes can actually be moved on the fly.

This capability invites the following question: given that the robot platforms are mobile, how should they be deployed in order to maximize the quality of the estimates returned by the team? This is a particularly important question in the context of robots equipped with vision sensors since most of the estimation techniques of interest in this case are based on some form of triangulation.

Similar questions arise when one considers the problem of integrating information from a sea of distributed sensors. Given that there is some cost associated with transmitting and processing data, which sensor readings should one use to form an estimate for the parameters of interest?

This paper presents a theoretical framework for discussing such questions and a practical computational approach, inspired by work on particle filtering, for tackling them. The suggested approach could be viewed as an application of the theory of games since the problem of controlling the robots' configuration is reformulated as the problem of optimizing a quality function that reflects the expected value of assuming a particular formation. Results obtained by applying this approach to a target tracking task are presented in Section 3.

It is important to note that while the approach was developed to handle the problems faced by teams of robots equipped with vision sensors, it could also be used to deploy robots equipped with other types of sensors like laser range finders or sonar systems.

1.1 Related Work

The focus of this research is a probabilistic framework which exploits the degrees of freedom afforded by robot mobility to actively manage sensor positions for improved state estimation. We demonstrate its effectiveness in an optimal target tracking task. "Optimal tracking" can be defined using various metrics. We choose to minimize the expected error in tracking target positions. Since the measurements of multiple robots are combined to estimate target pose, this relates strongly to work in sensor fusion.

In our target tracking task, robots rely on omnidirectional cameras for tracking groups of targets. Merging measurements from multiple vision sensors for improved

state estimation was considered by Bajcsy and others under the heading of *Active Perception*. Improvements were seen in various performance metrics, including ranging accuracy [3, 4]. Our framework can be viewed as an extension of this paradigm to distributed mobile robotics.

Durrant-Whyte *et al* pioneered work in sensor fusion and robot localization. This yielded significant improvements to methods used in mobile robot navigation, localization and mapping [5, 6]. Thrun *et al* have also contributed significant research to these areas [7, 8]. The work of both groups has emphasized probabilistic techniques for data fusion - with a recent focus on particle filtering methods. Our approach is also probabilistic, and it too leverages particle filtering methods. However, our work distinguishes itself from traditional data fusion techniques in that the sensors themselves are actively managed to improve the quality of the measurements obtained *prior* to the data fusion phase, resulting in corresponding improvements in state estimation.

Since the sensors are actively managed, our work relates to research in on-line sensor planning as well. Relevant to our approach was a methodology for distributed control proposed by Parker [9]. This framework - Cooperative Multi-Robot Observation of Multiple Moving Targets (CMOMMT) - attempted to maximize the collective time that each target was being observed by at least one robot.

The theory of games has also provided inspiration for similar research in target tracking. The pursuit-evasion problem was investigated by LaValle *et al* and Fabiani *et al* [10, 11]. Both examined the task of maintaining target visibility in a cluttered environment known *a priori* to the pursuer. LaValle's approach generated trajectories that minimized a loss function which grew when the target became occluded. Fabiani's motion strategy was based on the expected maximum value of a corresponding utility function U . Of the two, Fabiani's work is more relevant. The value of U was affected by uncertainty in the target's position, which was indirectly affected by uncertainty in the pursuer's position. As a result, the pursuer trajectory was influenced not just by target position, but also by known landmark positions in the environment which could be used to reduce uncertainty in target pose.

In all three of these cases, the optimization criterion was based on maintaining target observability, rather than the quality of the observation. Additionally, the work of LaValle and Fabiani was limited to the case of a single pursuer/evader. In theory, both could be extended to multiple agents. However, in practice the resulting explosion in computational complexity would be prohibitive. In contrast, the complexity of our framework is implementation dependent, and may be tuned by the user to scale very efficiently in terms of the number of robots and targets.

In the Next Best View (NBV) problem, sensor placement is of primary concern [12, 13]. Given, for example, previous range scans of an object, an NBV system attempts to determine the next best position of the scanner for acquiring the object's complete surface geometry. As in our framework, the emphasis is optimizing sensor placement. However, NBV is intended for use in a static environment. Inherent in our approach is the ability to handle dynamic scenes which makes it more akin to a control law for distributed sensors.

2 Developing the Framework

2.1 A Theoretical Approach

This section describes the theoretical framework used to discuss the problem of sensor deployment. In order to ground the terminology, we will describe how various elements in the framework relate to the scenario depicted in Figure 2. In this example, three robots are tasked with localizing a single moving target.




