

EXPERIMENTS WITH A TEAM OF MINIATURE ROBOTS

NIKOLAOS PAPANIKOLOPOULOS*, SASCHA A.
STOETER, PAUL RYBSKI, MARIA GINI, DEAN F.
HOUGEN, MICHAEL ERICKSON

Center for Distributed Robotics, Department of Computer Science
and Engineering, University of Minnesota, Minneapolis, MN 55455

Abstract. This paper presents several experiments with a large team of heterogeneous robots. The team consists of two types of robotic agents. The first type is a larger, heavy-duty robotic platform, called the “ranger.” Rangers are used to transport, deploy, and supervise a number of small, mobile sensor platforms called “scouts,” the second type of robotic agent. In an example scenario, the scouts are deployed into an office/lab environment, navigate towards dark areas, and position themselves to detect moving objects using their cameras. A ranger communicates with each of the scouts and determines whether there are objects of potential interest within the observed area. The paper also includes experimental results for individual scout and ranger-scout activities.

Key Words. Distributed robotics, miniature robots, robotic teams.

1. INTRODUCTION

For security and surveillance applications, an area is typically observed either by (1) multiple remote sensing devices that report to a coordination agent or (2) a mobile agent that patrols the required area. In both cases, the problem of adequate sensor coverage exists. In case 1, the problem is spatial: Are there enough sensors in the right locations? In case 2, the problem is temporal: Will the mobile agent be in the right place at the right time?

One possible solution is to combine the two approaches into one. A mobile agent that is capable of long distance travel can cover a large area and deploy smaller, less mobile agents in vari-

ous locations. The smaller agents can be given the responsibility of sensing a small area and can have the flexibility to change their vantage points to make sure that all of their local area is observed. A coordination agent can then communicate with the sensing agents, query them for information, and move them remotely to increase the area viewed by them.

This is the solution that this paper suggests. The robots that are used are customized RWI ATRV-JrTM-based robots called “rangers” and a group of extremely small custom mobile sensor platforms called “scouts.” Rangers are capable of navigating long distances without needing to recharge their batteries and are capable of navigating off-road terrain. Due to their size, however, not all areas may be accessible to them. The small size of the scouts makes them much easier

* Corresponding Author, email: npapas@cs.umn.edu
phone: 612-625-0163, fax: 612-625-0572.

to operate in these areas but presents a different set of problems including decreased range, battery lifetime, computational power, and sensing ability. By putting both kinds of robots into a team, the benefits of both can be achieved.

In this team, rangers are used as the primary navigational and computational resources. Their responsibilities include traversing the environment, selecting appropriate locations that are to be observed, and deploying the scouts into those areas. Once the scouts reach their designated locations, a ranger contacts each of the scouts in turn to analyze the area.

The paper is organized as follows: Section 2 describes the hardware components of the team, and Section 3 describes the software aspects. Finally, experimental results are presented in Section 4.

2. HARDWARE

The two robotic platforms used in this work are the rangers and the scouts. The next section describes the focus of our system, the scouts.

2.1. Scouts

Scouts are custom cylindrical robots 40 mm in diameter and 110 mm in length (see Figure 1) possessing a unique combination of locomotion modes. A scout can roll using its wheels (one on each end of its body) and a leaf spring “foot” mounted underneath for stabilization. It is also able to hop by winching its spring foot around its body and releasing it in a sudden motion.

For the scenario presented in this paper, each scout possesses a miniature video camera and a wireless video transmitter. The camera consists of a monochrome single chip CMOS video sensor and a miniature pinhole lens. Video data is broadcast back to a receiver via a 900 MHz analog video transmitter. Each scout also possesses a miniature RF data transceiver for receiving commands from rangers and transmitting status information back. Scouts are discussed more fully in Hougen et al. [1, 2].




Figure 1. Two scout robots (shown next to a quarter for scale).

