A multiple UAV system for vision-based search and localization

John Tisdale†, Allison Ryan†,
Zu Kim†, David Törnqvist§, and J. Karl Hedrick†

Abstract— The contribution of this paper is an experimentally verified real-time algorithm for combined probabilistic search and track using multiple unmanned aerial vehicles (UAVs). Distributed data fusion provides a framework for multiple sensors to search for a target and accurately estimate its position. Vision based sensing is employed, using fixed downward-looking cameras. These sensors are modeled to include vehicle state uncertainty and produce an estimate update regardless of whether the target is detected in the frame or not. This allows for a single framework for searching or tracking, and requires non-linear representations of the target position probability density function (PDF) and the sensor model. While a grid-based system for Bayesian estimation was used for the flight demonstrations, the use of a particle filter solution has also been examined.

Multi-aircraft flight experiments demonstrate vision-based localization of a stationary target with estimated error covariance on the order of meters. This capability for real-time distributed estimation will be a necessary component for future research in information-theoretic control.

I. INTRODUCTION

Unmanned Aerial Vehicles (UAVs) have proven successful in information gathering tasks, such as target search or environmental sensing. To robustly estimate a target state, an information gathering system must incorporate models for uncertainty in the sensor, target model and a model for its own state. Robust estimation becomes more important as teams of agents use these estimates to plan their own motion, as in [1–3]. In the case where an agent has noisy sensors, or a limited sensor field of view, the use of probabilistic estimation becomes more important to develop accurate target estimates.

Traditionally, probabilistic 'search' problems represent the unknown state of a target by decomposing the search space into as many regions as is computationally feasible and updating each region as it is observed. 'Track' problems generally represent target state as a multivariate parameterized PDF (often Gaussian), but the choice of representation precludes incorporation of sensor readings that do not detect the target. In this work, we present a unified framework that accounts for 'search' by updating the target probability density function even when the target is not detected, but also achieves localization or tracking by updating an estimate using multiple observations of the target.

Campbell *et. al.* [4] demonstrate a UAV based system for geolocation; a sigma point filter is used to filter sensor data from a gimbaled camera to estimate target location. Low-level target tracking guarantees that the target will be observed at all times. The work presented here differs from the work of [4] in that the search for targets is addressed in addition to target localization. Additionally, this work uses a decentralized multi-sensor approach, in which sensor readings are fused across a team of vehicles.

Probabilistic maps for search have been used extensively; recently the authors in [5] used probabilistic maps for pursuit evasion games using unmanned helicopters. While search was explicitly addressed, localization was not examined in the framework of this paper.

The development of a unified framework for search and track was first proposed in [6]. Furukawa *et. al.* use a grid-based PDF to represent target position in the plane; this approach is effective for search, but falls short computationally once the target has been effectively localized. The precision of the estimate is limited by the resolution of the grid, and states with more than two dimensions (such as position and velocity in two dimensions) are difficult to model. In this work, a particle filter is used, in addition to a grid, in the hope that better localization results may be obtained in a search and localization framework, with similar computational effort.

The major contributions of this paper are the development and demonstration of a UAV based system for decentralized Bayesian search and localization. A grid-based Bayesian filter and a particle filter are used to maintain a target estimate. This system uses off the shelf hardware components to search for targets, and once detected, localize their position. High localization accuracy is achieved using low-cost UAVs by modeling the sensor uncertainty and fusing multiple observations using a Bayesian estimator.

II. DECENTRALIZED FILTERING FOR SEARCH AND TRACK

A decentralized estimator provides a method to update an estimate using multiple observations and a process model of the target state evolution. Recursive Bayes estimation is a method to update the filtering density $p(x_k|z_{1:k})$ using a sensor model $p(z_k|x_k)$, and a process model $p(x_{k+1}|x_k)$. The observation z_k is applied in the update step (1a) and the process model is applied in the prediction step (1b).