Figure 2: Target localization by a robot team.

Let \mathcal{C}_r denote the configuration space of the robot platforms. In this case, one can consider the vector formed by concatenating the positions and orientations of the three platforms with respect to some base frame of reference $[x_1, y_1, \theta_1, x_2, y_2, \theta_2, x_3, y_3, \theta_3]^T$. Let $\rho \in \mathcal{C}_r$ denote an element of this configuration space. Similarly let \mathcal{C}_w denote the configuration space of the parameters under consideration. In Figure 2, this space is particularly simple since we need only consider the position of the moving target with respect to the base frame denoted by the vector $[x_t, y_t]^T$. In general, however, this space can be much more complicated. Let $\omega \in \mathcal{C}_w$ denote an element of this configuration space.

Let \hat{Z} represent the set of all possible sensor measurements, and let $\hat{z} \in \hat{Z}$ denote our measurement vector. In this example, \hat{z} corresponds to the vector formed by concatenating the three angles measured by the robots $[\alpha_1, \alpha_2, \alpha_3]^T$. The hat serves to remind us that these measurements are corrupted by noise. It is assumed that the designer has some model for the noise process, which is given in the form of a conditional probability density function $P(\hat{z}|\rho, \omega)$. This function allows us to predict the distribution of measurements as a function of ρ and ω .

Let $Est : \mathcal{C}_r \times \hat{Z} \rightarrow \mathcal{C}_w$ denote a function that produces an estimate of the target's position, $\hat{\omega}$, from the noisy measurements, \hat{z} , and the robots' configuration, ρ . $Disp : \mathcal{C}_w \times \mathcal{C}_w \rightarrow \mathcal{R}^+$ is a function which returns an indication of the disparity between an estimated value $\hat{\omega}$ and the actual value ω . For our target tracking example, an appropriate function might be the Euclidean distance $Disp(\omega, \hat{\omega}) = \|\omega - \hat{\omega}\|_2$. Other applications may require more sophisticated error metrics.

$P(\omega)$ denotes a probability density function on the configuration space \mathcal{C}_w which can be used to model prior information about the values of the parameters of interest. For example, one may have some information about where the target could be based on prior measurements or a dynamic model for the target's motion. Given this

terminology, one can define a quality function $Q(\rho)$ as follows:

$$Q(\rho) = \int_{\mathcal{C}_w} \int_Z Disp(\omega, Est(\rho, \hat{z})) P(\hat{z}|\rho, \omega) P(\omega) d\hat{z} d\omega \quad (1)$$

This function captures how the expected error in the estimate, $\hat{\omega}$, varies as the robots configuration changes. Note that there are, of course, several alternative definitions for this quality function that are equally reasonable. One could consider the maximum expected error in the estimate or the median expected error. Different choices for $Q(\rho)$ may be more appropriate in certain situations.

With these notions in place, one can formulate the problem of choosing an appropriate configuration for the robots as an optimization problem as shown below:

$$\min_{\rho \in \Delta} Q(\rho) \quad (2)$$

The goal in this case is to find a choice of $\rho \in \Delta$, where $\Delta \subset \mathcal{C}_r$, which minimizes the quality function $Q(\rho)$. Limiting the optimization to a subset of \mathcal{C}_r , Δ , allows us to model situations where certain configurations cannot be achieved due to obstacles in the environment, sensor constraints or limitations on the range of motion of the robots.

Note that even though the approach is being discussed in the context of target tracking, the framework is general enough to be applied to a wide range of sensor planning problems. The specifics of the task would be reflected in the definitions of \mathcal{C}_r , \mathcal{C}_w , \hat{z} , Est and $Disp$.

2.2 A Computational Solution

For most interesting systems the optimization problem given in equation 2 is difficult to solve analytically. However, it is possible to approximate this process computationally. To do this we draw inspiration from prior work on particle filtering [14].

In particle filtering, probability distributions such as $P(\omega)$ are approximated by sets of tuples (ω_j, π_j) , where ω_j is a single sample from \mathcal{C}_w and π_j a weight that reflects the likelihood of ω_j representing the state ω . By making use of this approximation, we can replace the integrals of equation 1 with a summation:

$$Q(\rho) \approx \frac{1}{N} \sum_{j=1}^N Disp(\omega_j, Est(\rho, \hat{z})) \quad (3)$$

where N corresponds to the number of samples selected at random from $P(\omega)$. By choosing N sufficiently large, this single summation can reflect the effects of both integrations in equation 1.

