Enlisting Rangers and Scouts for Reconnaissance and Surveillance

A Hierarchical Team of Robots Capable of Carrying Out Complex Missions in a Wide Variety of Environments, Both Military and Civilian

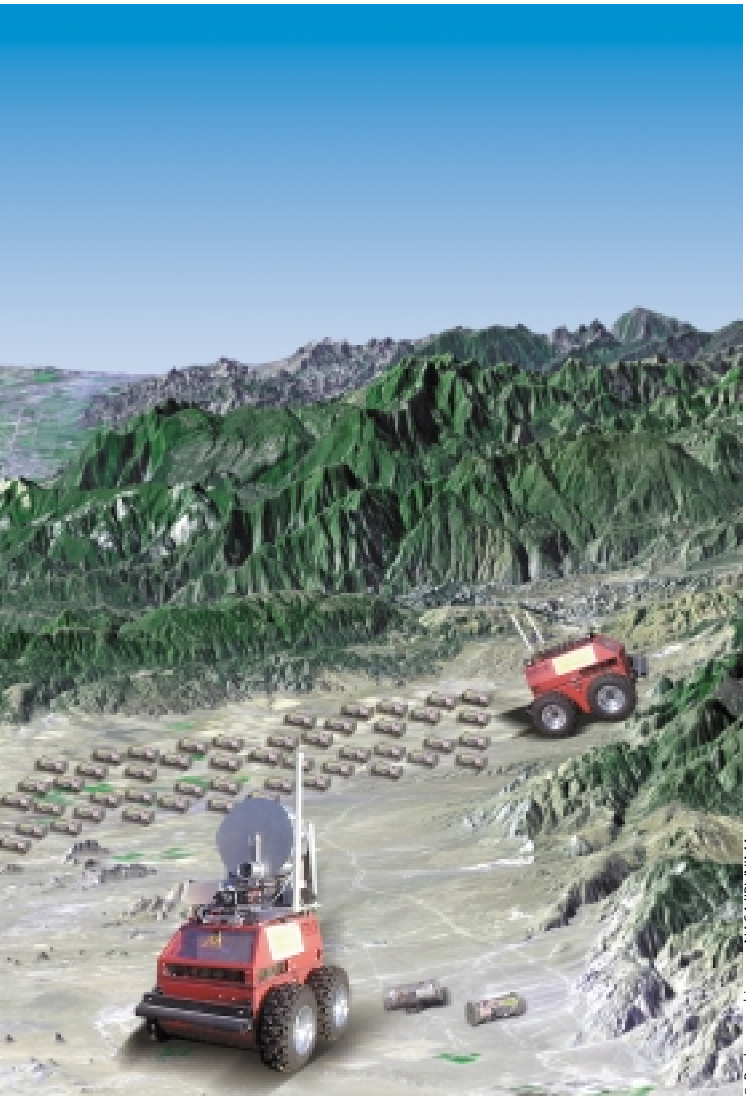

© Background image: NASA/JPL/NIMA

by PAUL E. RYBSKI,
NIKOLAOS P. PAPANIKOLOPOULOS,
SASCHA A. STOETER,
DONALD G. KRANTZ, KEMAL B. YESIN,
MARIA GINI, RICHARD VOYLES,
DEAN F. HOUGEN, BRAD NELSON,
and MICHAEL D. ERICKSON

Reconnaissance and surveillance are important activities for both military and civilian organizations. Hostage and survivor rescue missions, illicit drug raids, and responses to chemical or toxic waste spills are just some of the operations requiring a reconnaissance or surveillance component. To address these needs, we have developed a distributed heterogeneous robotic team that is based mainly on a miniature robotic system.

Because some of these operations require covert action, we have chosen to make most of the robots on the team extremely small so that they will evade detection. This small size also allows them to be easily transported and allows for a greater number (dozens) to be brought into use for a single operation. The small size and large number also make individual robots expendable without jeopardizing the overall mission. We call these small robots *scouts,* and they act as the roving eyes, ears, noses, etc. of our system.

The small size of the scouts creates great challenges, however. The individual components must all be exceedingly small and the overall design of the scout must make maximum use of all available space. Further, the scouts must make efficient use of resources (e.g., batteries) in order to survive for a

1070-9932/00/$10.00©2000IEEE

useful period of time. We meet these challenges with an innovative scout design and creative use of additional support. We team the scouts with larger *ranger* robots. The rangers can transport the scouts over distances of several kilometers and deploy the scouts rapidly over a large area. Rangers also serve to coordinate the behaviors of multiple scouts as well as to collect and present the data in an organized manner to the people who will ultimately make use of it.

In this article, we present the scouts and rangers, discuss the capabilities of each (in the case of the ranger the emphasis is on its role as a utility platform for the scouts) along with the associated software, and describe demonstrations conducted to test the innovative aspects of the system. We also discuss related work, analyze our results, and draw conclusions about our system.

Related Work

Automatic security and surveillance systems using cameras and other sensors are becoming more common. These typically use sensors in fixed locations, either connected ad hoc or, increasingly, through the shared communication lines of "intelligent buildings" [14]. These may be portable to allow for rapid deployment [15] but still require human intervention to reposition when necessary. This shortcoming is exacerbated in cases in which the surveillance team does not have full control of the area to be investigated, as happens in many law-enforcement scenarios. Static sensors have another disadvantage. They do not provide adaptability to changes in the environment or in the task. In case of poor data quality, for instance, we might want the agent to move closer to its target in order to sense it better.

