September 2006

Cooperative Air and Ground Surveillance

Ben Grocholsky
University of Pennsylvania

James Keller
University of Pennsylvania

R. Vijay Kumar
University of Pennsylvania, kumar@grasp.upenn.edu

George J. Pappas
University of Pennsylvania, pappasg@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/meam_papers

Recommended Citation

Grocholsky, Ben; Keller, James; Kumar, R. Vijay; and Pappas, George J., "Cooperative Air and Ground Surveillance" (2006). *Departmental Papers (MEAM)*. 74.
https://repository.upenn.edu/meam_papers/74

Copyright 2006 IEEE. Reprinted from *IEEE Robotics and Automation Magazine*, Volume 13, Issue 3, September 2006, pages 16-25.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/meam_papers/74
For more information, please contact repository@pobox.upenn.edu.

Cooperative Air and Ground Surveillance

Abstract

Unmanned aerial vehicles (UAVs) can be used to cover large areas searching for targets. However, sensors on UAVs are typically limited in their accuracy of localization of targets on the ground. On the other hand, unmanned ground vehicles (UGVs) can be deployed to accurately locate ground targets, but they have the disadvantage of not being able to move rapidly or see through such obstacles as buildings or fences. In this article, we describe how we can exploit this synergy by creating a seamless network of UAVs and UGVs. The keys to this are our framework and algorithms for search and localization, which are easily scalable to large numbers of UAVs and UGVs and are transparent to the specificity of individual platforms. We describe our experimental testbed, the framework and algorithms, and some results.

Comments

Copyright 2006 IEEE. Reprinted from *IEEE Robotics and Automation Magazine*, Volume 13, Issue 3, September 2006, pages 16-25.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Cooperative Air and Ground Surveillance

A Scalable Approach to the Detection and Localization of Targets by a Network of UAVs and UGVs

BY BEN GROCHOLSKY, JAMES KELLER, VIJAY KUMAR, AND GEORGE PAPPAS

Unmanned aerial vehicles (UAVs) can be used to cover large areas searching for targets. However, sensors on UAVs are typically limited in their accuracy of localization of targets on the ground. On the other hand, unmanned ground vehicles (UGVs) can be deployed to accurately locate ground targets, but they have the disadvantage of not being able to move rapidly or see through such obstacles as buildings or fences. In this article, we describe how we can exploit this synergy by creating a seamless network of UAVs and UGVs. The keys to this are our framework and algorithms for search and localization, which are easily scalable to large numbers of UAVs and UGVs and are transparent to the specificity of individual platforms. We describe our experimental testbed, the framework and algorithms, and some results.

Introduction

The use of robots in surveillance and exploration is gaining prominence. Typical applications include air- and ground-based mapping of predetermined areas for tasks such as surveillance, target detection, tracking, and search and rescue operations. The use of multiple collaborative robots is ideally suited for such tasks. A major thrust within this area is the optimal control and use of robotic resources to reliably and efficiently achieve the goal at hand. This article addresses this very problem of coordinated deployment of robotic sensor platforms.

Consider the task of reliably detecting and localizing an unknown number of features within a prescribed search area. In this setting, it is highly desired to fuse information from all available sources. It is also beneficial to proactively focus the attention of resources, minimizing the uncertainty in detection and localization. Deploying teams of robots working towards this common objective offers several advantages. Large environments preclude the option for complete sensor coverage. Attempting to increase coverage leads to tradeoffs

between resolution or accuracy and computational constraints in terms of required storage and processing. A scalable and flexible solution is therefore desirable.

In this article, we present our approach to cooperative search, identification, and localization of targets using a heterogeneous team of fixed-wing UAVs and UGVs. There are many efforts to develop novel UAVs and UGVs for field applications. Here, we assume standard solutions to low-level control of UAVs and UGVs and inexpensive off-the-shelf sensors for target detection. Our main contribution is a framework that is scalable to multiple vehicles and decentralized algorithms for control of each vehicle that are transparent to the specificity of the composition of the team and the behaviors of other members of the team. In contrast to much of the literature that addresses the very difficult planning problems for coverage and search (see [1], for example) and for localization (see [2], for example), our interests are in reactive behaviors that 1) are easily implemented; 2) are independent of the number or the specificity of vehicles; and 3) offer guarantees for search and for localization.

A key aspect of this work is the synergistic integration of aerial and ground vehicles that exhibit complementary capabilities and characteristics. Fixed-wing aircraft offer broad field of view and rapid coverage of search areas.

However, minimum limits on operating airspeed and altitude, combined with attitude uncertainty, place a lower limit on their ability to resolve and localize ground features. Ground vehicles, on the other hand, offer

UNMANNED AIR VEHICLES

high-resolution sensing over relatively short ranges, with the disadvantage of obscured views and slow coverage.

The use of aerial- and ground-based sensor platforms is closely related to other efforts to exploit the truly complementary capabilities of air and ground robots. Examples of such initiatives include the DARPA PerceptOR program [3] and Fly Spy project [4]. Pursuit-evasion strategies with ground vehicles and helicopters are described in [5]. The use of aerial-vehicle-mounted cameras or fixed ground cameras to guide ground robots is discussed in [6]. However, these approaches don't readily lend themselves to scaling up to large numbers or to tasks other than navigation. Further, none of these approaches incorporates the level of integration across aerial and ground vehicles that is captured here.

Our framework and algorithms are built on previous work in decentralized data fusion using decentralized estimation algorithms derived from linear dynamic models with assumptions of Gaussian noise [7]. We use the architecture proposed here and in [8]. In [9], we developed control algorithms that refine the quality of estimates, addressing both the detection and the localization problems. Our approach to active sensing and localization with UAVs and UGVs, briefly summarized in this article, is discussed in greater detail in [10]. Our work on scalable coordinated coverage with UAVs is also discussed in a previous paper [11].

This article is organized as follows. "Experimental Testbed" describes the demonstration system. "Framework for Scalable Information-Driven Coordinated Control" details the technical approach taken and system architecture. "Air-Ground Coordination" describes the application to our network of aerial and ground vehicles. We describe the characteristics of the UAV and UGV platforms and comparative qualities of feature observations from onboard cameras, deriving measurement uncertainty for features observed by vision sensors with uncertain state. These elements are combined and applied to an illustrative example of collaborative ground feature detection and localization. Concluding remarks follow.

Experimental Testbed

Figure 1 illustrates the UAVs in use at the GRASP Laboratory of the University of Pennsylvania. Each UAV consists of an airframe and engine, avionics package, onboard laptop, and additional sensing payload. We briefly describe the basic components of our UAVs and UGVs as well as the overall system architecture. Utilizing off-the-shelf airframe and autopilot components allows for effort to be directed at mission-level control schemes. A formation flight experiment is described in [12].

UAV Airframe and Payload

The airframe of each UAV is a quarter-scale Piper Cub J3 model airplane with a wingspan of 104 in (~ 2.7 m). The powerful glow fuel engine has a power rating of 3.5 hp, resulting in a maximum cruise speed of 60 kn (~ 30 m/s), at altitudes up to 5,000 ft ($\sim 1,500$ m), and a flight duration of 15–20 min.

The airframe-engine combination enables having significant scientific payload on board. Figure 1 shows pods that have been installed underneath each side of the wing containing high-resolution cameras and inertial measurement units (IMUs) as well as deployable sensors, beacons, and landmarks. More precisely, each UAV can carry the following internal and external payloads:

- ◆ onboard embedded PC
- ◆ IMU 3DM-G from MicroStrain
- ◆ external global positioning system (GPS): Superstar GPS receiver from CMC electronics, 10 Hz data
- ◆ camera DragonFly IEEE-1394 1024 $\times$ 768 at 15 frames/s from Point Grey Research
- ◆ custom-designed camera-IMU Pod includes the IMU and the camera mounted on the same plate. The plate is soft mounted on four points inside the pod. Furthermore, the pan motion of the pod can be controlled through an external-user PWM port on the avionics.

Figure 1. PennUAVs: Two Piper J3 Cub model airplanes fitted with external payload pods.

Figure 2. Multi-UAV and ground station functional architecture.

- ◆ custom-designed deployable Pod could be used to carry sensors, beacons, landmarks, or even robotic agents.

UAV Avionics and Ground Station

Each UAV is controlled by a highly integrated, user-customizable Piccolo avionics board, which is manufactured by CloudCap Technologies [13]. The avionics board comes equipped with the core autopilot, a sensor suite (which includes GPS), and an IMU consisting of three gyros, three accelerometers, and two pressure ports, one for barometric altitude and one for airspeed. A 40-MHz embedded Motorola MPC 555 Power PC receives the state information from all sensors and runs core autopilot loops at a rate of 20 Hz, commanding the elevator, ailerons, and rudder and throttle actuators as well as external-user payload ports.

Each UAV continuously communicates with the ground station. The communication occurs at 1 Hz and the range of the communication can reach up to 6 mi. The ground station performs differential GPS corrections and updates the flight plan, which is a sequence of three dimensional (3-D) waypoints connected by straight lines. The UAVs can also be commanded in a similar way from a supervisory controller (residing on board the UAV laptop), allowing further decentralization in the physical layer of the architecture (see Figure 2).

The ground station can concurrently monitor up to ten UAVs. Direct communication between UAVs can be emulated through the ground or by using the local communication channel on the UAVs (80211b—wireless network card). The ground station has an operator interface program (shown in Figure 3), which allows the operator to monitor flight progress, obtain telemetry data, or dynamically change the flight plans using georeferenced maps. Furthermore, the operator interface program can act as a server and enable multiple instances of the same software to communicate over a TCP/IP connection. This allows us to monitor or command and control the experiment in real time, remotely.

The UGV Platform

The ground vehicles, shown in Figure 4, are commercial four-wheel-drive model trucks modified and augmented with

Figure 3. Ground station operator Interface showing the flight plan and actual UAV position (August 2003, Fort Benning, Georgia).

Figure 4. Ground robot platforms.

onboard computers, stereo firewire cameras, GPS, and odometric and inertial sensors. Communication between ground vehicles and to the aerial platform base station is through an ad hoc 802.11b network.

Framework for Scalable Information-Driven Coordinated Control

In this section, we briefly discuss our framework for modeling and control that leads to an information-driven framework for the execution of multirobot sensing missions. We use the active sensor network (ASN) architecture proposed in [14]. The key idea is that the value of a sensing action is marked by its associated reduction in uncertainty. Mutual information [15] captures formally the utility of sensing actions in these terms. Dependence of the utility on robot and sensor state and actions allows us to formulate the tasks of coverage, search, and localization as optimal control problems.

Target Detection

Following this approach, detection and estimation problems are formulated in terms of summation and propagation of formal information measures. The feature-detection and feature-location estimation processes are now presented along with descriptions of the action utility, control strategy, and architecture network node structure.

We use certainty grids [16] as the representation for the search and coverage problems. The certainty grid is a discrete-state binary random field in which each element encodes the probability of the corresponding grid cell being in a particular state. For the feature detection problem, the state x of the i th cell C_i can have one of two values: *target* and *no target*. This is written as $s(C_i) = \{target | no\ target\}$. The information measure $\hat{y}_{d,i}(k|k)$, where subscript d denotes detection, stores the accumulated target detection certainty for cell i at time k

$$\hat{y}_{d,i}(k|k) = \log P(x) = \log P(s(C_i) = target). \quad (1)$$

Information associated with the likelihood of sensor measurements z —which, again, take one of two values, *target* or *no target*—is given by

$$\mathbf{i}_{d,s}(k) = \log P(z(k)|x). \quad (2)$$

The information measure that incorporates the current probabilities of detected targets is updated by the log-likelihood form of Bayes rule:

$$\hat{y}_{d,i}(k|k) = \hat{y}_{d,i}(k|k-1) + \sum_s \mathbf{i}_{d,s}(k) + C, \quad (3)$$

where C is a normalization factor.

Target Location Estimation

The coverage algorithm described above allows us to identify cells that have an acceptably high probability of containing features or targets of interest. The localization of features or targets is the second part of the task. This problem is

posed as a linearized Gaussian estimation problem. As in [9], the information form of the Kalman filter is used. New target location filters are instantiated as the detection process reaches a set threshold.

In this problem, we redefine the state vector $\mathbf{y}_f$ to be the coordinates of all the features detected by the target detection algorithm, with $\mathbf{y}_{f,i}$ denoting the (x, y) coordinates of the feature in a global coordinate system. Note that the target detection algorithm can run concurrently, updating the state vector with new candidate features and coordinates. The information filter maintains an information state vector $\hat{\mathbf{y}}_{f,i}(k|k)$ and matrix $\mathbf{Y}_{f,i}(k|k)$, distinguished by subscript f for each feature i , that relate to the feature estimate mean $\hat{\mathbf{x}}_{f,i}(k|k)$ and covariance $\mathbf{P}_{f,i}^{-1}(k|k)$ by

$$\hat{\mathbf{y}}_{f,i}(k|k) \triangleq \mathbf{P}_{f,i}^{-1}(k|k) \hat{\mathbf{x}}_{f,i}(k|k) \quad (4)$$

$$\mathbf{Y}_{f,i}(k|k) \triangleq \mathbf{P}_{f,i}^{-1}(k|k). \quad (5)$$

Each sensor measurement z contributes an information vector and matrix that captures the mean and covariance of the observation likelihood $P(z_s(k)|x) \sim \mathcal{N}(\mu_s, \Sigma_s)$:

$$\mathbf{i}_{f,s}(k) \triangleq \sum_s^{-1}(k) \mu_s(k), \quad \mathbf{I}_{f,s}(k) \triangleq \sum_s^{-1}(k). \quad (6)$$

The fusion of N_s sensor measurements with accumulated prior information is simply

$$\begin{aligned} \hat{\mathbf{y}}_{f,i}(k|k) &= \hat{\mathbf{y}}_{f,i}(k|k-1) + \sum_{j=1}^{N_s} \mathbf{i}_{f,j}(k) \\ \mathbf{Y}_{f,i}(k|k) &= \mathbf{Y}_{f,i}(k|k-1) + \sum_{j=1}^{N_s} \mathbf{I}_{f,j}(k), \end{aligned} \quad (7)$$

from which the state estimate for the i th target and the covariance associated with it can be easily recovered.

Decentralization, in the sense that nodes maintain local knowledge of aggregate system information, is made possible by the additive structure of the estimate update (3) and (7). This characteristic allows all nodes in a network to be updated through propagation of internodal information differences. A communications manager known as a *channel filter* implements this process at each interconnection [7].

Uncertainty Reducing Control

Equations (2)–(7) detail how sensing processes influence estimate uncertainty. An entropy-based measure [15] provides a natural quantitative measure of information in terms of the compactness of the underlying probability distributions. Mutual information measures the information gain to be expected from a sensor *before* making an observation. Most importantly, this allows a priori prediction of the expected information outcome associated with a sequence of sensing actions.

The control objective is to reduce estimate uncertainty. Because this uncertainty directly depends on the system state and action, each vehicle chooses an action that results in a

maximum increase in utility or the best reduction in the uncertainty. New actions lead to an accumulation of information and a change in overall utility. Thus, local controllers that direct the vehicle and sensors according to the mutual information gradient with respect to the system state are implemented on each robotic sensor platform. Analytic gradient expressions are available for the models used here in terms of the sensor quality, observer state, and estimate uncertainty. This is referred to as *information surfing* since the vehicles are, in essence, driven by information gain contours.

Scalable Proactive Sensing Network

The network of aerial and ground sensor platforms can now be deployed for searching for targets and for localization. Both the search and localization algorithms are driven by information-based utility measures and, as such, are independent of the source of the information, the specificity of the sensor obtaining the information, and the number of nodes that are engaged in these actions. Most importantly, these nodes automatically reconfigure themselves in this task. They are proactive in their ability to plan trajectories to yield maximum information instead of simply reacting to observations. Thus, we are able to realize a *proactive sensing network* with decentralized controllers, allowing each node to be seamlessly

aware of the information accumulated by the entire team. Local controllers deploy resources accounting for and, in turn, influencing this collective information. Coordinated sensing trajectories result that transparently benefit from complementary subsystem characteristics. Information aggregation and source abstraction result in nodal storage, processing, and communication requirements that are independent of the number of network nodes. The approach scales to indefinitely large sensor platform teams. This scalability is achieved at a potential performance cost through limiting control to local decision making. Alternative distributed and anonymous control approaches that seek global cooperation and assignment are pursued in [8].

Air-Ground Coordination

We have implemented our approach to active sensing on our network of robotic platforms described earlier. We present further detail of the sensing and control schemes used along with experimental results. The search and localization task consists of two components. First, detection of an unknown number of ground features in a specified search area $\hat{\mathbf{y}}_d(k|k)$. Second, the refinement of the location estimates for each detected feature $\mathbf{Y}_{f,i}(k|k)$. Feature observation uncertainty is investigated, confirming the complementary characteristics

of air and ground vehicles. Refinement of the location estimates requires the development of a reactive controller that is based on visual feedback. This is discussed followed by experimental results for a fixed UAV search pattern. Finally, a reactive controller for generating coordinated UAV search trajectories is presented.

Feature Observation Uncertainty

Figure 5 shows the onboard cameras, which are the primary air and ground vehicle mission sensors enabling detection and localization of features in operational environments. Figure 6 provides example images of a ground feature observed from air- and ground-based cameras. The accuracy of feature observations depends on the uncertain camera calibration and platform pose. A linearized error model developed in [11] is used here.

We consider points across the image and illustrate how their corresponding uncertainties on the ground plane vary. We also compare the uncertainties in ground feature localization using a UAV and UGVs. This comparison reveals the pros and cons of either platform and highlights the advantage of combining sensor information from these different sources for reliable localization.

Figure 5. Onboard cameras on the (a) UGVs and (b) UAVs.

Figure 6. A ground feature observed by (a) a UAV and (b) a UGV.

We simplify the target detection and localization problem by using colored targets and a simple color blob detection algorithm with our cameras. Figure 6 shows a typical 1.1-m $\times$ 1.4-m ground target as seen from a UAV and from a UGV. We use the geometric information specific to our sensor platforms. The UAV camera looks down from an altitude of 50 m, having a typical pitch angle of $\theta = 5^\circ$. The UGV cameras nominally look horizontally and are positioned 0.32 m above the ground plane. The variance in roll and pitch was estimated to be 4° while that for heading to be 25° . The variance in GPS coordinates is 25 m². Figure 7 displays ground feature position confidence ellipses associated with different points in the air and ground

imagery with these parameters. Thus, we can visualize how uncertainties in target localization vary across the field of view of the camera and from the different perspectives provided by the vehicles.

This comparison confirms the complementary character of the air and ground vehicles as camera platforms. Airborne cameras offer relatively uncertain observations over a wide field of view. Ground vehicles offer high relative accuracy that degrades out to an effective range of approximately 5 m. The ground vehicle field of view encompasses the aerial observation confidence region. This allows feature locations to be reliably handed off to ground vehicles, alleviating any requirement for ground vehicles to search for ground features.

Figure 7. Ground feature observation uncertainty for (a) air and (b) ground camera installations. The UAV camera looks down 5° off vertical at 50 m altitude. The UGV camera is mounted horizontally 0.32 m above the ground plane. Comparative feature observation accuracy is illustrated by ground plane confidence ellipses associated with uniformly spaced pixels in the imagery.

Figure 8. Ground vehicle information gain utility measure and iso-utility contours.

Figure 9. An illustration of the gradient controller for deployment of a UGV to localize a ground feature. The control law schedules the robot's heading Ψ to be in the direction of steepest mutual information gradient. The resulting trajectory and mutual information contours at two intermediate points are shown in (a) and (b). This strategy is referred to as information surfing since the platform is driven by the information gain contours.

Optimal Reactive Controller for Localization

Our proactive sensing network includes a reactive optimal controller that actively seeks to improve the quality of estimates of features or targets. While all vehicles (UAVs and UGVs) can run this controller, it is particularly relevant to the operation of UGVs, whose sensors are better equipped for precise localization. Therefore, we describe this controller and its implementation on the ground vehicles next. The controller is a gradient control law,

which automatically generates sensing trajectories that actively reduce the uncertainty in feature estimates by solving

$$\mathbf{u}_i(k) = \arg \max_{\mathbf{u} \in \mathcal{U}} \mathcal{I}_f(\mathbf{u}_i(k)), \quad (8)$$

where $\mathcal{U}$ is the set of available actions, and $\mathcal{I}_{f,i}(\mathbf{u}_i(k))$ is the mutual information gain for the feature location estimates given action $\mathbf{u}_i(k)$. For Gaussian error modeling of N_f features,

$$\mathcal{I}_f(\mathbf{u}_i(k)) = \sum_{j=1}^{N_f} \log \frac{|\mathbf{Y}_{f,j}(k|k-1) + \mathbf{I}_{f,i}(\mathbf{u}_i(k))|}{|\mathbf{Y}_{f,j}(k|k-1)|}. \quad (9)$$

This utility measure is illustrated in Figure 8. The controller involves forward motion at a fixed speed while choosing the steering velocity to enable heading toward the direction of steepest gradient. Figure 9 illustrates an example UGV trajectory generated by (8).

When implemented on a nonholonomic robot with constraints imposed on the vehicle turn rate and sensor field of view, this controller may result in the robot circling a feature while unable to make observations. To resolve this, the controller is disengaged when the expected feature location is within the turn constraint and outside the field of view, as illustrated in Figure 10.

Figure 10. Handling ground vehicle sensing field of view and control constraints. The controller is deactivated, allowing the vehicle to simply pursue the minimum turning radius (dotted trajectory), and reactivated when the direction of steepest descent is within the field of view (solid trajectory).

Experimental Results

Results are presented for an experimental investigation of a collaborative feature localization scenario. Three rectangular orange features, each measuring 1.1 m $\times$ 1.4 m, were placed in a 50-m $\times$ 200-m search area. Figure 3 details a typical UAV trajectory generated to cover a search area in multiple

passes. The elapsed time for each pass was approximately 100 s. A sequence of images captured from an altitude of 65 m is shown in Figure 11. The feature estimates are seamlessly made available to all vehicles.

Figure 12 illustrates the initial feature uncertainty and the trajectory taken by the ground vehicle to refine the quality of these estimates. Detailed snapshots of the active sensing process are shown in Figure 13. These indicate

Figure 11. Aerial images of the test site captured during a typical UAV flyover at 65 m altitude. Three orange ground features highlighted by white boxes are visible during the pass.

Figure 12. Figures indicating (a) initial feature confidence and UGV active sensing trajectory and (b) σ_x and σ_y components of feature estimate standard deviation over time.

A key aspect of this work is the synergistic integration of aerial and ground vehicles that exhibit complementary capabilities and characteristics.

the proposed control scheme successfully positioning the ground vehicle to take advantage of the onboard sensor characteristics.

It is important to note the performance benefit obtained through collaboration. Assuming independent measurements, in excess of 50 passes (about 80 min of flight time) are required by the UAV to achieve this feature estimate certainty. It would take in excess of half an hour for the ground vehicle with this speed and sensing range to cover the designated search area and achieve a

Figure 13. Snapshots of the active feature location estimate refinement by an autonomous ground robot equipped with vision, GPS, and inertial and odometric sensors. This corresponds to the second feature indicated in Figure 12(a). The initial confidence region obtained through aerial sensing alone is indicated in (a). Any need for an extensive search by the ground vehicle is alleviated since this confidence region is slightly smaller than the ground vehicle onboard camera effective field of view. Compounded error sources in the ground vehicle sensor system result in feature observations that provide predominantly bearing information as shown in (b)–(c). The controller successfully drives the ground robot to sensing locations orthogonal to the confidence ellipse major axis that maximize the expected reduction in estimate uncertainty. False feature detections are rejected as indicated in (d).

high probability of detecting the features. The collaborative approach of using aerial cues to active ground sensing completes this task in under 10 min. Thus, the proactive sensing network has a performance level well in excess of the individual system capabilities.

Coordinated Multi-UAV Area Coverage

In the previous experiment, the UAV trajectory followed a fixed search pattern selected to provide coverage of the target area. In this section, the reactive information-gathering control concept previously used for UGV feature localization is applied to generate online UAV trajectories for search area coverage. Summation over the sensor field of view captures the value of performing a sensing action at a given location. Directing the UAV camera platforms according to the gradient of this utility measure provides a reactive scheme for area coverage. Coordinated trajectories arise due to coupling through accumulated measurement information, without knowledge of the state of other UAVs. The information objective drives platforms apart and towards unexplored regions as detailed in Figure 14.

Concluding Remarks

We presented our experimental testbed of aerial and ground robots and a framework and a set of algorithms for coordinated control with the goal of searching for and localizing targets in a specified area. While many details were omitted because of space constraints, the details of the hardware and software integration are presented in [10] and [12]. The methods described here lend themselves to decentralized control of heterogeneous vehicles without requiring any tailoring to the specific capabilities of the vehicles or their sensors. The unique features of our approach are as follows. First, the methodology is transparent to the specificity and the identity of the cooperating vehicles. This is because vehicles share a common representation, consisting of a certainty grid that contains information about the probability of detection of targets and an information vector/matrix pair that is used in the information form of the Kalman filter. Observations are propagated through the network, changing both the certainty grid and the information vector/matrix. Second, the

The coverage algorithm allows us to identify cells that have an acceptably high probability of containing features or targets of interest.

computations for estimation and control are decentralized. Each vehicle chooses the action that maximizes the utility, which is the combined mutual information gain from

Figure 14. Two snapshots in time (a) and (b) of the UAV trajectories and information measures relating to the probability of ground features remaining undetected over a specified search area. Each has four displays indicating the detection probability entropy, extent of coverage, current information gain, and observation information. The UAVs are directed by the information gain gradient. Each UAV camera is considered to have an 80% detection probability over its projected field of view. The task is terminated when 99.9% detection confidence is reached.

onboard sensors towards the detection and localization processes. Finally, the methodology presented here is scalable to large numbers of vehicles. The computations scale with the dimensionality of the representation and are independent of the number of vehicles. Our experiments demonstrate the performance benefit obtained through collaboration and illustrate the synergistic integration of ground and aerial nodes in this application.

Acknowledgments

This work was in part supported by DARPA MARS NBCH1020012, ARO MURI DAAD19-02-01-0383, and NSF CCR02-05336. The authors would like to acknowledge Daniel Gomez Ibanez, Woods Hole Oceanographic Institute, and Selcuk Bayraktar, Massachusetts Institute of Technology, for their help in deploying the UAVs.

Keywords

Decentralized collaborative control, heterogeneous robot teams, active perception.

References

- [1] H. Choset, K.M. Lynch, S. Hutchinson, G.A. Kantor, W. Burgard, L.E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*. Cambridge, MA: MIT Press, June 2005.
- [2] N. Roy, W. Burgard, D. Fox, and S. Thrun, "Coastal navigation – Mobile robot navigation with uncertainty in dynamic environments," in *Proc. IEEE Conf. Robotics Automation (ICRA)*, May 1999, pp. 35–40.
- [3] T. Stentz, A. Kelly, H. Herman, P. Rander, and M.R., "Integrated air/ground vehicle system for semi-autonomous off-road navigation," in *Proc. AUVSI Symp. Unmanned Systems*, 2002.
- [4] R. Vaughan, G. Sukhatme, J. Mesa-Martinez, and J. Montgomery, "Fly spy: lightweight localization and target tracking for cooperating ground and air robots," in *Proc. Int. Symp. Distributed Autonomous Robot Systems*, 2000, pp. 315–324.
- [5] R. Vidal, O. Shakernia, H.J. Kim, D.H. Shim, and S. Sastry, "Probabilistic pursuit-evasion games: Theory, implementation and experimental evaluation," *IEEE Trans. Robot. Automat.*, vol. 18, no. 5, pp. 662–669, Oct. 2002.
- [6] R. Rao, V. Kumar, and C. Taylor, "Experiments in robot control from uncalibrated overhead imagery," in *Proc. 9th Int. Symp. Experimental Robotics (ISER'04)*, 2004.
- [7] J. Manyika and H. Durrant-Whyte, *Data Fusion and Sensor Management: An Information-Theoretic Approach*. Englewood Cliffs, NJ: Prentice Hall, 1994.
- [8] B. Grocholsky, "Information-theoretic control of multiple sensor platforms," Ph.D. dissertation, Univ. Sydney, 2002 [Online]. Available: <http://www.acfr.usyd.edu.au>
- [9] B. Grocholsky, A. Makarenko, T. Kaupp, and H. Durrant-Whyte, "Scalable control of decentralised sensor platforms," in *Proc. Information Processing Sensor Networks: 2nd Int. Workshop, IPSN03*, 2003, pp. 96–112.
- [10] B. Grocholsky, S. Bayraktar, V. Kumar, C. Taylor, and G. Pappas, "Synergies in feature localization by air-ground robot teams," in *Proc. 9th Int. Symp. Experimental Robotics (ISER'04)*, 2004, pp. 353–362.
- [11] B. Grocholsky, R. Swaminathan, J. Keller, V. Kumar, and G. Pappas, "Information driven coordinated air-ground proactive sensing," in *Proc. IEEE Int. Conf. Robotics Automation (ICRA'05)*, 2005, pp. 2211–2216.
- [12] S. Bayraktar, G. Fainekos, and G.J. Pappas, "Experimental cooperative control of fixed-wing UAVs," in *Proc. 43rd IEEE Conf. Decision and Control*, 2004, pp. 4292–4298.
- [13] B. Vaglienti and R. Hoag, *Piccolo System User Guide*, 2003 [Online]. Available: <http://www.cloudcaptech.com/downloads.htm>
- [14] A. Makarenko, A. Brooks, S. Williams, H. Durrant-Whyte, and B. Grocholsky, "An architecture for decentralized active sensor networks," in *Proc. IEEE Int. Conf. Robotics Automation (ICRA'04)*, New Orleans, Louisiana, 2004, pp. 1097–1102.
- [15] T. Cover and J. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [16] A. Makarenko, S. Williams, and H. Durrant-Whyte, "Decentralized certainty grid maps," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems (IROS)*, 2003, pp. 3258–3263.

Ben Grocholsky is a project scientist at the Robotics Institute, Carnegie Mellon University. He received his Ph.D. from the University of Sydney in 2002 and was previously a postdoctoral researcher at the University of Pennsylvania GRASP Laboratory. His research spans active sensor networks, decentralized cooperative control, and unconventional operator interfaces.

James Keller joined the University of Pennsylvania GRASP Laboratory in 2002 as a project engineer and Ph.D. student. Before this, he enjoyed a 20-year career in the helicopter industry with the Boeing Company. He is currently working on autonomous path planning for aerial and underwater vehicles.

Vijay Kumar received his M.Sc. and Ph.D. in mechanical engineering from the Ohio State University in 1985 and 1987, respectively. He has been on the faculty in the Department of Mechanical Engineering and Applied Mechanics with a secondary appointment in the Department of Computer and Information Science at the University of Pennsylvania since 1987. He is currently the UPS Foundation professor and chair of mechanical engineering and applied mechanics. His research interests lie in the area of robotics and networked multiagent systems. He is a Fellow of the IEEE and ASME.

George Pappas received his Ph.D. degree from the University of California at Berkeley in 1998. In 2000, he joined the University of Pennsylvania Department of Electrical and Systems where he is currently an associate professor and director of the GRASP Laboratory. He has published over 100 articles in the areas of hybrid systems, hierarchical control systems, distributed control systems, nonlinear control systems, and geometric control theory, with applications to flight management systems, robotics, and unmanned aerial vehicles.

Address for Correspondence: Prof. Vijay Kumar, Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, 3330 Walnut St., Levine Hall, GRW 470, Philadelphia, PA 19104 USA. E-mail: kumar@grasp.upenn.edu.