3. SOFTWARE

In order for the rangers and the scouts to coordinate their efforts and work together properly, a proxy-processing system has been developed which allows for the scout’s control programs to run on a computer separate from the scout hardware. This is important because the small size of the scout’s design severely restricts the speed and computation power of its on-board computer. The scout’s limited on-board computational resources can only handle the most basic low-level control routines, such as setting the pulse width modulation frequency of the motor controllers, handling the sensor payload, and decoding information received on the inter-robot RF data link. High-level scout control is achieved by executing the scout motion control algorithms on a different computer. In the experiments described here, all of the scout programs run as separate processes on the ranger’s on-board computers. These scout control programs send commands to the scout hardware through an RF data link.

While this proxy-processing scheme means that each physical scout *robot* is dependent on a ranger robot, this does not mean that the corresponding scout *agent* is not autonomous. Rather, the scout agent can be seen as distributed, with its sensors and effectors located in the scout robot and its behavioral controller located on board a ranger’s computer. This takes advantage of the nature of artificial intelligent agents—they are

not limited to a single physical location.

We have developed behaviors for a scenario in which rangers will find interesting areas to explore and deploy scouts into them. In our scenario a ranger is placed in a building to traverse the corridor and launch scouts into rooms that it finds along its path. A second ranger is used as a communication agent to coordinate the actions of the deployed scouts. The scouts must find dark areas in which to conceal themselves and watch for moving entities (such as people).

3.1. Scout Behaviors

Several simple behaviors have been implemented for the scouts. The only environmental sensor available to the scout is its video camera, the use of which presents several problems. First, the scout's proximity to the floor severely restricts the area it can view at a time. Secondly, since the video is broadcast over an RF link to the ranger for processing, the quality of the received video often degrades due to of range limitations, proximity of objects that interfere with transmission, and poor orientation of the antennas.

The scout behaviors are:

Locate Goal: Determining the location of the darkest (or lightest) area of the room is accomplished by spinning the scout in a circle and checking the mean value of the pixels in the image. Since the scout has no encoders on its wheels to determine how far (or even if) it has moved, frame differencing is used to determine whether motion took place. The circular scan is accomplished in a number of discrete movements. Before each move, the scout must determine the quality of the video and set a threshold to filter out RF noise. This is accomplished by doing image differencing on a stream of video and increasing a difference threshold until RF noise is filtered out. Once the threshold is set, the robot takes an image, rotates for half a second, takes a new image, and subtracts the new image from the old one. A large difference indicates movement. There are several instances where this approach can fail, however. First, if the transmitted image quality is so low that motion in the image can-

not be distinguished from noise. Second, if the robot is operating in an area of very low light or very uniform color, there may not be enough detail in the images to generate significant differences.

Drive Towards Goal: Identifying a dark area to move towards is a simple matter of scanning across the image at a fixed level on or about the level of the horizon and determining the horizontal position of the darkest area in the image. The mean pixel values in a set of overlapping windows in the image are determined. The scout selects the darkest window and drives in that direction. The scout knows that it should stop when its camera is either pressed up against a dark object, or the scout is in shadows. Scout motion in this behavior is continuous and the scout does not check its movements by frame differencing (unlike the discrete movements of the previous behavior). This is because the scout is unable to move very quickly. The difference between subsequent frames captured during forward motion is minimal, making it very difficult for the robot to detect forward motion.

Detect Motion: Detecting moving objects is accomplished using frame differencing. Once the scout has been placed in a single location, it sets its frame differencing noise threshold in the same way as described in the Locate Goal behavior. The scout then subtracts sequential images in the video stream and determines whether the scene changes at all (caused by movement in the image.)

Handle Collisions: If the scout drives into an obstacle, all motion in the image frame will stop. If no motion is detected after the scout attempts to move, it will invoke this behavior and start moving in random directions in an attempt to free itself. In addition to freeing the scout from an object that it has driven into, this random motion has an additional benefit. If the scout is in a position where the video reception quality is extremely bad, the static in the image will prevent the scout from detecting any motion

(regardless of whether it is hung up on an object). Moving the scout changes the orientation of the antenna which may help improve reception.