$$p(x_k|z_{1:k}) = \frac{p(z_k|x_k)p(x_k|z_{1:k-1})}{p(z_k|z_{1:k-1})} \quad (1a)$$

$$p(x_{k+1}|z_{1:k}) = \int p(x_{k+1}|x_k)p(x_k|z_{1:k})dx_k \quad (1b)$$

The sensor fusion problem is also easily solved in this framework using the sensor model $p(z_k|x_k)$ to relate the sensors. In our problem, the measurements (z_k^1 and z_k^2) from the two airplanes are independent so

$$p(z_k|x_k) = p(z_k^1, z_k^2|x_k) = p(z_k^1|x_k)p(z_k^2|x_k). \quad (2)$$

In the case where our distributions are Gaussian and the models are linear, (1) can be solved analytically and results in the well-known Kalman Filter [7]. For almost all other cases, analytical solutions cannot be obtained and approximations have to be used. The most common approximation is the Extended Kalman Filter (EKF) [8], that linearizes the system around the current state estimate and assume Gaussian distributions. To treat more complex distributions, the Unscented Kalman Filter (UKF) can be used [9]. This filter uses sigma points to represent the distribution undergoing a non-linear transform. While this approach has been used successfully in [4, 10], the EKF and UKF are unable to model highly multi-modal or skewed distributions. Other methods for treating arbitrary distributions with non-linear models are approximating the distributions spatially as a grid or with movable gridpoints (particles). In the grid-case, (1) is computed for each grid point. The case with particles result in the particle filter, see [11].

In this paper, where a search task is treated, the distributions are far from Gaussian. A grid-based or particle filter approach are therefore more suited, since they better captures the successive application of 'no detection' updates on a prior target distribution that might be uniform over a wide area.

Numerous approaches to data fusion have been explored by Durrant-Whyte *et. al.* [12, 13]. These techniques, such as the information filter and channel filter, often exploit the filtering technique to minimize the information communicated between nodes. In the work presented here, a parameterization of the likelihood function $p(z_k|x_k)$ is available and the parameters are distributed among the aircraft. The estimate is then updated using (1a). The likelihood function is given by one of two functional forms for the vision sensor: the case in which the target is observed at some camera coordinates and the case where it is not observed.

III. MODELING AND IMPLEMENTATION

An accurate model that captures the inherent uncertainty of an imperfect sensor is a crucial component of any real-world estimation system. Including observations in which the target is not detected enables probabilistic search and localization using all available information and requires sensor models for both the 'detect' (D) and 'no detect' ($\bar{D}$) cases. The 'detect' sensor model $p(D_k, z_k|x_k)$ is used to update the estimate when the target is detected at image coordinates $z_k = (u, v)$. The 'no detect' sensor model $p(\bar{D}_k|x_k)$ provides the update when the target is not detected in the frame. The time subscript will be dropped, assuming that all values correspond to the present time.

The sensor is modeled as a standard pinhole camera and the probability of detecting a target in the field of view is based on the image resolution at the target location. The visibility of a target in the field of view depends on the
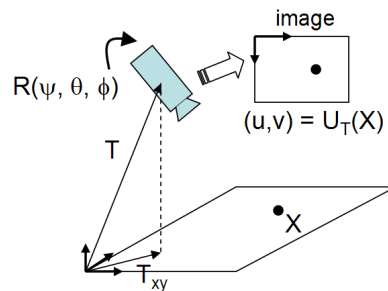

Fig. 1. Coordinate systems. There is a fixed coordinate system and an image coordinate system, the rotation between them is described by the Euler angles ψ, θ, ϕ.

aircraft position T and orientation defined by the Euler angle vector $q = [\psi\ \theta\ \phi]$. When the aircraft state is unknown, the visibility is a random variable which can be conditioned on the aircraft state estimates. We derive a probability of true detection, $p(D, V|x)$, and assume that the probability of false detection is constant per frame, $p(D, \bar{V}|x) = P_{false}$. We also assume that the location of a false detection is uniformly distributed over the image area (3).

$$
\begin{aligned}
p(D, z|x) &\\
&= p(z|x, D, V)p(D, V|x) + p(z|x, D, \bar{V})p(D, \bar{V}|x) \\
&= p(z|x, D, V)p(D, V|x) + (\text{image area})^{-1} P_{false} \quad (3)
\end{aligned}
$$

The detection likelihood $p(D, V|x)$ and localization likelihood $p(z|x, D, V)$ components of the model are derived separately.

A. Target Localization Model

The aircraft translation $T = [T_{XY}\ h]$ and Euler angle vector q define a transformation $U_T(x, T, q) : \mathbb{R}^2 \to \mathbb{R}^2$ from ground to image coordinates, as shown in Fig. 1. This transformation also depends on the intrinsic camera parameters such as focal length, which are estimated using standard calibration procedures, see e.g., [14]. If T and q are known, we expect to observe the target at $z = U_T(x, T, q)$, and so $p(z|x, D, V) = \delta_z(z - U_T(x, T, q))$, where δ_z is a two-dimensional delta function in image coordinates.

When the aircraft state is unknown, $p(z|x, D, V)$ requires marginalizing over state estimates (4). We assume that each of the Euler angles, altitude, and ground position are estimated independently, breaking down $p(T, q)$ by individual state estimates.

$$
\begin{aligned}
p(z|x, D, V) &= \int_{T, q} p(z|x, D)p(T, q) \\
&= \int_{T, q} \delta_z(z - U_T(x, T, q))p(T, q) \\
&= \int \delta_z(z - U_T(x, T, q))p(T_{XY}) \\
&\quad p(h)p(\psi)p(\theta)p(\phi)\, dT_{XY}\, dh\, d\psi\, d\theta\, d\phi \quad (4)
\end{aligned}
$$

For small estimation errors, the transformation $U_T(x, T, q)$ can be approximated by a first order Taylor expansion about

the state estimate, so that $P(z|x, D, V)$ for uncertain aircraft state becomes a series of one-dimensional convolutions. The partial derivatives of $U_T(x, T, q)$ with respect to translation are simply components of the rotation matrix, and derivatives with respect to the Euler angles are obtained by differentiating the rotation matrix.

When the UAV state estimation errors can be approximated as Gaussian or Gaussian mixture models (GMMs), the PDF $p(z|x, D, V)$ will be Gaussian (or GMM) in image coordinates z because the set of Gaussians is closed under convolution. However, the likelihood function on ground coordinates x will not be Gaussian or GMM due to the non-linear dependance on the Euler angles. This is similar to the effect seen in range-bearing sensors, where the likelihood function has a characteristic banana-like shape due to the bearing angle estimate uncertainty.

B. Target Detection Model

The probability of a true detection $p(D, V|x)$ is a function of the image resolution $r(z, T, q)$ at the target's image coordinates, $z = U_T(x, T, q)$, as well as the visibility of the target. The product of four step functions $s_i(z)$ along the edges of the image defines the visible region in image coordinates. For uncertain aircraft state, marginalization is required as in the previous section.

$$p(D, V|x) = \int_{T,q} p(D|V, x, T, q)p(V|x, T, q)p(T, q)$$

$$= \int_{T,q} p(D|V, r(z, T, q)) \prod_{i=1}^{4} s_i(z)p(T, q) \quad (5)$$

The dependence on resolution can be fitted to experimental data including true and missed detections using Bayes' rule (6). $p(D)$ is the detection rate for the data set and $p(r(z, T, q)|D)$ and $p(r(z, T, q)|\bar{D})$ can each be fitted as a parametric distribution (in this case Gaussian) from the data set.

$$p(D|r(z, T, q)) =$$
$$\frac{p(r(z, T, q)|D)p(D)}{p(r(z, T, q)|D)p(D) + p(r(z, T, q)|\bar{D})p(\bar{D})} \quad (6)$$

With the resolution and camera coverage expressed in image coordinates, the aircraft state uncertainties will be modeled independently as in the previous section and represented as a two-dimensional convolution mask in image coordinates. Therefore the marginalization in (5) corresponds to application of the convolution mask $M(\hat{T}, \hat{q})$. We convolve the aircraft state uncertainty with the nominal likelihood function $p(D|x; \hat{T}, \hat{q})$ based on the estimated state $(\hat{T}, \hat{q})$. Recall that detection likelihood function $p(D|x)$ includes both true and false detection.

$$p(D|x) = p(D, V|x) + p(D, \bar{V}|x)$$

$$= p(D|V, r(z, \hat{T}, \hat{q})) \prod_{i=1}^{4} s_i(z) * M(\hat{T}, \hat{q}) + P_{false} \quad (7)$$

TABLE I
DETECTION MODELING BASED ON CALIBRATION DATA SET.

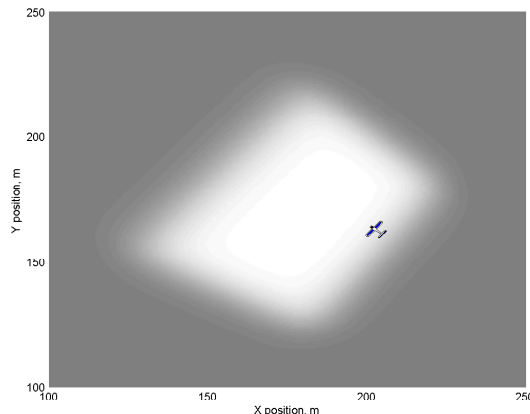|	
$p(res|D)$	$\mathcal{N}(60, 40.06)$
$p(res|\bar{D})$	$\mathcal{N}(60, 62.03)$
$P(D|V)$	0.96
P_{false}	0.004

Fig. 2. An example of a no-detect likelihood function. Dark areas indicate values close to 1, whereas the lighter areas indicate values closer to zero.

For Gaussian state estimates, the form of the detection likelihood is convolution of a Gaussian and a step function in two dimensions, resulting in cumulative Gaussian distributions. As these have no closed form equation, Sigmoid functions are used to approximate the image boundaries, where the Sigmoid function parameters may vary for best fit along each edge.

Fig. 2 shows an example no-detect likelihood function used in the flight experiments. Dark areas of the plot indicate values close to 1, whereas the lighter areas indicate values closer to zero. The effect of decreasing camera resolution may be seen in the upper right side of the field of view. The camera and aircraft parameters for this example are shown in Tables I and II, respectively. Representations such as a particle filter or grid based filter are better suited than EKF methods for this type of likelihood, which is obviously neither Gaussian or convex.

C. Target Model

In this work, it is assumed that only a single target is in the area of interest. Multiple targets have been consider in [15]. The primary difficulty in dealing with multiple targets is dealing with the problem of data association, or determining which sensor measurement comes from which target. In this work, we will not examine the data association problem, as there is already a rich literature on the topic. Also, the target is assumed to be static; no update step is necessary.

IV. EXPERIMENTAL RESULTS AND ANALYSIS

Flight tests of a two-aircraft team with fixed downward-looking cameras were performed at Camp Roberts, CA in August of 2007. The airframe and infrastructure have been previously described in [16, 17]. A series of flight tests were performed to verify sensor modeling, localization
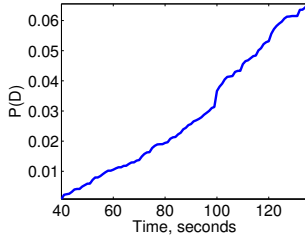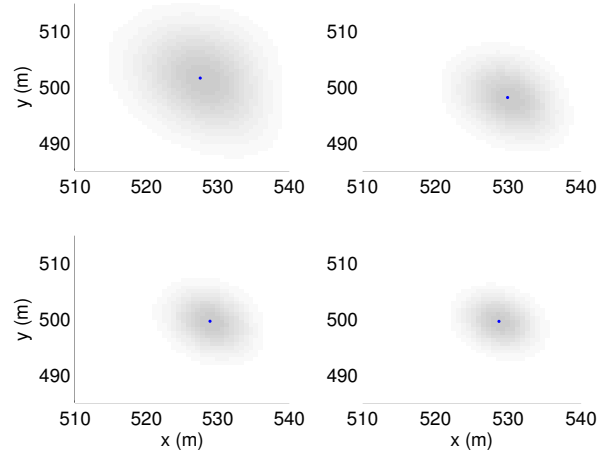
Fig. 3. Cumulative Probability of Detection.

Fig. 4. Evolution of filtering PDF over four consecutive observations, beginning with the first observation of the target. Dark area represents high probability, and the mean of each PDF is marked.

accuracy and system performance. Both grid-based Bayesian estimation and particle filtering were tested as representations for the estimated target state.

The grid representations were tested in flight, to update the target position in real-time. Real-time target position data was relayed over an 802.11b wireless link to a ground-based visualization GUI. The particle filter representation was tested post-flight, on the same data collected from the grid-based flight; while the particle filter was not run in real time, testing indicates that is fast enough to do so. Video was processed at roughly four frames per second.

For the sake of real-time data fusion, likelihoods were distributed between aircraft via an ad-hoc wireless network. Flight tests indicated that the wireless network was sufficient to ensure that all likelihoods were communicated between aircraft.

A. Grid-based Estimation

Fig. 5 shows the evolution of the searching process with dark areas representing high probability. Snapshots of the search density function and aircraft trajectories are shown during the 'search' portion of the mission, at 70 and 105 seconds, before the target has been detected. UAV 2 detects and localizes the target at roughly 120 seconds. For this flight experiment, the prior was uniform over a one kilometer by one kilometer square. Each grid cell has a width of 12.5 meters. In this flight, the system effectively localized the target to within the resolution of the filter; a red star is used to indicate target position when the target has been localized to a single grid cell. Clearly, the granularity of the grid is a limitation on the possible localization accuracy.

For the purposes of path planning for target search, [18] demonstrated that cumulative probability of detection, $p(D)$ is an appropriate metric for determining the a priori value of a search trajectory. When the target is not detected over a series of k steps, the probability of detection may be found by (8). Fig. 3 shows the cumulative probability of detection.

$$P(D) = 1 - \int p(x_k|z_{1:k} = \bar{D}) \, dx_k$$
$$= 1 - \int p(x_k) \prod_{j=1}^{k} p(z_j = \bar{D}|X) \, dx_k \quad (8)$$

Fig. 4 shows post processing of the same flight data using a grid cell size of 0.67 meters. Four consecutive filtering PDFs are shown, beginning with the first observation of the

target. Coverage of the entire search area at this grid density would require over two million cells, which is not feasible for real time performance on the current flight computer. However, this post processing provides an idea of the system performance that could be achieved with more computing power using the same sensors. Because the quantization error for the fine grid is greatly reduced, this implementation can be used as a benchmark for comparing the accuracy of the original grid and particle filter implementations.

B. Particle Filter Implementation

The particle filter used is the standard Sampling Importance Resampling (SIR) particle filter first introduced by [11]. For a more recent overview, see [19]. The SIR version of the particle filter resamples the particles after each time step, maintaining an equal weight for all particles. This makes the probability proportional to the density of particles in a certain area. The number of particles used is 6400, which is equal to the number of cells in the one square kilometer grid.

The initial probability is again uniform as in the previous section. The target position is modeled as a random walk according to

$$x_{k+1} = x_k + w_k, \quad (9)$$

with $\mathrm{Cov}(w_k) = \mathrm{diag}(1, 1)$ because the resampling step requires process noise. Fig. 5(b) shows the target location estimate using a particle filter.

C. Error Analysis and Comparison

The implemented grid cell size of 12.5 meters limits the localization accuracy through quantization effects, and cannot be significantly reduced using the current flight computer. The interpretation of the particle filter as an adaptive grid suggests that bias in the estimate due to quantization can be reduced without increasing computation cost by moving particles to represent areas of interest in greater detail. This is demonstrated in Fig. 5, where the particles concentrate
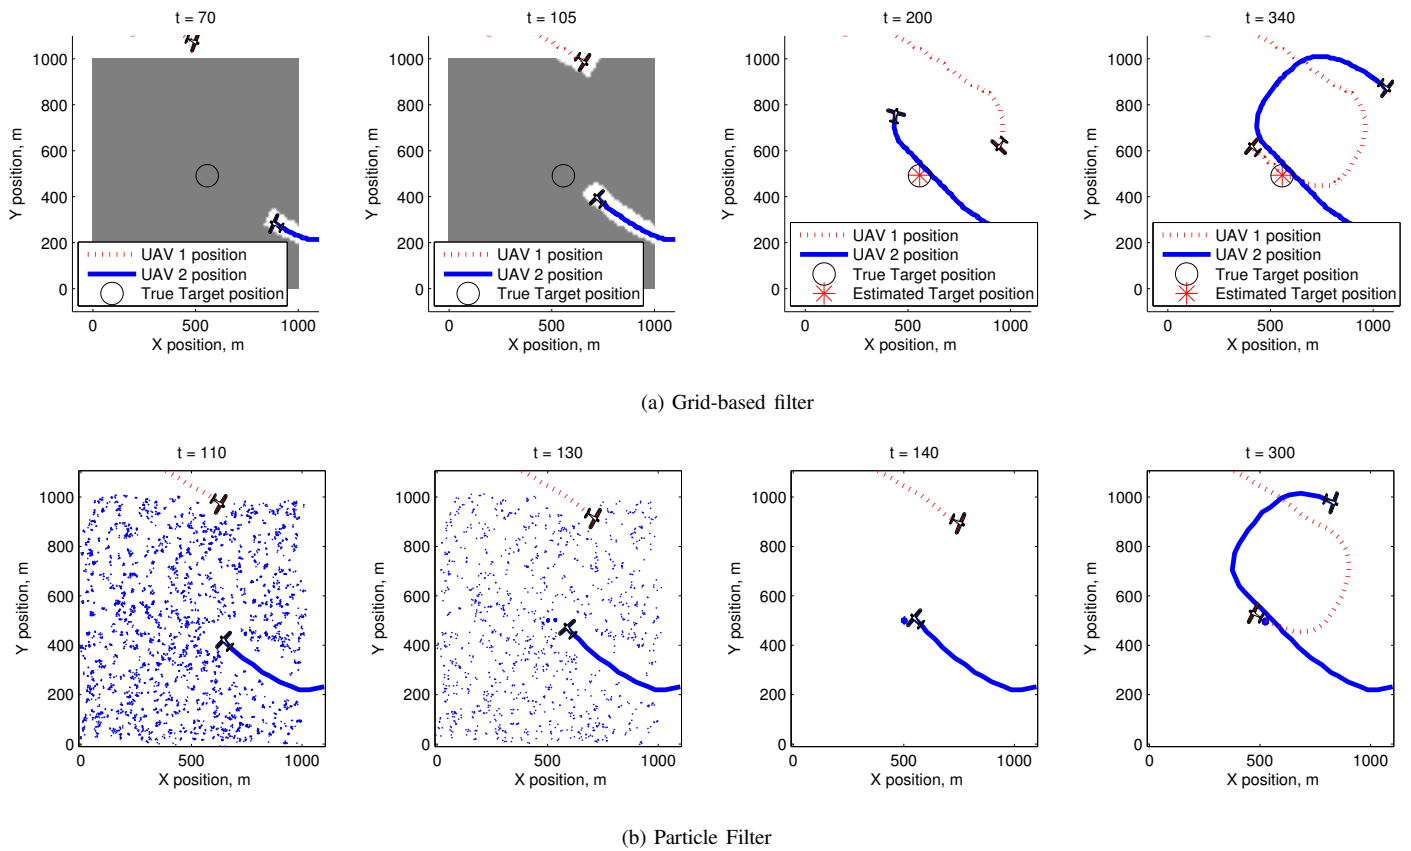
(a) Grid-based filter

(b) Particle Filter

Fig. 5.   (a) Evolution of the grid-based PDF throughout the search and localization process, with high probability shown as dark (b) Evolution of the particle filter throughout the search and localization process.

around the observed location of the target, and in Fig. 6 where the particle filter produces an estimate which closely corresponds to that of the post-processed high-resolution grid. Covariance ellipses corresponding to the filtering PDFs also show the effect of a large grid size.

Fig. 7 shows time evolution of a scalar measure of covariance for the three implementations. Again considering the fine grid as a benchmark, we observe that the estimate uncertainty decreases with the observations from each sensor and otherwise remains constant for the stationary target. The initial covariance for the fine grid represents a reduced search area. After the target is observed, the square kilometer grid always has higher uncertainty due to quantization. The covariance of the particle filter becomes very small when the target is observed, but then increases due to the required artificial process noise. This increased uncertainty is an artifact of the filter implementation and is a drawback of particle filtering of a stationary target. The uncertainty ellipse shown in Fig. 6 represents the minimum particle filter uncertainty, which is achieved immediately after the final observation.

In addition to quantization, aircraft state estimate uncertainty also contributes to the estimation error. To analyze the contribution of aircraft state uncertainty, we assume that the camera is aligned with the aircraft $\mathbf{e}_3$ (local down) axis and the target appears in the center of the image.

Therefore, the target position $\mathbf{x}$ lies along the $\mathbf{e}_3$ direction with $\mathbf{x} \cdot \mathbf{e}_3 = 0$, where $\{\mathbf{e}_i\}$ are fixed world coordinates. The partial derivatives of the resulting expression (10) for $\mathbf{x}$ result in the target position estimate uncertainties shown in Table II, based on linearization at a set of nominal parameter values.

The largest contributor to target estimate uncertainty is the UAV's position along the direction of motion, which has high uncertainty due to 4 Hz GPS updates. This uncertainty will be reduced in future work by filtering the GPS using angular data. Pitch and roll angle also contribute significantly, especially for larger nominal angles or low quality INS sensors.

$$
\begin{aligned}
\mathbf{x} = {} & \hat{T}_X \mathbf{e}_1 + \hat{T}_Y \mathbf{e}_2 + \\
& \frac{h}{\cos\phi\cos\theta}(\cos\phi\sin\theta\cos\psi + \sin\phi\sin\psi)\mathbf{e}_1 + \\
& \frac{h}{\cos\phi\cos\theta}(\cos\phi\sin\theta\sin\psi - \sin\phi\cos\psi)\mathbf{e}_2 \quad (10)
\end{aligned}
$$

V. CONCLUSIONS

This work presents a decentralized approach to target search and localization using low-cost UAVs and sensors. Using low-resolution fixed cameras, a team of UAVs can successfully search a significant area and localize a target with high accuracy by filtering multiple observations. This

TABLE II

ERROR MODELING FOR LOCALIZATION

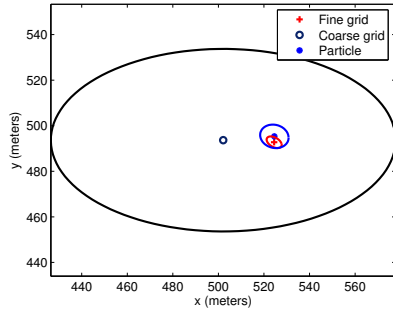| Parameter	nominal	σ	σ localization (m)
T_X (along-track)	n/a	10.5 m	10.5
T_Y (cross-track)	n/a	5.5 m	5.5
h	90 m	5.5 m	1.77
ψ	0 r	0.01 r	0.29
θ	0.17 r	0.05 r	4.65
ϕ	0.26 r	0.02 r	1.96

Fig. 6. Mean and covariance ellipse (sigma squared) for fine grid, coarse grid, and particle filter final estimates. Note that the particle filter representation achieves smaller bias error with computation cost similar to the coarse grid.
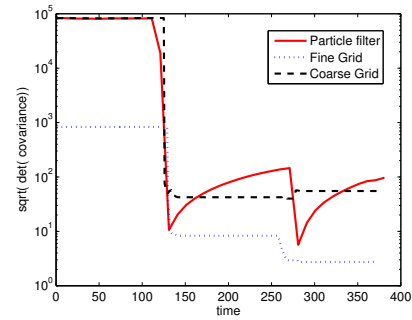

Fig. 7. Time evolution of covariance for fine grid, coarse grid, and particle filter. The particle filter covariance increases over time, due to a random walk target model.

work has demonstrated the importance of choosing an appropriate representation for the target PDF. The grid-based distribution is effective for the search problem, the particle filter approach was shown to produce superior localization under the same computational constraints.

REFERENCES

[1] G. Hoffmann, S. Waslander, and C. Tomlin, "Distributed cooperative search using information-theoretic costs for particle filters, with quadrotor applications," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, August 2006.
[2] B. Grocholsky, "Information driven coordinated air-ground proactive sensing," in *Proceedings of the IEEE international conference on robotics and automation*, April 2005.
[3] G. Mathews and H. Durrant-Whyte, "Decentralised optimal control for reconnaissance," in *In Proceedings of the Conference on Information, Decision and Control*, Adelaide, Australia, Feb 2007.
[4] M. E. Campbell and M. Wheeler, "A vision based geolocation tracking system for UAVs," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, August 2006.
[5] R. Vidal, O. Shakernia, H. J. Kim, H. Shim, and S. Sastry, "Multi-agent probabilistic pursuit-evasion games with unmanned ground and aerial vehicles," *IEEE Transactions on Robotics and Automation*, vol. 18, pp. 662–669, 2002.
[6] T. Furukawa, F. Bourgault, B. Lavis, and H. Durrant-Whyte, "Recursive bayesian search-and-tracking using coordinated uavs for lost targets," in *Proceedings of the IEEE Conference on Robotics and Automation*, Orlando, Florida, May 2006.
[7] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the American Society of Mechanical Engineering — Journal Basic Engieering*, vol. 82, no. Series D, pp. 35–45, Mar. 1960.
[8] T. Kailath, , A. H. Sayed, and B. Hassibi, *Linear Estimation*. Prentice Hall, 2000.
[9] S. J. Julier, "The scaled unscented transformation," in *Proceedings of American Control Conference*, vol. 6, Anchorage, AK, USA, May 2002, pp. 4555–4559.
[10] E. Frew, "Receding horizon control under uncertainty using optimal input design and the unscented transform," in *Proceedings of the IEEE conference on decision and control*, San Diego, CA, december 2006.
[11] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," *IEE Proceedings-F*, vol. 140, no. 2, pp. 107–113, Apr. 1993.
[12] B. P. Grocholsky, H. F. Durrant-Whyte, and P. W. Gibbens, "Information-theoretic approach to decentralized control of multiple autonomous flight vehicles," in *Proc. SPIE (International Society for Optical Engineering), Sensor Fusion and Decentralized Control in Robotic Systems III*, Oct 2000, pp. 348–359.
[13] B. Grocholsky, A. Makarenko, T. Kaupp, and H. Durrant-Whyte, "Scalable control of decentralised sensor platforms," in *2nd International Workshop on Information Processing in Sensor Networks*, 2003, pp. 96–112.
[14] B. Horn, *Tsai's camera calibration method revisited*, 2000.
[15] E. Wong, F. Bourgault, and T. Furukawa, "Multi-vehicle bayesian search for multiple lost targets," in *ICRA*. IEEE, 2005.
[16] A. Ryan, X. Xiao, S. Rathinam, J. Tisdale, D. Caveney, R. Sengupta, and J. Hedrick, "A modular software infrastructure for distributed control of collaborating unmanned aerial vehicles," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, August 2006.
[17] J. Tisdale, A. Ryan, M. Zennaro, X. Xiao, D. Caveney, S. Rathinam, J. K. Hedrick, and R. Sengupta, "The software architecture of the berkeley uav platform," in *Proceedings of the Conference on Control Applications*. IEEE, 2006.
[18] F. Bourgault, T. Furukawa, and H. Durrant-Whyte, "Optimal search for a lost target in a bayesian world," in *Proc. Int. Conf. on Field and Service Robotics*, 2003.
[19] M. Arulampalam, S. Maskel, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, Feb. 2002.