The computation of $Q(\rho)$ is outlined in Algorithm 1. For a more concrete explanation, we refer to Figure 3. To estimate $Q(\rho)$, a sample ω is chosen at random from $P(\omega)$, and projected from target space to each robot's image frame (3a). The corresponding image measurements are then corrupted with noise from our sensor model, and projected back into target space using our Est function (3b). We can obtain the disparity between this new sample $\hat{\omega}$ and our original sample from $Disp(\omega, \hat{\omega})$ (3c). This procedure is repeated for the N samples. Using $\|\omega - \hat{\omega}\|^2$ as our $Disp$ function, $Q(\rho)$ then reflects the expected mean-squared error in target position.




Figure 3: Evaluating $Q(\rho)$. A sample ω is taken randomly from $P(\omega)$ and projected into the sensor frames (a), corrupted with noise from our sensor model and reprojected as $\hat{\omega}$ (b), so that $Disp(\omega, \hat{\omega})$ can be calculated (c). This is repeated for N random samples.

A simple but effective approach to optimizing the robot configuration is to approximate the gradient of the quality function, $\frac{\partial Q}{\partial \rho}$, by sampling its value in the vicinity of the current robot configuration. The control law governing robot motion would then be

$$\dot{\rho} \propto -\frac{\partial Q}{\partial \rho} \quad (4)$$

Gradient based approaches can often lead to local minima - and consequently sub-optimal performance - in traditional optimization problems, and local minima do in fact exist when tracking multiple targets. However, in a dynamic environment $Q(\rho)$ is actually redefined whenever the estimated positions of the targets with respect to one another change. As a result, we often have only a single time-step for optimizing a given $Q(\rho)$, and in such cases can generate “piece-wise optimal” trajectories using a gradient approach. Alternatively one could employ standard optimization techniques, like the Simplex method [15] to choose the best achievable robot configuration in the vicinity for the next time instant.

Algorithm 1 Calculate $Q(\rho)$ for a given $P(\omega)$

```

err ← 0
for i = 1 to N do
  ωi ←∈ P(ωi) {Choose a sample at random from P(ωi)}
  ẑi ←∈ P(ẑ|ρ, ω) {Choose a sample at random from P(ẑ|ρ, ω)}
  ω̂i ← Est(ρ, ẑi) {Generate an estimate for ω̂}
  err ← err + Disp(ωi, ω̂i)
end for
Q ←  $\frac{err}{N}$ 

```

Note that it is possible to incorporate knowledge of system dynamics into this framework in the usual manner. In the CONDENSATION algorithm described by Isard and Blake [14], a particle distribution $P(\omega)$ is propagated at each time-step according to a known dynamic model. This same $P(\omega)$ serves as the assumed input for our framework, and establishes a complementary relationship between sensing and control, as the same particle sets used for tracking targets are also used to control the robot team for improving future tracking estimates.

3 Experimental Results

3.1 Simulation Experiments

In order to demonstrate the utility of the proposed framework, we first apply it to three sensor planning problems in simulation: tracking a single point target, tracking multiple point targets, and tracking a box. We then extend the point target tracking problem by incorporating a dynamical model for the target. Finally, we integrate motion planning techniques for local obstacle avoidance and demonstrate target tracking in a cluttered workspace. Each of these scenarios is explained in more detail below.

For the sake of simplicity we have assumed in all of these scenarios that the robots can accurately measure their positions and orientations with respect to one another, since it is the robot positions relative to the targets that are of interest. Note that we could consider error in the positioning of the robots within this framework by adding extra noise terms to the measurements or by including the robots configuration as part of the state to be estimated.




Figure 4: Generated trajectories (left) and disparity measurements for 2 robots tracking a static point target.

3.1.1 Tracking a single point target

For the first scenario we consider 2 robots equipped with omnidirectional cameras, and tasked with tracking a single target. \mathcal{C}_r represents the concatenation of the robot positions, \mathcal{C}_w the target position, and \hat{z} the 2 angles to the target measured by the members of the robot team. We assume \hat{z} to be corrupted with random noise generated from our sensor model. $Est(\rho, \hat{z})$ returns an estimate for the target position, $\hat{\omega}$, which minimizes the squared disparity with the measurements, \hat{z} , and $Disp(\omega, \hat{\omega})$ simply returns the Euclidean distance between the estimated target position and the actual value.

In our simulations, robot motions are constrained by the maximum robot velocity and the robots positions are limited by mandating a minimum standoff distance to the target. These serve to define the valid configuration space for the robots, Δ . Results from Matlab simulations for two robots with both static and dynamic targets are provided below. For these trials, 100 exemplars were used to approximate $P(\omega)$, and the sensor model (for all trials) was assumed to be Gaussian noise with $\sigma = 1^\circ$.




Figure 5: Generated trajectories (left) and disparity measurements for 2 robots tracking an unpredictable moving point target.

Figure 4 shows the static target case for two robots. Trajectories for this symmetric case are predictable and consistent with simulation results, as are the dramatic drops in estimation error over time. Similar results are obtained for the case of an unpredictably moving target, as shown in Figure 5 (Extension 1).

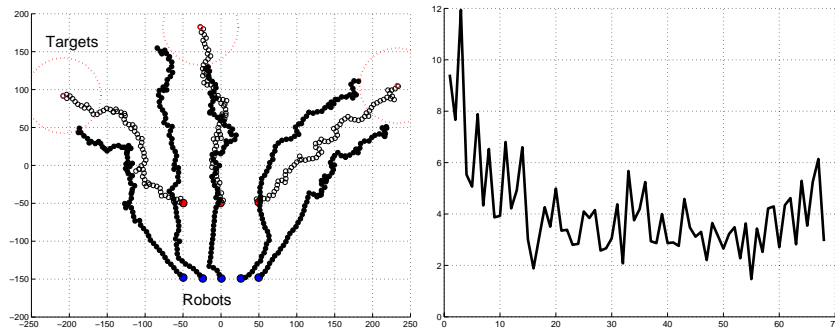


Figure 6: Generated trajectories and summed disparity measurements for 5 robots tracking 3 point targets.

3.1.2 Tracking multiple point targets

For the second scenario, we examine the more interesting problem of n robots tracking m independently moving, unpredictable point targets. This problem can be tackled in much the same manner. \mathcal{C}_w now represents the concatenation of possible target positions, and \hat{z} the corresponding $n \times m$ angles measured from robots to targets. $Est(\rho, \hat{z})$ approximates the position of every target, and $Disp(\omega, \hat{\omega})$ returns the summed disparities between estimated and true target positions.

Results from a pair of simulation runs can be found in Figures 6 and 7 (Extension 2). In these trials, 3 unpredictable targets were tracked by 5 and 4 robots, respectively.

Note the behavior of the robots as they move from their original positions to more advantageous vantage points. The robots automatically split off to track targets with-




Figure 7: Generated trajectories and summed disparity measurements for 4 robots tracking 3 point targets.

out any need for sophisticated switching rules for arbitrating robot-target assignment. The final configuration is simply a consequence of the definition of the $Q(\rho)$ function that the system attempts to optimize. Note also that it is not possible in these scenarios to assign two robots to every target so the robots distribute themselves automatically to come up with the best composite estimate. This is significant, as relatively complex tracking behaviors can be implicitly encoded into $Q(\rho)$, and the need for explicit switching controllers is mitigated.

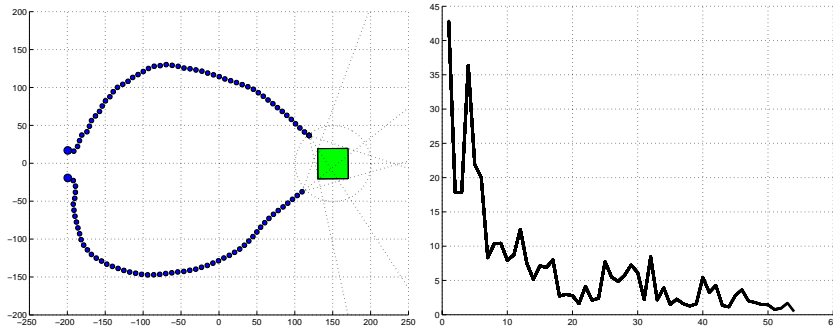


Figure 8: Generated trajectories and summed disparity measurements for 2 robots tracking a static box.

3.1.3 Tracking a box

For the third case, we consider the problem of using the measurements from the robots to estimate the configuration of a box in the scene. This example demonstrates how the proposed framework can be applied to scenarios where the state estimate is not simply the concatenation of a set of point locations. Here the configuration space \mathcal{C}_w is identified with $SE(2)$ and elements of this set denote the position and orientation of the box. The robots can measure the angles to all of the visible box corners, \hat{z} . The estimation function $Est(\rho, \hat{z})$ as always is nonlinear, and minimizes the disparity between the predicted angles to the visible corners and the actual measurements.

$Disp(\omega, \hat{\omega})$ returns the sum of the distances between the recovered box corners and the actual corners. For these trials, 20 “box” exemplars were used to estimate $P(\omega)$. Sample simulation results can be found in Figures 8 and 9.

In both cases we can see that the robots not only migrate to positions more advantageous for tracking the corner features, but also for maximizing the *number* of visible features. The latter effect is a result of the *Est* function using only the visible corners to estimate the box pose. Inherently better estimates are obtained when more features are available.




Figure 9: Generated trajectories and summed disparity measurements for 3 robots tracking a static box.

3.1.4 Incorporating the Dynamical Model

Integrating target dynamics into sensor planning often provides significant improvements in tracking performance. Dynamical models can be obtained using an approximation of target dynamics, or through “learned” models as demonstrated in [14]. For our simulations, we employed the former approach.

Consider the case of n observers on the ground tracking a ball traveling through the air with some unknown initial velocity V_i . We again model these observers as robots equipped with omnidirectional cameras. In this case, \mathcal{C}_r represents the concatenation of the robot positions which are constrained to operations in the x - y plane, $\mathcal{C}_w \subset R^3$ represents the space of target positions. The measurement vector \hat{z} denotes the n azimuth and elevation angle pairs to the target measured by members of the robot team.

We approximated the dynamical model for the ball $\frac{d\hat{\omega}}{dt}$ by assuming constant acceleration under gravity g and estimating its velocity from position measurements over time. Actual ball dynamics $\frac{d\omega}{dt}$ in the simulation were slightly more realistic and also modeled drag effects. The corresponding equations of motion are given in Equations 5 and 6, where μ is density, and m , d , C_D , and v_0 are the mass, diameter, drag coefficient, and current velocity of the ball, respectively.

$$\frac{d\hat{\omega}}{dt} = v_0 t + \frac{gt^2}{2} \quad (5)$$

$$\frac{d\omega}{dt} = v_0 t + \frac{\pi d^2 t^2 (4\mu_{ball} g d - 3C_D \mu_{air} v_0^2)}{24m} \quad (6)$$




Figure 10: Ground observer trajectories optimally tracking an aerial target.

As before, $P(\omega)$ was initially approximated by a randomly generated set of exemplars that were constrained to lie within the intersection of the sensors' error cones, and all of the particles were given equal weight. The distribution was then propagated using standard particle filtering techniques. At each time step t , deterministic drift was applied to $P(\omega)$ based on this dynamic model, followed by stochastic diffusion to account for model uncertainty. We sampled $P(\omega)$ at this point, as it allowed the robot pose ρ_{t+1} to be optimized over the expected target position ω_{t+1} - *not* over the current target position ω_t . That is, the robots moved so as to optimize their ability to localize the target at the next time instant based upon where they predicted the target would move.

In our simulations, robot motions were constrained by the maximum robot velocity $V_r \ll V_t$. This served to define the limits of the set over which the optimization occurs, Δ . Results from a sample Matlab simulation for three robots are provided below. For this trial, 100 exemplars were used to approximate $P(\omega)$, and the sensor model was assumed to be Gaussian noise with $\sigma = 1^\circ$.

Figure 10 (Extension 3) shows a representative simulation run of three robots tracking a single target. Robot trajectories are inefficient from a “distance-traveled” viewpoint, as they attempt to optimize position estimates over the target's entire flight rather than its endpoint. Figure 11 shows the error in measured target position for the same target trajectory from both stationary (dashed line) and moving (solid line) observers. When viewed in this light, the benefits of the otherwise curious robot trajectories become readily apparent. Reductions in measurement errors by a factor of 4-5 over the stationary case clearly demonstrate the effectiveness of the integrated optimization/dynamical modeling approach.




Figure 11: Measurement errors from stationary (dashed line) and moving (solid line) robot observers. Reductions in the latter case are significant across the entire target trajectory.

3.1.5 Tracking targets in a cluttered workspace

In the simulation results we have presented thus far, constraints to C_r were limited solely to pursuer dynamics and a mandatory target standoff distance. This is adequate for operations in an uncluttered workspace, but does not handle the more generic case where obstacles are present. To address the resulting additional constraints on C_r (and C_w), we assumed that the robots were able to obtain accurate information about obstacles in their immediate vicinity. This was consistent with our approach of generating locally optimal trajectories, and did not require *a priori* information of obstacle locations or a global map of the environment. Δ was then defined by the local obstacle-free configuration space.




Figure 12: Tracking a point target in a cluttered environment. Significant reductions to target position error were still realizable even in the presence of obstacles.

Next, we applied standard motion planning techniques for collision avoidance in this local neighborhood [16]. More specifically, the trajectory was modeled as the sum

of attractive and repulsive force vectors F_{att} and F_{rep} , respectively. F_{att} corresponded to the velocity vector \dot{q} generated from our optimization approach. We allowed local obstacles detected by the robot to impose a repulsive force vector F_{rep} onto this desired trajectory. The magnitude of F_{rep} was proportional to the robot velocities and inversely proportional to the distances from obstacles. The resultant force $F = F_{att} + F_{rep}$ represented the compromise robot trajectories as influenced by the presence of obstacles. This effectively constrained the optimization of $\rho \in \Delta$. A representative simulation trial can be found in Figure 12 (Extension 4).

While the presence of obstacles in this example constrained the robots' motion, the control law automatically adjusted their trajectories in order to compensate for these limitations and provide improved state estimates.

3.2 Experiments with the Clodbusters

The framework was implemented on our team of Clodbuster robots shown in Figure 1. These use omnidirectional vision for sensing, on-board Pentium III computers, PXC200 framegrabbers, and 802.11b wireless networking for inter-robot communications. In these experiments, a pair of pursuers tracked a third robot serving as a moving target. Two sets of trials were conducted to demonstrate operations in both cluttered and uncluttered environments.

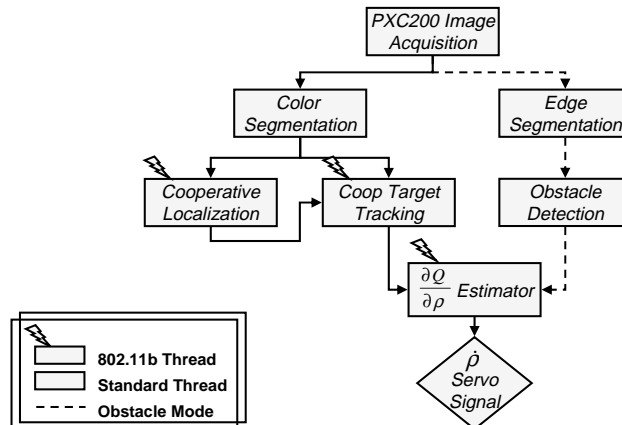


Figure 13: Experimental architecture used by each pursuer. The 802.11b blocks correspond to threads where cooperation - and as a result communication - with other robots is required.

The system architecture used by each pursuer robot is illustrated in Figure 13. It employed our *Live-Object* framework - a programming paradigm where objects encapsulate not only relevant algorithms and data, but also a thread within which these algorithms execute and communicate. The robots relied on YUV color segmentation to isolate one another in the camera images. Since our model assumes that the positions of the tracking robots are known, the pursuers communicated their relative tracking vectors to each other. This cooperation allowed each to estimate the other's relative position *and* orientation. With the pursuers localized, each then communicated its relative bearing to the target which was then cooperatively tracked. The complete localization process ran at 15Hz.

In the cluttered workspace trials, it was also necessary for each pursuer to estimate the position of obstacles. This was done by generating a rangemap to edge features in the environment as outlined in our previous work [17]. Target and pursuer robots were then discriminated from obstacles using their relative pose as estimated during the localization phase.

Experimental implementation followed closely with that used in the corresponding simulation experiment. Derivative estimation techniques were used to approximate the gradient of $Q(\rho)$ for optimizing the pursuers' headings. The maximum robot speed and a prescribed standoff distance served to define Δ for a given time-step. The sensor model assumed angular measurements obtained by the robots were corrupted with additive errors drawn from a normal distribution with a standard deviation of $\sigma = 0.5^\circ$. For the cluttered workspace trials, obstacles exhibited repulsive forces when the separation was less than 1 meter. Using 100 particles to approximate the probability $P(\omega)$ over the target configuration space, we were able to compute locally optimal robot configurations at a rate of 15Hz.



Figure 14: Trajectory for two pursuer robots tracking a moving target robot in an obstacle-free environment.

A representative trial from our obstacle-free experiments is shown in Figures 14 and 15 (Extension 5). The former shows a series of images from an overhead view of the scene, while the latter shows the corresponding position error estimates. Both the trajectory and the dramatic drop in the error estimate correlate well with the corresponding simulation results presented previously in Figure 5.

Figures 16 and 17 (Extension 6) show the corresponding trial for a cluttered workspace. The effect on the motion of the right pursuer robot was significant. In contrast to the obstacle-free case, its motion was constrained to a much narrower region. However, the control scheme automatically adjusted the path of the left pursuer to compensate for this limitation. As a result, the estimated target tracking error still fell dramatically.

It should again be noted that no explicit controllers were needed for maneuvering the formation. Trajectories were implicitly generated by $Q(\rho)$, which captured the notion of a good configuration.




Figure 15: Estimated RMS position error (cm) vs. time for the single target case.




Figure 16: Trajectory for two pursuer robots tracking a moving target robot in a cluttered workspace. The left pursuer adapts its trajectory to the right pursuer's mobility constraints.

4 Complexity Analysis

Referring to Algorithm 1, we can see that the computational complexity will depend heavily on the Est function used and the number of times $Q(\rho)$ must be computed by our $Q - optimizer$ function. Thus, the complexity is implementation dependent and can be tuned by the user based upon available computational resources and desired estimator robustness.

For our target tracking example, let m represent the number of robots, and n the number of targets. Our implementation employed a least-squares Est , function which runs in $O(m^3)$ time. Finite difference techniques approximated $\frac{\partial Q}{\partial \rho}$, and ran in $O(m)$ time for a total complexity of $O(m^4n)$. The framework lends itself to distributed computation. Taking advantage of this, the workload can be divided among the m




Figure 17: Estimated RMS position error (cm) vs. time for the single target case with obstacles. Results are comparable to the obstacle-free case.

robots so that the time complexity for each would scale linearly in the number of targets, and cubically in the number of robots.

Alternately, our *Est* function could employ a weighted average of the $O(m^2)$ stereo pair position estimates for each target. A straightforward weight factor is $w_{ij} = \frac{r_i r_j}{\sin \theta_{ij}}$, where r_i is the estimated distance from robot i to the target, and θ_{ij} is the vergence angle. This is a first-order approximation for stereo error from the determinant of the Jacobian. Using a constant number of simultaneous perturbations to estimate $\frac{\partial Q}{\partial \rho}$, this scheme results in an $O(m^2 n)$ complexity, or $O(mn)$ when distributed.

These are only two possible implementations, but they demonstrate the flexibility of our approach. Robust estimators that scale potentially exponentially could be used for small numbers of robots and/or targets, or more expedient estimation techniques could be applied when computational resources are at issue. The choice is left to the user's discretion.

5 Conclusions and Discussion

This paper presents an approach to the problem of controlling the configuration of a team of mobile agents so as to optimize the quality of the estimates derived from their measurements. The ideas were applied to target tracking tasks, where a team of robots was charged with optimally estimating the positions of a group of targets. This was demonstrated both in simulation and on robot platforms. Results indicate that this approach can solve fairly difficult sensor planning problems online without requiring excessive amounts of computational resources.

Our approach has several nice attributes. It provides an actual measurement reflecting the expected error in the estimated targets' positions. This is a side-effect of using a numerical representation for $P(\omega)$. Also, implicit rules for switching between robots-target assignments are embedded in the optimization of $Q(\rho)$, which negates the need for explicit switching controllers.

Perhaps the most attractive feature of the approach is its flexibility. It can be used with heterogeneous sensors, as the Est and $Q(\rho)$ functions abstract away specific sensor characteristics. It is also scalable in the number of robots and targets, and the computational effects of this scaling can be regulated by choosing appropriate estimation and optimization techniques.

In spite of these attributes, room for improvement still exists. The robots must combine measurements to estimate the position of targets. So while distributed computation is readily realizable, some level of centralization is inherent.

Our work to this point assumed that the robot positions ρ were known *a priori*, and there are means to estimate the formation pose as demonstrated in Section 3.2. However these estimates themselves are subject to uncertainty. We are currently reposing the problem within our framework as a simultaneous localization and target tracking task to address this.

The approach was applied to target tracking tasks in both open and cluttered workspaces. However, the latter was accomplished by merging with traditional motion planning techniques. As a result, it was subject to similar shortcomings (e.g. becoming trapped in local minima). Additionally, our work in cluttered environments only addressed issues relating to motion planning and not occluding obstacles. The latter topic the subject of ongoing research.

Lastly, to this point we have assumed a sensor model with an omnidirectional field of view (FOV). Adapting our approach to limited FOV sensors involves assimilating optimal assignment techniques with trajectory generation. This is also the topic of ongoing work.

Acknowledgments: This material is based upon work supported by the National Science Foundation under a CAREER Grant (Grant No. 9875867) and by DARPA under the MARS program.

Appendix

Index to multi-media Extensions The multi-media extensions to this article can be found online following the hyperlinks from *www.ijrr.org*.

References

- [1] R. Alur et al, "A framework and architecture for multirobot coordination," in *Proc. Seventh International Symposium on Experimental Robotics*, Honolulu, Hawaii, Dec. 2000.
- [2] J. Spletzer et al, "Cooperative localization and control for multi-robot manipulation," in *International Conference on Intelligent Robots and Systems*, Maui, Hawaii, Oct 2001.
- [3] R. Bajcsy, "Active perception," in *Proceedings of the IEEE, Special Issue on Computer Vision*, August 1988, vol. 76-8, pp. 996–1005.
- [4] E. Krotkov and R. Bajcsy, "Active vision for reliable ranging: Cooperating focus, stereo, and vergence," *International Journal of Computer Vision*, vol. 11-2, pp. 187–203, Oct 1993.

Extension	Media Type	Description
1	Video	Simulation of two robots optimally tracking an unpredictable point target.
2	Video	Simulation of four robots optimally tracking three unpredictable point targets.
3	Video	Simulation of three ground observers using a dynamical model to optimally track an aerial target.
4	Video	Simulation of tracking a point target in a cluttered workspace.
5	Video	Experimental trial with two pursuer robots tracking a third target robot in an obstacle-free workspace.
6	Video	Experimental trial with two pursuer robots tracking a third target robot in a cluttered workspace.

Table 1: Multi-media Extensions

- [5] S. Majumder, S. Scheduling, and H. Durrant-Whyte, “Multi-sensor data fusion for underwater navigation,” *Robotics and Autonomous Systems*, vol. 35, no. 1, pp. 97–108, 2001.
- [6] G. Dissanayake, P. Newman, H. Durrant-Whyte, and M. Csorba, “A solution to the simultaneous localization and map building,” *IEEE Transactions on Robotics and Automation*, vol. 17, no. 3, pp. 229–241, 2001.
- [7] S. Thrun, “A probabilistic online mapping algorithm for teams of mobile robots,” *International Journal of Robotics Research*, vol. 20, no. 5, pp. 335–363, 2001.
- [8] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, “Robust monte carlo localization for mobile robots,” *Artificial Intelligence*, vol. 128, no. 1-2, pp. 99–141, 2000.
- [9] L. Parker, “Cooperative robotics for multi-target observation,” *Intelligent Automation and Soft Computing*, 1999.
- [10] S. LaValle, H. Gonzalez-Banos, C. Becker, and J. Latombe, “Motion strategies for maintaining visibility of a moving target,” in *Proceeding of the IEEE Int. Conference on Robotics and Automation*, Albuquerque, NM, April 1997, pp. 731–736.
- [11] P. Fabiani, H. Gonzalez-Banos, J. Latombe, and D. Lin, “Tracking a partially predictable object with uncertainties and visibility constraints,” *Journal of Autonomous Robots*, 2001.
- [12] R. Pito, “A solution to the next best view problem for automated surface acquisition,” in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, October 1999, vol. 21.
- [13] I. Stamos and P. Allen, “Interactive sensor planning,” in *Computer Vision and Pattern Recognition Conference*, Santa Barbara, CA, June 1998, pp. 489–495.
- [14] M. Isard and A. Blake, “Condensation – conditional density propagation for visual tracking,” *International Journal of Computer Vision*, vol. 29-1, pp. 5–28, 1998.
- [15] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C*, Cambridge University Press, 1993.
- [16] J. Latombe, *Robot Motion Planning*, Kluwer Academic Publishers, 1991.

- [17] A. Das et al, “Real-time vision based control of a non-holonomic robot,” in *Proceeding of the IEEE Int. Conference on Robotics and Automation*, Seoul, Korea, May 2001, pp. 1714–1719.