4. EXPERIMENTS



Figure 2. Top view of experiment 1.

Three different experiments were devised to test the scout and ranger's ability to function in an environment, report back useful data and operate successfully as a team. The first two experiments tested the ability of the scout to locate useful goals in various environments and move towards them. The third experiment tested the ability of the rangers and scouts to work together to achieve a useful goal.

4.1. Experiment 1

The first experiment was to determine, in a controlled environment, how well the scout could locate and move towards an appropriately dark area. This experiment was designed to examine the scout's behaviors in an analytical fashion.

For the first experiment, a controlled environment was constructed. This environment consisted of a 2.5 m \times 3 m enclosed rectangle with uniformly-colored walls. A 1 m \times 0.5 m black

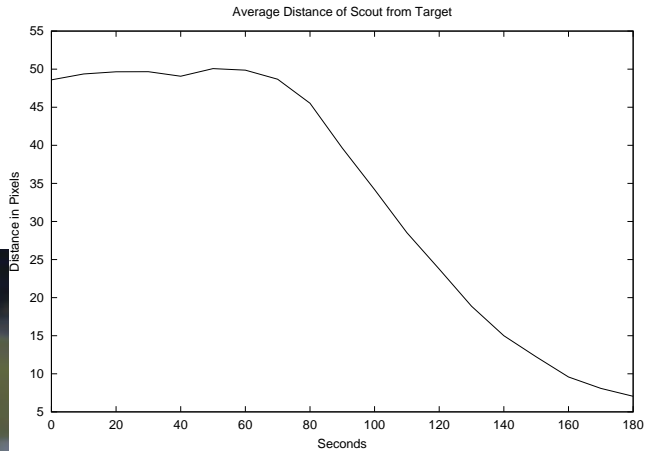


Figure 3. Experiment 1: Average distance (nine runs) of the robot from the target. Distance is in pixels, determined in Figure 2. 1 pixel is approximately 3cm.

rectangle set up on one side of the environment as the target for the scout. The scout was started 1.5 m away from the center of the target.

Nine experiments were run to see how long it would take the scout to locate the black target object and move itself towards it. A camera was mounted on the ceiling of the room and was used to view the progress of the robot from above. A simple tracking algorithm was used to automatically chart the progress of the scout as it moved towards the target. Figure 2 shows the view from this camera as well as a superimposed plot of the path that the scout took to reach its objective during one of its nine runs. In each case, the scout successfully located the target and moved towards it.

Figure 3 shows a plot of average distance the scout was away from the target vs. time for all of these runs. In the first 70-80 seconds, the scout uses its Locate Goal behavior to find the dark spot. Once it identifies it, the scout starts its Drive Towards Goal behavior until it comes in contact with the goal, somewhere between 150 and 160 seconds.




Figure 4. Experiment 2: Lab environment, showing locations of the scout for all five runs. X marks the starting position used in all runs and numbered arrows correspond to final position and orientation for individual runs. Ovals represent chairs under which scouts may hide. Chairs are positioned at a table and a lab bench, both of which also provide hiding opportunities. Other objects are impassable.

4.2. Experiment 2

The second experiment was set up to determine how well the scout could position itself in a more “real world” environment— a somewhat cluttered office or lab space. For these experiments, the scout’s ability to locate a dark area was combined with the ability to turn towards the lighter areas and search for moving objects.

Two environments were used for this experiment. One was a lab environment with chairs, a table, lab benches, cabinets, boxes, and miscellaneous other materials (see Figure 4). The other was an office environment with chairs, a table, desks, cabinets, and boxes (see Figure 5). The floor of the lab is a shiny, light tile surface of relatively uniform color whereas the floor of the office




Figure 5. Experiment 2: Office environment, showing locations of the scout for all five runs. X marks the starting position used in all runs and numbered arrows correspond to final position and orientation for individual runs. Ovals represent chairs under which scouts may hide. Chairs are positioned at a table and at two desks, all of which also provide hiding opportunities. Other objects are impassable.

is a carpet of medium and dark piles providing a high localized contrast. This difference in surface brightness and contrast were accounted for in the scout’s vision behaviors which were effectively self-calibrating. Five runs were conducted in each environment, using a fixed starting point for the scout in each room (shown as an X in Figures 4 and 5).

In four of the five runs in the lab environment, the scout chose the same chair under which to hide (locations 1 & 3-5 in Figure 4). On run number 2, however, the scout wound up roughly 0.5m out from under the chair in a relatively exposed position (location 2 in Figure 4). In all five runs the scout ended up facing towards a relatively bright area of the room. However, in run 4 this happened to be towards the rear of the room.

Similarly, in four of the five runs in the office environment, the scout chose the same chair as its destination (locations 1-3 & 5 in Figure 5).

On run 4 the scout chose the other nearby chair (location 4 in Figure 5). In four of the five runs the scout wound up facing brightly lit areas roughly towards the door of the office. On run 1, though, the scout became physically stuck under the chair, forcing it to face the somewhat darker area towards the back of the room.

While collecting data, several experiments had to be aborted and restarted due to problems with radio communication. In these experiments, the scout found itself surrounded by objects that interfered with RF transmissions. In other experiments, the scout's batteries ran low and had to be replaced before data collection could continue.

4.3. Experiment 3

The third experiment was designed to determine if the combined scout/ranger team could carry out an entire surveillance mission. This mission combines all behaviors described above. The scouts are initially manually loaded into the launcher, mounted on Ranger 1. Ranger 2 acts as a communication relay. From there on the actions of the team are autonomous. Ranger 1 moves down the hall, finds doors, and launches the scouts through doorways. Each scout, through proxy processing with Ranger 2, finds the darkest area visible from its landing site, drives to the dark area, turns around to face the more brightly-lit room, and begins watching for motion. In all of the experiments, the scouts were able to detect the motion of a person walking through the areas, either the lab or the office space.

5. SUMMARY

The system as presented in this paper handles a task where cooperation increases performance by increasing reliability. By having its sensors spread throughout the environment with several agents, rather than concentrated on a single agent, there is less chance of an observation being missed. Further, because some of the agents are small and more easily hidden, even persons attempting to avoid detection by the system are more likely to be detected than in the case of

a single, large robotic security guard. The controlling agent architecture is distributed in nature, allowing the controlling algorithms for the smaller "computationally challenged" robots to run on a computer separate from the physical body of the robot. This provides for greater flexibility and power for interpreting the environmental information provided to the robot.

6. ACKNOWLEDGEMENTS

This material is based upon work supported by the Defense Advanced Research Projects Agency, Microsystems Technology Office (Distributed Robotics), ARPA Order No. G155, Program Code No. 8H20, issued by DARPA/CMD under Contract #MDA972-98-C-0008.

References

- [1] D. F. Hougen, S. Benjaafar, J. C. Bonney, J. R. Budenske, M. Dvorak, M. Gini, D. G. Krantz, P. Y. Li, F. Malver, B. Nelson, N. Papanikolopoulos, P. E. Rybski, S. A. Stoeter, R. Voyles, and K. B. Yesin. A miniature robotic system for reconnaissance and surveillance. In *Proc. of the IEEE Int'l Conference on Robotics and Automation*, pages 501–507, San Francisco, CA, Apr. 2000.
- [2] D. F. Hougen, J. C. Bonney, J. R. Budenske, M. Dvorak, M. Gini, D. G. Krantz, F. Malver, B. Nelson, N. Papanikolopoulos, P. E. Rybski, S. A. Stoeter, R. Voyles, and K. B. Yesin. Reconfigurable robots for distributed robotics. In *Government Microcircuit Applications Conference*, Anaheim, CA, Mar. 2000.