Mobile robotics can overcome these problems by giving the sensor wheels and autonomy. This research has traditionally focused on single, large, independent robots designed to replace a single human security guard as he makes his rounds [7]. Such systems are now available commercially and are in place in, for example, factory, warehouse, and hospital settings 8, 10, 13], and research continues along these lines [3, 11, 16]. However, the single mobile agent is unable to be in many places at once—one of the reasons why security systems were initially developed. Unfortunately, large mobile robots are unable to conceal themselves, which they may need to do in hostile or covert operations. They may also be too large to explore tight areas.

Multiple robots often can do tasks that a single robot would not be able to do or do them faster, as described in the extensive survey by Cao et al. [2]. The tasks traditionally studied with multiple robots are foraging [9], which involves searching and retrieving items from a given area; formation marching [1], which involves moving while maintaining a fixed pattern; map making [6]; and janitorial service [12], where robots have to clean a room in an unfamiliar building by emptying the garbage, dusting the furniture, and cleaning the floor.

A nice overview of miniature single robots for space and urban warfare is given by C. Weisbin et al. in [17], where vehicles such as the Sojourner, the Sample Return Rover, and the Nanorover are presented and discussed. Finally, multiple mobile robots for security have recently been investigated [5]. In this case, the robots were meant to augment human security guards and fixed sensor systems in a known and semi-tailored environment.

Scouts

Scouts must be small yet highly capable robots. They must be easily deployable and able to move efficiently yet traverse obstacles or uneven terrain. They must be able to sense their environment, act on their sensing, and report their findings. They must be able to be controlled in a coordinated manner.

To support all of these requirements, we have designed a robot 40 mm in diameter and 110 mm in length (see Fig. 1). Its cylindrical shape allows it to be deployed by launching it from an appropriate barreled device (see the "Launcher" section below). Once deployed, it moves using a unique combination of locomotion types. It can roll using wheels (one on each end of the cylinder body) and jump using a spring "foot" mechanism. The rolling allows for efficient traversal of relatively smooth surfaces, while the jumping allows it to operate in uneven terrain and pass over obstacles.

Besides the mechanical components, scouts contain a vast array of electronic devices. Each scout is provided with a sen-

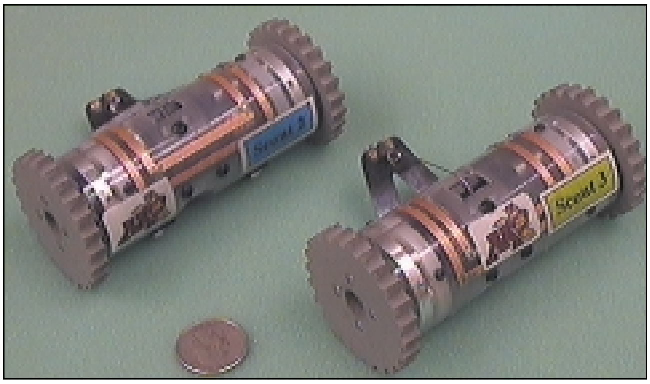

Figure 1. *Two scout robots (shown next to a quarter for scale).*

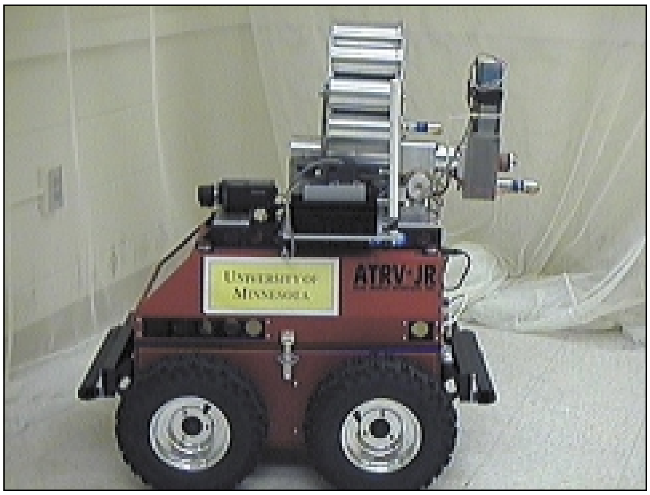

Figure 2. *A ranger with scout launcher.*

sor suite, which may vary with the scout's mission. All scouts contain magnetometers and tiltometers. Scouts may also contain some combination of a CMOS camera, a passive infrared sensor, a microphone, a MEMS vibration sensor, a MEMS gas sensor, and other sensors. The camera may be mounted on a pan-tilt mechanism or in a fixed position within the scout body. The scouts contain transmitters/receivers for transmitting video and audio signals and other sensed data and for receiving commands. The scouts contain microcontrollers for radio/network management and implementation of autonomous behaviors.

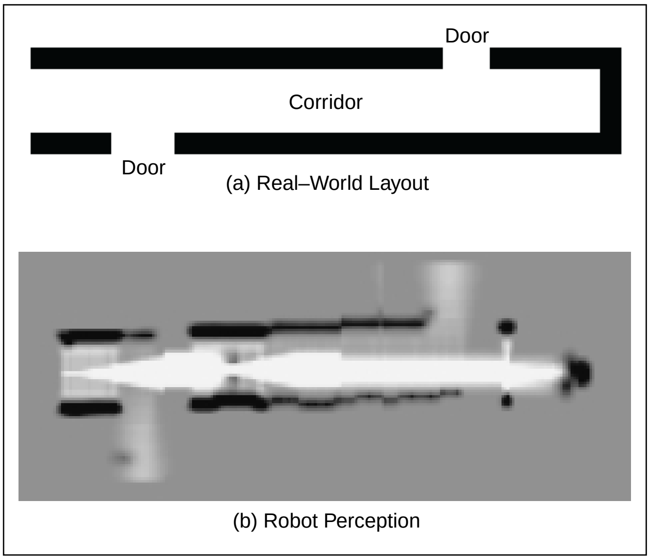

Figure 3. *The environment.*

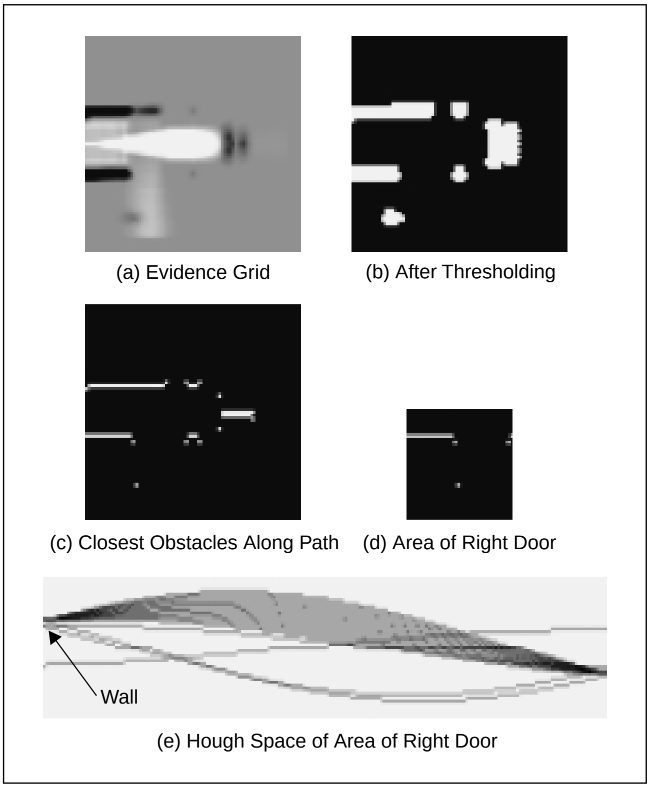

Figure 4. *Processing of an evidence grid.*

Scout Video

A video reconnaissance module that consists of a miniature video camera, an optional pan-tilt mechanism, and a wireless video transmitter was built to provide visual feedback from the scouts. The camera consists of a single-chip CMOS video sensor and a miniature pinhole lens. Unlike the common CCD-type video sensors, the CMOS sensor is able to integrate all functionality within a single chip, reducing its size dramatically. Additionally, CMOS sensors typically consume three to five times less power than a comparable CCD sensor. The overall dimensions of the camera are 15 mm × 15 mm × 16 mm, and the power consumption is 100 mW.

A pan-tilt mechanism provides increased field of view to the camera. Micromotors, which became recently available, are utilized for actuation. These brushless DC motors are only 3 mm in diameter and 15 mm in length.

Scout Communication

For other communication with the scouts, a miniature transceiver has been developed that employs on/off keying (OOK) modulation and operates at 434 MHz. The communications make use of a media-access control (MAC) protocol implemented on a Scenix SX processor using RF Monolithics Virtual Wire components. The MAC's reliable delivery scheme is augmented with a version of Architecture Technology Corporation's Source Adaptive Routing Algorithm (SARA) that has been simplified to operate within the confines of the SX processor's 2 KB program ROM and 128 B RAM. The simplified SARA implementation allows RF nodes (rangers or scouts) to act as routers in order to increase end-to-end communications range. As the nodes move, routing information is dynamically updated throughout the wireless network.

Rangers

The scouts function in conjunction with the rangers, which act as utility platforms. Rangers move the team rapidly into place and deploy the scouts. They process the sensor data as needed for scout behaviors and group behaviors and act as coordinators for the team. Finally, they organize the data streams for presentation to users.

Our rangers are based on the ATRV-Jr.™ platform from the RWI Division of iRobot. The "Junior" was developed by RWI with input from our team and others wanting a smaller platform than the existing ATRV-2™, suitable for both outdoor and indoor operation.

Rangers can carry the scouts into position over distances of up to 20 km, giving the scouts a much greater effective range than they would have if they needed to transport themselves. Further, by mounting a novel launching mechanism on each ranger (see Fig. 2), scouts may be deployed more rapidly and into places rangers might have difficulty accessing. Rangers also are equipped with significant computing resources that allow for proxy processing for scout behaviors and for mission coordination.

Launcher

The launcher system is used to deploy the scouts around the field of operation. The launcher can selectively launch any of the ten scouts in its magazine, at a selectable elevation angle and propulsion force, up to a range of 30 m. Scouts are launched with a compressed spring. A stepper motor is used to compress the spring via a rack-and-pinion setup. The pinion is engaged to the motor when the spring is to be cocked and is disengaged when the scout is launched. The mechanical energy efficiency is about 45% due to the weight of the piston and the rack-and-pinion mechanism. The indexing of the magazine to select a particular scout is achieved accurately with an internal Geneva mechanism, without the need of an encoder.

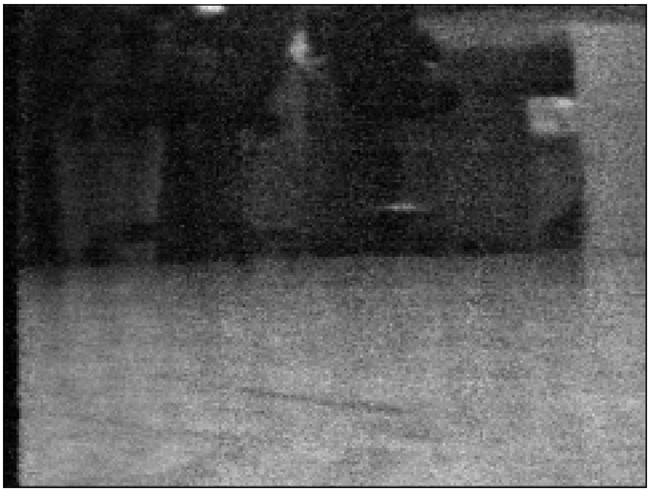

Figure 5. *The world from the scout's point of view. Here the scout is viewing a lab bench and two chairs at a range of 2 m.*

Ranger Computer Resources

Each ranger is equipped with a Pentium 233 MHz-based PC running Red-Hat Linux, which is linked to the robot's sensors and actuators with RWI's rFLEX™ control interface. The PC runs RWI's Mobility™ (an object-oriented, CORBA-based modular robotics control architecture). Ranger-to-ranger data communications are implemented using a 2.4 GHz frequency-hopping wireless LAN system.

Software Components

In order for the rangers and the scouts to coordinate their efforts and work together properly, a proxy processing system has been developed that allows the ranger to control the scouts. The scout's limited computational resources restrict it to handling only the most basic low-level control routines (such as pulsewidth modulation control of the motors). High-level control is achieved through this proxy-processing scheme in which the individual control algorithms that direct the motion of the scout are run as separate threads on-board the ranger's computer. This control is accomplished through the use of a client/server style of architecture where the ranger sends command packets through an RF data link to the scouts.

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).

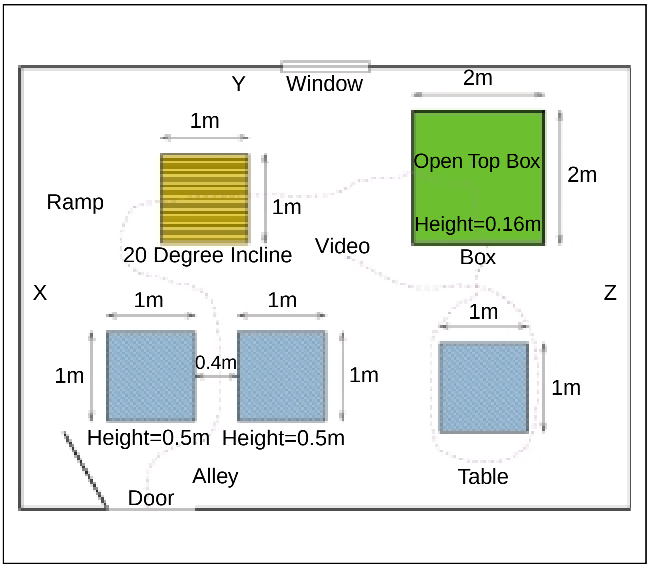

Figure 6. *The obstacle course.*

Ranger Behaviors

Door detection and motion control are solely based on sonar input. Concurrent to the ranger's motion, sonar readings from two side-sonars and one front-sonar are integrated into an evidence grid [4]. Evidence grids partition the world into a number of discrete cells. Each cell carries a probability value describing the likelihood of that part of the world being free space. A sensor model expresses the effect of a reading from that sensor on the evidence grid. This allows for readings from different sensor sources to be combined into a unified model of the world. Here, the evidence grid covers an area of 4 m × 4 m centered around the robot where each cell is 6.25 cm on a side. The environment in Fig. 3(a) is perceived by the ranger as depicted in Fig. 3(b). White areas are considered free whereas black spots are likely to contain obstacles. Gray regions indicate insufficient knowledge to assume either state.

To identify doors or any other opening in a wall, the evidence grid surrounding the ranger [Fig. 4(a)] is treated as a grayscale image. Note that the ranger moves to the right of the image. First, a threshold is applied to retain occluded regions resulting in Fig. 4(b). Figure 4(c) shows the cells containing obstacles closest to the axis of motion. The remaining pixels to the left and right of the ranger are split into two subimages, shown in Fig. 4(d), and then projected into Hough space to find lines in the image that correspond to the walls. Figure 4(e) shows the Hough space of the right side of the ranger. The darkest pixel in this image corresponds to the location of the wall with respect to the ranger. Lastly, openings are searched for along these lines within a dynamically chosen strip. If the opening is wide enough (i.e., about 1 m), then it is classified as a door.

The ranger moves back to center itself in the door frame, turns to face the door, and launches a scout. After successful deployment, it continues to seek out further rooms until all scouts have been exhausted from the magazine.

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 one 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 range limitations, proximity of objects that interfere with transmission, and poor orientation of the antennas. Figure 5 shows an example image received from the scout's camera.

The 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. It does so by doing image differencing and increasing the difference threshold until noise is no longer detected. Once the threshold is set, the robot rotates for half a second and compares the current image against the previous image. A large difference indicates movement. There are several instances where this approach can fail, however. First, if the transmitted image quality is too low, motion in the image cannot be distinguished from noise. Second, if the robot

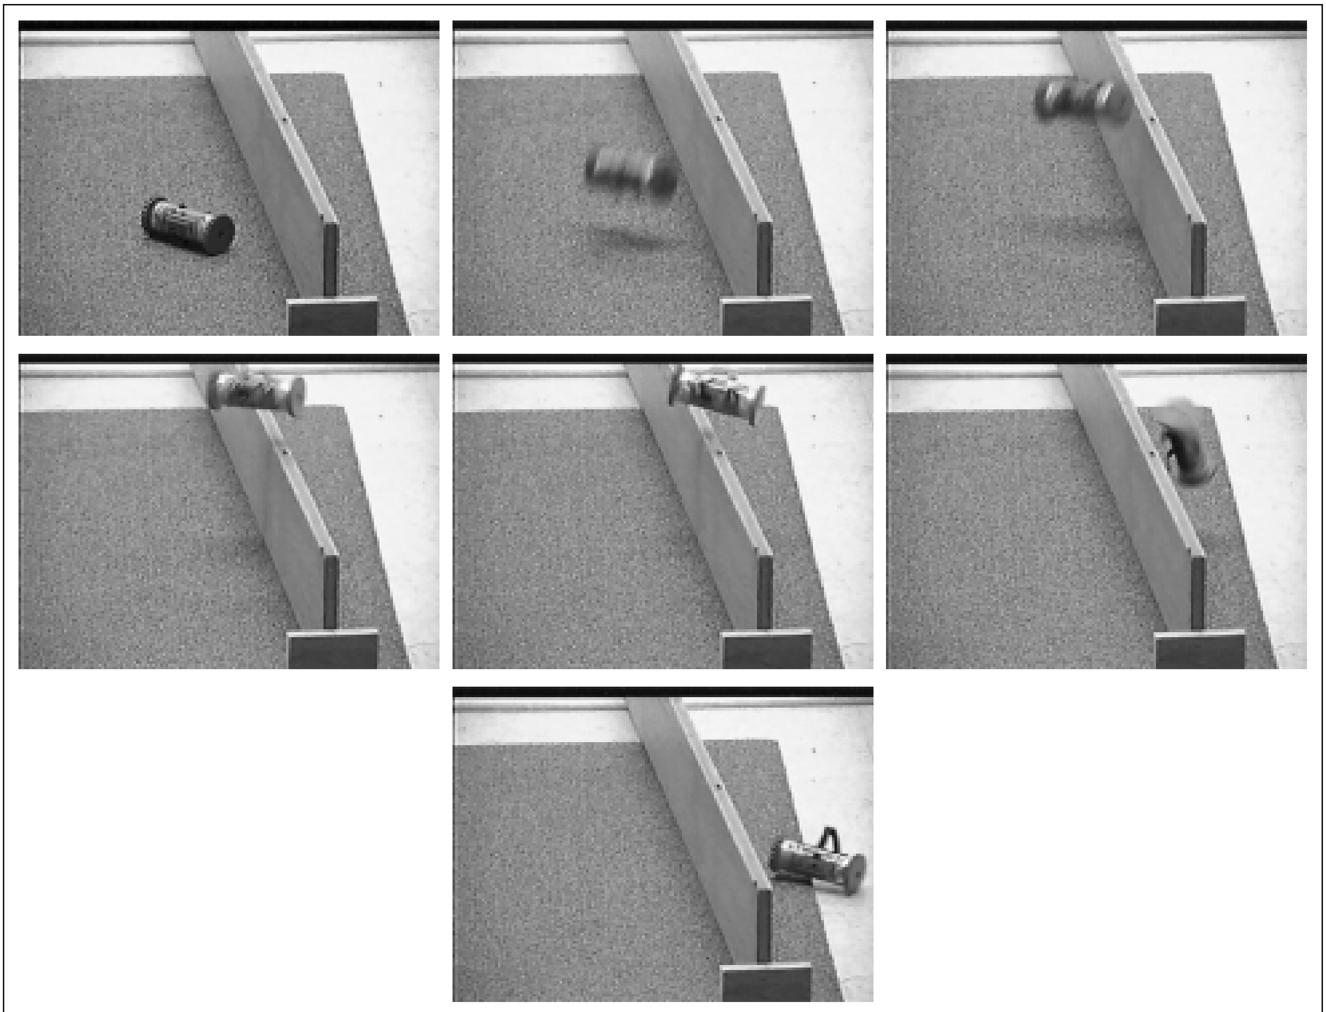

Figure 7. A scout jumping over a barrier (sequence starts from the upper left corner).

Figure 8. Launching a scout through a window.

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:** Moving towards a dark area 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 robot selects the darkest window and drives in that direction. The robot knows that it should stop when it is either pressed up against a dark object, in which case the entire image is uniformly colored, or it is in shadows and the floor is reflecting roughly the same illumination as what is coming down from above the robot. 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. From there, the ranger can activate the scout's camera and determine if there is motion in the field of view of the scout.
- **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 unsticking 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.

Table 1. Scout Performance on Obstacle Course

Trial	Alley	Ramp	Box	Table	Video	Time
1	ok	ok	ok	ok	ok	6:30
2	ok	ok	2^3	ok	ok	6:59
3^1	ok	ok	ok	ok	ok	7:02
4	ok	ok	ok	ok	ok	7:30
5^2	ok	ok	1^3	ok	ok	8:45
6^2	ok	1^4	ok	ok	ok	5:50
7	ok	ok	ok	ok	ok	4:28
8	ok	ok	ok	ok	ok	4:52
9	ok	ok	ok	ok	ok	7:28
10	ok	ok	ok	ok	ok	7:18
11^2	ok	ok	ok	ok	ok	7:12
12^2	ok	ok	ok	ok	ok	7:52
13^2	ok	ok	1^3	ok	ok	10.05
14^1	1^3	ok	1^3	2^3	ok	9:15
15	1^3	ok	ok	ok	ok	5:55

Time is in minutes. Numbers in middle five columns indicate frequency of specific errors, "ok" indicates perfect performance. Possible errors: for alley and table, collision; for ramp, falling off side; for box, missed jump. The superscripts indicate reasons for suboptimal performance. [1]Faulty reassembly of scout after battery change. [2]Manual reset of scout required. [3]Operator error. [4]Loss of communication link.

Figure 9. Top view of experiment 1.

Experimental Results

To test the innovative aspects of our system, we conducted three basic sets of tests. The first set was aimed mainly at testing the hardware capabilities of our scout robots alone, the second was aimed at testing the survivability of our scouts when deployed and controlled by the rangers, and the last set was aimed at testing the operational capabilities of the scouts.

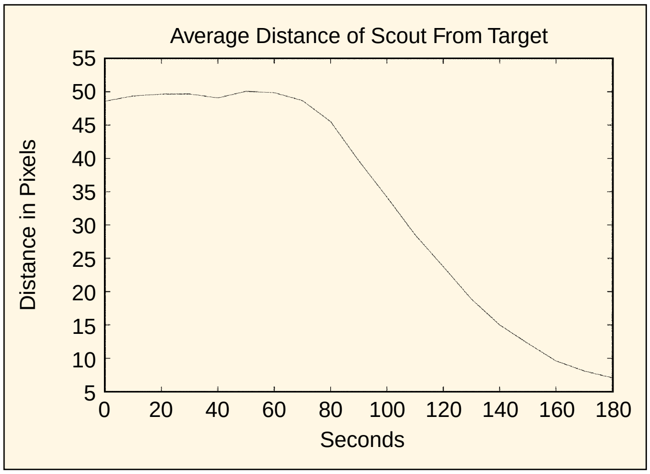

Figure 10. *Experiment 1: Average distance (nine runs) of the robot from the target. Distance is in pixels, determined in Fig. 9. 1 pixel is approximately 3 cm.*

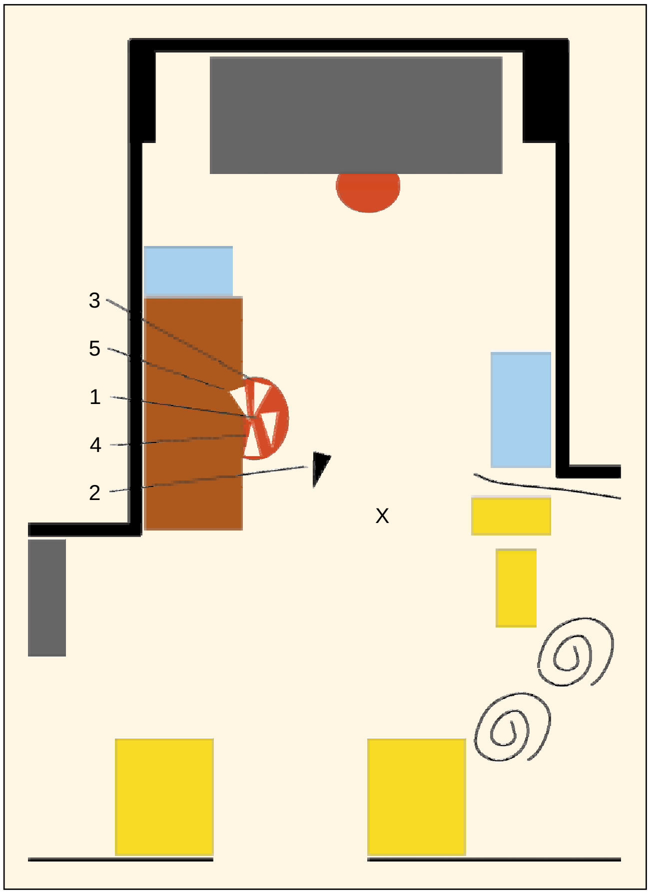

Figure 11. *Experiment 2: Lab environment, showing locations of scouts for all five runs. X marks the starting position used in all runs and numbered arrows correspond to final position and orientation for individual runs. Ovals represent chairs under which scouts may hide. Chairs are positioned at a table and a lab bench, both of which also provide hiding opportunities. Other objects are impassable.*

Scout Hardware Capabilities

To test the capabilities of individual scout robots, we constructed an obstacle course (see Fig. 6). The five major components of the obstacle course are:

- **Alley.** The scout must follow a straight path between two large obstacles without hitting either. The obstacles form an alley 1 m in length and 0.4 m wide.
- **Ramp.** The scout must roll up a ramp at a 20° incline then, from the top of the ramp, jump or fall back to the floor. The scout must not fall from the sides of the ramp nor roll back down the incline.
- **Box.** The scout must jump into and out of an open box (see Fig. 7). The height of the box's walls is 0.16 m.
- **Table.** The scout must circle around a 1 m square table without touching it.
- **Video.** The scout must drive to the center of the room and look for items of interest around the room. These were large letters affixed to the walls of the room at a height of 1.5 m (*X*, *Y*, and *Z* in Fig. 8).

Besides success or failure at completing each component of the obstacle course, we measured the time required to complete the entire course. During this test, the scout was teleoperated by the same human operator. The results are given in Table 1. In the table, superscripts indicate different events during the trial while the numbers in the middle five columns indicate the frequency of a specific event. The maximum jump had a height of 0.25 m.

We also tested the audio unit and were able to receive audio even when the audio scout was 20 m away. The pan/tilt module was tested five times. The deployment lasted 10 s on the average and we were able to perform a tilt of 90° and a pan of 360°. The average time for pan was 4 s and for tilt was 2 s.

We also tried to find the narrowest corridor (0.2 m) that the scout could master. It traversed this corridor (length 1.2 m) without touching its walls in 9 s. We also checked the maximum range of communication (indoors) for the ranger-scout model. We found this range to be 9.5 m. With a scout as a repeater, this distance almost doubles.

Launching and Survivability Tests

The objective was to test the basic functionality of the ranger-scout system and the durability of the scout. First, the ranger launched a scout through a glass window (Fig. 8). The launching distance was 5 m and the height was 2 m. The scout needed to survive both the launch and landing as demonstrated by rolling and hopping when powered on after impact. We repeated this experiment twice. In both cases, the scout survived the impact. In two other tests, a human tossed a scout 21 m and 25 m. Similarly to the previous trials, the scout was functional when powered on after impact. In a final experiment, the ranger launched four scouts in a single room. In this experiment the scouts were launched already powered on. Three out of the four scouts were functional immediately after the launch while the fourth needed a manual power cycling to restart.

Operational Capabilities of the Scouts

In order to examine the scout's ability to hide itself in an environment and report back useful data, three different experiments were run.

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. These experiments were designed to examine the scout's behaviors in an analytical fashion.

For the first experiment, a controlled environment consisting of uniformly colored walls and a single dark object was constructed. An area, roughly a 2.5 m × 3 m rectangle, was defined. The target was a 1 m × 0.5 m black rectangle set up on one side of the environment. The robot was started 1.5 m away from the center of the target.

Nine experiments were run to see how long it would take the robot to locate the black target object and place itself next to 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 toward the target. Figure 9 shows the view from this camera as well as a superimposed plot of the path that the scout took to reach its objective. Figure 10 shows a plot of average distance the scout was away from the target versus time for all of these runs. In each case, the robot successfully located the target and moved toward it.

EXPERIMENT 2

The second experiment was set up to determine how well the scout could position itself in a more "real world" environment, meaning that of 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 toward 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 Fig. 11). The other was an office environment with chairs, a table, desks, cabinets, and boxes (see Fig. 12). The floor of the lab is a shiny, light tile surface of relatively uniform color whereas the floor of the office is a carpet of medium and dark piles providing a high localized contrast. These differences in surface brightness and contrast were accounted for in the scouts 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 Figs. 11 and 12).

In four of the five runs in the lab environment, the scout chose the same chair under which to hide (location 1 in Fig. 12). On run number 2, however, the scout wound up roughly 0.5 m out from under the chair in a relatively exposed position (location 2 in Fig. 11). In all five runs the scout ended up facing toward a relatively bright area of the room. However, in run 4 this happened to be toward the rear of the room. Times required for these runs are given in column 2 of Table 2.

Similarly, in four of the five runs in the office environment, the scout chose the same chair as its destination (location 1 in Fig. 12). On run 4 the scout chose the other nearby chair (location 2 in Fig. 12). In four of the five runs the scout wound

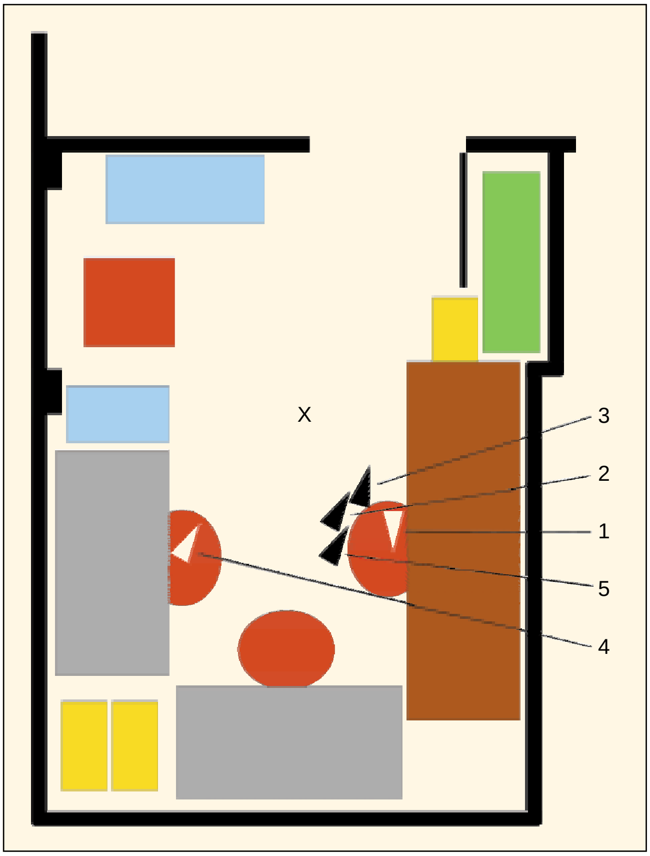

Figure 12. Experiment 2: Office environment, showing locations of scouts for all five runs. X marks the starting position used in all runs and numbered arrows correspond to final position and orientation for individual runs. Ovals represent chairs under which scouts may hide. Chairs are positioned at a table and at two desks, all of which also provide hiding opportunities. Other objects are impassable.

Table 2. Duration of Experiments 2 (Lab and Office Environments) and 3 (Coordinated Actions) (Time in Minutes)

Run	Lab Environment	Office Environment	Coordinated Actions
1	3	4	11
2	4	4	8
3	3	4	19
4	2	6	11
5	5	4	14

up facing brightly lit areas roughly toward the door of the office. On run 1, though, the scout faced a somewhat darker area toward the back of the room. Times required for these runs are given in column 3 of Table 2.

Problems due to poor radio communication between the scout and the ranger caused several runs to have to be aborted and restarted. Other times, the scout's batteries ran out and had to be replaced before the data collection could continue.

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. Rangers 1 and 2 are positioned as shown in Fig. 13(a). 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. The final positions of the rangers and scouts are shown in Fig. 13(b). Times required for these runs are given in column 4 of Table 2.

Analysis

One of the major issues of the whole system is the power consumption. We can currently perform eight jumps of 0.25 m on a set of nine 3 V batteries. A single scout in idle mode has a power draw of 1.725 W, while rolling over a level surface it has a power draw of 2.145 W, and during video transmission

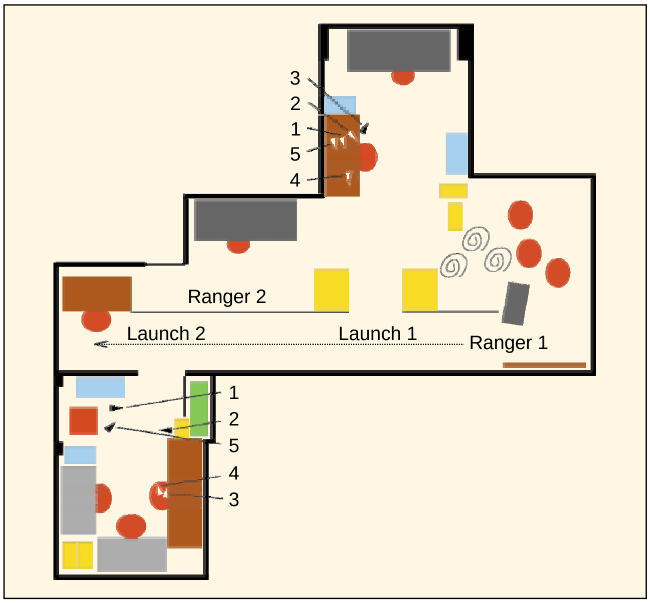

Figure 13. *Experiment 3: Coordinated agent behaviors, showing path taken by Ranger 1, launching sites, and final location of scouts on all five runs (numbered arrows showing orientation and position). Ovals represent chairs under which scouts may hide. Chairs are positioned at tables, desks, and lab benches, all of which also provide hiding opportunities. Other objects are impassable.*

the power draw is 2.040 W. However, the most expensive action is "winching in" the spring foot with a power draw of 3.465 W.

We plan to address the issue of power consumption in the following ways. (1) By reducing the speed of both CPUs to 20 MHz from 50 MHz we can save ~1 W. (2) By using 3.3 V parts we can save another 0.23 W. (3) By using a chopper power supply (versus linear) we can save another 0.15 W. Our goal is to have a power draw of 0.345 W in idle mode, which would give a five-fold or greater increase in battery life.

Due to the power requirements for jumping, we plan to revisit the design of the jumping mechanism. Hardening the scout is important in order for the scout to survive long-distance launching. Finally, we need to extend the communication range between the ranger and the scout, to improve the human-ranger-scout interaction, and address the miniaturization of other sensors that are useful payloads for the scout.

With respect to the software, major emphasis will be given to localization algorithms, improvements in the autonomous operation of the scout, the development of complex coordinated behaviors, and the human/computer interface issues.

Conclusions

We have presented an innovative miniature robotic system called the "scout," which is the basis of a large heterogeneous distributed robotic team. The scout effectively combines rolling and jumping locomotion capabilities. It has a small size that makes it suitable for several challenging reconnaissance and surveillance tasks. It can be deployed manually or by a launcher. The scout functions in conjunction with a larger utility platform, which is called the "ranger." Both systems carry a large number of processing, communication, control, and sensing modules that make the whole system an effective reconnaissance tool.

Acknowledgments

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

Keywords

Miniature robots, distributed robotics, mobile robots, teams of robots, reconnaissance and surveillance

References

[1] T.R. Baich and R.C. Arkin, "Behavior-based formation control for multiagent robot teams," *IEEE Trans. Robot. Automat.*, vol. 14, no. 6, pp. 926-939, Dec. 1998.

[2] Y. Gao, A. Fukunaga, and A. Kahng, "Cooperative mobile robotics: Antecedents and directions," *Auton.Robots*, vol. 4, no. 1, pp. 7-27, 1997.

[3] R. Dillmann, M. Kaiser, F. Wailner, and P. Weckesser, "PRIAMOS: An advanced mobile system for service, inspection, and surveillance tasks," in *Modelling and Planning for Sensor Based Intelligent Robot Systems* (volume 21 of *Series in Machine Perception and Artificial Intelligence)* Singapore: World Scientific, 1995.

[4] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *IEEE Computer*, vol. 22, no. 6, pp. 46-57, 1989.
[5] H. Everett and D. