

DEPLOYING AIR-GROUND MULTI-ROBOT TEAMS IN URBAN ENVIRONMENTS

L. Chaimowicz, A. Cowley, D. Gomez-Ibanez, B. Grocholsky, M. A. Hsieh, H. Hsu, J. F. Keller, V. Kumar, R. Swaminathan, C. J. Taylor
GRASP Laboratory – University of Pennsylvania
Philadelphia – PA – USA

Abstract We present some of the work performed in the GRASP Laboratory with the objective of deploying multi-robot teams in urban environments. Specifically, we focus on three important issues in this type of mission: the development of tools for providing situational awareness, the use of air and ground vehicles for cooperative sensing and the construction of radio maps to keep team connectivity. We describe the main approaches that we have been using for tackling these issues and present some preliminary results from experiments conducted with our team of air and ground vehicles.

Keywords: Air-ground cooperation, situational awareness, radio mapping.

1. Introduction

Urban environments provide unique challenges for the deployment of multi-robot teams. In this type of environment, buildings pose 3-D constraints on visibility, communication network performance is difficult to predict and GPS measurements can be unreliable or even unavailable. We believe that a network of aerial and ground vehicles working in cooperation can have a better performance in these type of environments. By constructing a three-dimensional sensing network, teams of air and ground vehicles can obtain better and more complete information and thus be more robust to the challenges posed by these environments. For this, it is necessary to keep the network tightly integrated at all times since vehicles have to support each other in order to function with synergy. Also, it is important to provide ways for a human operator to command the whole network, and not individual vehicles.

In this paper, we present some of the efforts that have been done by the GRASP Laboratory – University of Pennsylvania for deploying teams of air and ground vehicles in urban environments as part of the MARS2020 project. Sponsored by DARPA, this project is focused on the development of critical technologies required to realize network-centric control of heterogeneous platforms that is strategically responsive, survivable and sustainable for reconnaissance, surveillance or search and rescue type missions. In this endeavor, the University of Pennsylvania is teamed with the Georgia Institute of Technology, the University of Southern California and BBN Technologies.

At the University of Pennsylvania, our MARS2020 research thrust is to establish the overall paradigm, modeling framework and the software architecture to enable a minimum number of human operators to manage a heterogeneous robotic team with varying degrees of autonomy. The central features of our approach are to organize the robotic platforms for network centric autonomous operations; to develop small team communication-sensitive behaviors, which allow robots to perform alongside humans as full team members; to enable the team to learn and adapt to changing terrain conditions that may impact communication network performance and localization information (e.g., GPS, line of sight sensing, etc.); and to develop computationally distributed strategies to provide an unprecedented level of situational awareness. The effort will result in an integrated team of UAVs and UGVs, in which the team and the network will adapt to the needs and commands of a remotely located human operator to provide situational awareness.

This paper is organized as follows: the next section presents our multi-robot team. Section 3 describes some of the basic capabilities of our team in terms of localizing and navigating in urban terrains. Sections 4 to 6 discuss our main research thrusts, namely: situational awareness, air-ground cooperation and cooperative radio mapping. Finally, section 7 brings the conclusion of this paper.

2. Hardware and Software Testbed

Our multi-robot team consists of 5 unmanned ground vehicles (UGVs), 2 fixed wing aircraft and a blimp (Figure 1). The UGVs employ a commercially available, radio controlled scale model truck with suspension and chassis modified for autonomous duty. The chassis is approximately 480 mm long and 350 mm high. Mounted in the center of the chassis is a Pentium III laptop computer. Each UGV contains a specially designed Universal Serial Bus (USB) device which controls drive motors, odometry, steering servos and a camera pan mount with input from the

PC. A GPS receiver is mounted on the top of an antenna tower, and an inertial measurement unit (IMU) is mounted between the rear wheels. A forward-looking stereo camera pair is mounted on a pan mount which can pivot 180 degrees to look left and right. A small embedded computer with 802.11 wireless Ethernet handles network connectivity. This junction box (JBox), developed jointly by the Space and Naval Warfare Systems Center, BBN Technologies, and the GRASP Lab, handles multi-hop routing in an ad-hoc wireless network. The UGV is powered by one or two hot swappable lithium polymer battery packs, each with 50 watt hour capacity.

The fixed wing aircraft are quarter scale Piper Cub model aircraft equipped with the Piccolo autopilot by Cloud Cap Technology. The autopilot provides innerloop attitude and velocity stabilization control allowing research to focus on guidance for mission level tasks. In addition to the sensors within the autopilot, the air vehicles carry a sensor pod containing a high resolution firewire camera, inertial sensors and a 10Hz GPS receiver. A spread-spectrum radio modem is used for communications between air vehicles and the operator base station. Ground based system components communicate through an Ad-Hoc 802.11b network. We also have a medium sized blimp (9 meter length) that has nearly a 3 kg payload for research equipment. It is equipped a GPS, an inertial measurement unit capable of sensing rates and attitudes, a video camera that can be slewed in azimuth and elevation, and the onboard computing and communication hardware to be autonomous but also capable of being dynamically redirected by a remote human operator.

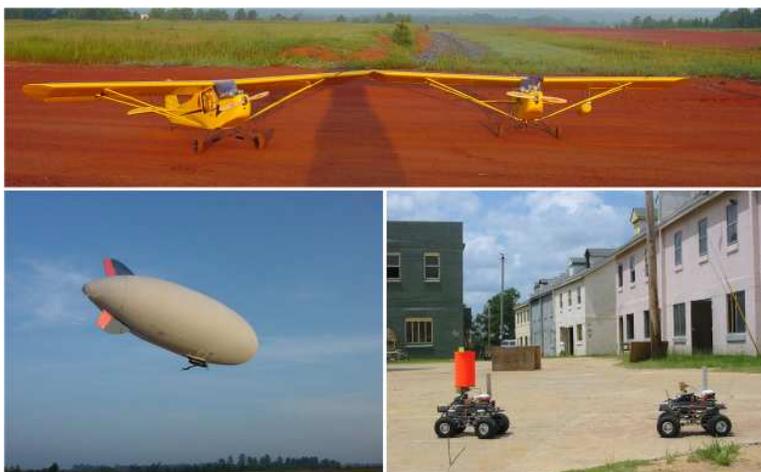


Figure 1. Multi-robot team composed of air and ground vehicles.

We are using ROCI (Remote Object Control Interface) (Chaimowicz et al., 2003; Cowley et al., 2004) for programming and tasking both the ground and air vehicles. ROCI is a high level operating system useful for programming and managing networks of robots and sensors. In ROCI, each robot is considered a node which contains several processing and sensing modules and may export different types of services and data to other nodes. Each node runs a kernel that mediates the interactions of the robots in a team. The kernel is responsible for handling program allocation and injection, managing the network and maintaining an updated database of other nodes in the ROCI network. The control functionality needed by such a kernel is made possible by self-contained, reusable modules. Each module encapsulates a process which acts on data available on its inputs and presents its results on well defined outputs. Thus, complex tasks can be built by connecting inputs and outputs of specific modules. A ROCI task is a way of describing an instance of a collection of ROCI modules to be run on a single node, and how they interact at runtime. It is defined in an XML file which specifies the modules that are needed to achieve the goal, any necessary module-specific parameters, and the connections between these modules. At runtime, these connections are made through a pin architecture that provides a strongly typed, network transparent communication framework.

3. Localization and Navigation

Two of the key requirements for the robots is the ability to localize themselves and navigate in urban environments. A Kalman filter framework is employed to estimate robot localization. Prediction is driven by wheel encoder odometry and inertial measurements from a low cost IMU. Appropriate observation models allow various sources of position information to be incorporated. These include on-board GPS, robot observations from external vision sensors and landmarks observed by the on-board camera.

Our robots navigate based on a list of desired waypoints. Each waypoint is a pair of georeferenced coordinates that specify a destination point for the robot. The waypoints can be specified manually through a user interface or can be input directly to the navigation module by other modules and tasks. Sometimes it may be necessary to automatically generate a list of intermediary waypoints between the robot current position and the desired destination, so that the robots will follow an specific path to their goal. One way of doing this is to create a graph based on a Voronoi Diagram of the environment and use it as a roadmap for planing the intermediary waypoints. The Voronoi Diagram can be

generated beforehand, using overhead imagery obtained by the air vehicles. Another possibility is to use “mission scripts”, which will be discussed in the next section.

A trajectory controller generates linear and angular velocities for the robot based on its current position and the next desired waypoint. A robot considers a waypoint reached when the distance between them is less than a threshold ϵ . Local obstacle avoidance is accomplished through the use of the robot’s stereo vision system. Images captured simultaneously from the two cameras are used to generate a medium-density depth map through a multi-pass process of confidence adjustment. This depth map is converted to a two-dimensional occupancy grid centered on one of the cameras in the stereo setup. Several trajectory arcs, corresponding to various vehicle movements (e.g. turn left, go straight, or turn right) can be compared against this grid to verify if a collision would result. The use of a finite number of discrete trajectories, rather than a more complete shortest path solver, lets the system run at the rate of the camera with very little processor overhead.

Figure 2 shows the trajectory performed by a robot following a sequence of waypoints in the MOUT (Military Operations on Urban Terrain) site at Fort Benning, the main test site for this project. The MOUT site is a replica of a small city consisting of 17 two and three store buildings, streets and access roads. It is configured with cameras that allow a multiple view tracking of training missions. It also features a small airfield, making it a suitable test ground for air-ground cooperation.



Figure 2. Trajectory of a robot navigating using a sequence of waypoints. The waypoints are depicted as ‘*’ and the robot executes the sequence twice.

4. Situational Awareness

In our framework, the main interface between a human operator and the robot team is the ROCI Browser. The browser displays the multi-robot network hierarchically: the human operator can browse nodes on the network, tasks running on each node, the modules that make up each task, and pins within those modules. The browser’s main job is to give a user command and control over the network as well the ability to retrieve and visualize information from any one of the distributed nodes.

Using the browser, the user can start and stop the execution of tasks in the robots remotely, change task parameters or send relevant control information for the robots. Elaborated missions can be constructed using *scripts*. Mission scripts can be generated online or offline, and specify a sequence of actions that should be performed by a team member. For example, capturing panoramic images at different waypoints, or navigating through multiple intermediate waypoints before reaching a target site. A synchronization mechanism allows for coordinated efforts between multiple robots, and a success/failure condition check on the outcome of each action makes limited branching possible. We are currently utilizing these capabilities to support a multi-robot signal strength mapping mission with intelligent recovery behavior if at any point any of the robots lose radio connectivity to the other members. This specific mission will be discussed in Section 6.

The visualization and exploitation of data generated by the multi-robot team is one of the main features of our situational awareness framework. To access and visualize the data, a human operator interacts with the ROCI Browser, part of which is a display generation runtime made up of a collection of plug-in Display Modules that convert incoming pin data to raster images. Data can be retrieved from different nodes equipped with various types of sensors such as GPS receivers, IMU readers, cameras, etc, and can be combined to give the user a rich view of the mission.

Figure 3 shows a snapshot of the browser during the execution of one of our experiments in Ft. Benning. It displays an image of the MOUT site taken previously from one of the UAVs overlaid with different information acquired from the robots. For example, the thick lines represent signal strength of the radio connection between the five nodes in our network. One robot can be seen in the center of the image surrounded by a octagon that indicates the uncertainty of its localization. As mention in Section 3 the robot position is estimated by fusing information from several sources, in this case an on board GPS, odometry, and an

external overhead camera. If desired, these individual observations can also be displayed on the browser.

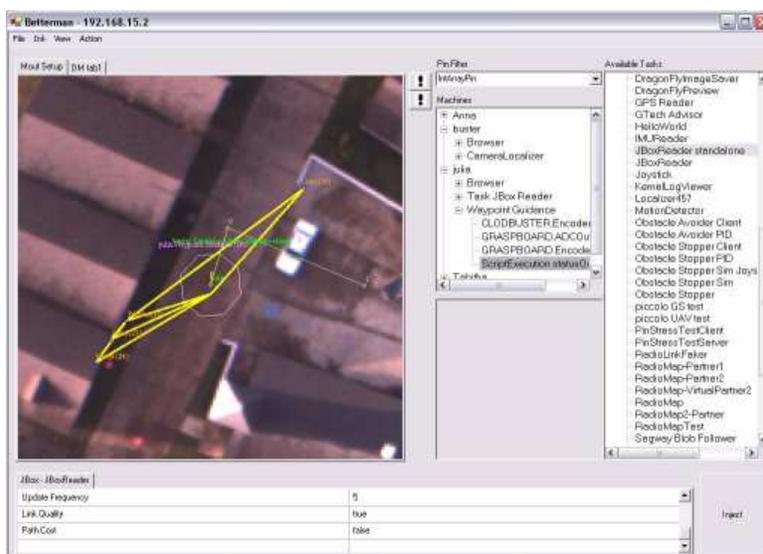


Figure 3. Snapshot of the ROCI Browser during the execution of a mission. Overlaid in the aerial picture, the thick lines represent radio connectivity between nodes and the octagon shows the uncertainty in the localization of one of the robots.

Both users and autonomous agents also have the ability to access sensor data through a distributed database infrastructure. The foundation of this system is made up of instances of a Logger module that is capable of receiving, indexing, and storing any type of pin data. Using these modules, a user can instrument a task with loggers listening to any pin communications. This data is primarily indexed by time (an index that is globally meaningful in the context of distributed sensing), but can also be indexed by more sensor-specific methods (such as position). By using shared indices to join multiple logs in an on-demand fashion, human operators and programmed agents are able to make use of data as it becomes available on the network without relying on any pre-defined schemas. A query architecture is used to interact with this distributed database. Active queries, made up of executable code, are sent to the nodes that are storing the relevant information in order to minimize the amount of data that must be transferred over the network. In other words, the query is moved to the data, rather than the other way around. This system helps to eliminate the distinction between low-level sensor data and high-level fused structures by removing the need to hard code every type of useful structure.

5. Air-Ground Cooperation

The use of air and ground robotic vehicles working in cooperation can be very important in tasks involving the reconnaissance and exploration of cluttered urban environments where communication and GPS localization may be unreliable. In this type of mission, groups of unmanned air vehicles (UAVs) could significantly help the ground vehicles (UGVs) by providing localization data and acting as communication relays.

One of our first experiments in air-ground cooperation was to try to localize our ground vehicles using a sequence of images taken from the blimp. Basically, for each image we had to compute the projection matrix M that relates the position of the robot in a global coordinate frame with its pixel coordinates in the image. We compared two complementary approaches for computing M : the first used a set of known landmarks in the image while the second relied on the measurements from the blimp's on board sensors (GPS/IMU) and the intrinsic parameters of the camera. The localization results were compared with measurements from a GPS on board of the ground vehicle.

As discussed in (Chaimowicz et al., 2004), these experiments demonstrated that none of these approaches could be applied alone if we need a localization system that is applicable, reliable, and accurate. In spite of being very precise, the air-ground localization using known landmarks can not always be applied because it requires the identification of a certain number of world locations to register the image, which is impractical in general situations. On the other hand, the approach based on the blimp's on board sensors did not performed well due to the combination of different sensor errors without an adequate fusion methodology. These preliminary results motivated us to pursue more sophisticated methods for performing the cooperative localization. The general idea is to fuse information from different sources in a systematic way in order to have a more reliable and accurate localization system.

Therefore, we developed a related approach in which air and ground vehicles can collaborate to localize ground features. As noted, when detecting ground features from images taken from a blimp or an airplane, their exact location on the ground is always subject to errors in attitude and location estimates. Thus, for robust localization of ground targets, it is imperative to know and reduce the uncertainty in their position.

This approach builds on previous endeavors in decentralized data fusion (DDF) (Manyika and Durrant-Whyte, 1994). DDF provides a decentralized estimation framework equivalent to the linearized Kalman filter. Decentralized active sensor networks (Grocholsky et al., 2003) extend this to include a control layer that refines the quality of the esti-

mates obtained. The established architecture and methodology is used here.

As detailed in (Grocholsky et al., 2004), our methodology combines UAV and UGV ground feature observations, actively deploying the sensor platforms to maximize the quality of the location estimates obtained. The different perspective provided to sensors on-board air and ground vehicles results in significantly different sensing accuracy and coverage qualities. A collaborative feature search and localization example is illustrated in Figure 4. The approach exploits the complementary character of UAV and UGV sensor platforms. The UAV rapidly covers the designated search area, detecting features and providing relatively uncertain location estimates. UGVs deploy to refine the feature location estimates. Localization accuracy beyond that achievable by UAV sensing alone is realized without requiring the UGVs to conduct an extensive area search.

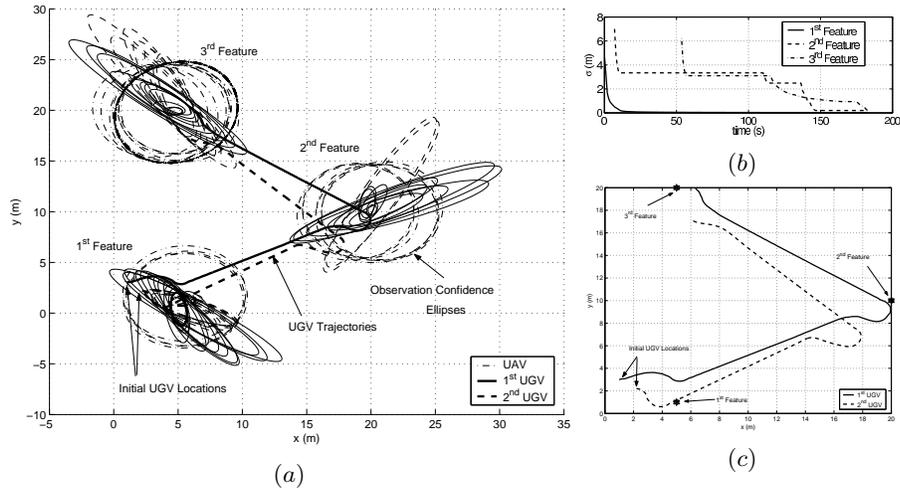


Figure 4. Details of the ground feature localization process. (a) Confidence ellipses associated with UAV and UGV feature observations. Indicating the significant uncertainty reduction with UGV sensing distance. (b) The estimate standard deviation over time. Constant values indicate the time between UAV spotting and UGV refinement. (c) UGV trajectories and true feature locations.

6. Cooperative Radio-Mapping

Communication is essential for coordination in most cooperative control and sensing paradigms. Successful deployment of multi-robot mapping and exploration, surveillance, and search and rescue relies in large

part on a reliable communication network. Often times, these tasks are executed in environments that are adverse to wireless communications. Radio propagation characteristics are difficult to predict a priori since they depend upon a variety of factors such as transmission power, terrain characteristics, 3-D geometry of the environment, and interference from other sources (Neskovic et al., 2000). The difficulty in accurately modeling radio propagation in urban environments makes it difficult to design multi-agent systems such that the individual agents operate within reliable communication range at all times. In this section, we consider the problem of acquiring information for radio connectivity maps in urban terrains that can be used to plan multi-robot tasks and also serve as useful perceptual information.

A radio connectivity map is a function that returns the signal strength between any two positions in the environment. The radio connectivity construction can be formulated as a graph exploration problem for small teams of robots. An overhead surveillance picture is used to automatically generate roadmaps for motion planning and determine a plan to obtain the desired signal strength measurements (Hsieh et al., 2004). The plan consists of a list of waypoints for each robot such that radio signal strength measurements for the connectivity map is obtained as team members simultaneously traverse through each of their respective waypoints.

Radio connectivity maps can be therefore used to plan multi-robot tasks to increase the probability of a reliable communication network during the execution phase. This however will not guarantee that radio signal strength measurements obtained during the exploration and execution phases will not differ due to the sensitivity of radio propagation to environmental factors. Therefore, to ensure reliable communication during task execution, additional recovery behaviors such as returning to the last position the robot was able to successfully communicate with other team members should be included. Ideally, the measurements obtained during the exploration phase can be used to construct a limited model for radio propagation in the given environment such that when coupled with additional reactive behaviors, a reliable communication network can be maintained during deployment.

Preliminary experiments were performed using our ground vehicles to test the radio connectivity at the Ft. Benning MOUT site. In these experiments, each robot is individually tasked with the corresponding list of waypoints determined by the algorithm described in (Hsieh et al., 2004). Team members navigate to their designated waypoints and broadcasts an “arrival” message. Once the robots have completed the radio signal strength measurements, they broadcast a “ready to go” to notify



Figure 5. Preliminary experimental radio connectivity map for the MOUT site obtained using our multi-robot testbed.

each other to move on to their next targeted location. This is repeated until every member has traversed through all the waypoints on their list. The waypoints are selected to minimize the probability of losing connectivity under line-of-sight conditions in the planning phase to ensure the success of the synchronization based on line-of-sight propagation characteristics that can be determined a priori. Figure 6 shows some measurements of the radio signal strength between pairs of positions in the environment. An edge between two pairs of positions shows that the signal strength between the two locations is above the desired threshold. The weights on the edges (barely visible) denote the signal strength that was measured between the two locations. In these experiments, the signal strength was measured using the JBox, described in Section 2.

In the future, we would like to incorporate more reactive behaviors so as to be able to do some on-line mapping instead of collecting data for a set of completely predetermined locations. Furthermore, it is often the case that the exploration of the radio map of the scene is being carried out concurrently with other activities such as environmental monitoring or situational awareness. Thus, another area which we plan to address is pursuing the radio mapping with other objectives and which must be effectively balanced against the other mission goals.

7. Conclusions

This paper presented some of the work that has been done in the GRASP Lab. as part of the DARPA-MARS2020 project for deploying teams of robots in urban environments. We introduced our hardware and software framework, discussed some important issues related to situational awareness, air-ground cooperation and cooperative radio mapping, and presented some preliminary results obtained during field tests. This project is scheduled to culminate with a demonstration of the performance of the integrated team of UAVs, UGVs, and a human leader in a reconnaissance type mission at the Fort Benning McKenna Range MOUT site in December 2004.

Acknowledgments

This work was in part supported by The Defense Advanced Research Projects Agency (DARPA MARS NBCH1020012).

References

- Chaimowicz, L., Cowley, A., Sabella, V., and Taylor, C. J. (2003). ROCI: A distributed framework for multi-robot perception and control. In *Proceedings of the 2003 IEEE/RJS International Conference on Intelligent Robots and Systems*, pages 266–271.
- Chaimowicz, L., Grocholsky, B., Keller, J. F., Kumar, V., and Taylor, C. J. (2004). Experiments in multirobot air-ground coordination. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, pages 4053–4058.
- Cowley, A., Hsu, H., and Taylor, C. (2004). Distributed sensor databases for multi-robot teams. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*.
- Grocholsky, B., Bayraktar, S., Kumar, V., and Pappas, G. (2004). Uav and ugv collaboration for active ground feature search and localization. In *Proc. of the AIAA 3rd "Unmanned Unlimited" Technical Conference*.
- Grocholsky, B., Makarenko, A., Kaupp, T., and Durrant-Whyte, H. (2003). Scalable control of decentralised sensor platforms. In *Information Processing in Sensor Networks: 2nd Int Workshop, IPSN03*, pages 96–112.
- Hsieh, M. A., Kumar, V., and Taylor, C. J. (2004). Constructing radio signal strength maps with multiple robots. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, pages 4184–4189.
- Manyika, J. and Durrant-Whyte, H. (1994). *Data Fusion and Sensor Management: An Information-Theoretic Approach*. Prentice Hall.
- Neskovic, A., Neskovic, N., and Paunovic, G. (2000). Modern approaches in modeling of mobile radio systems propagation environment. *IEEE Communications Surveys*.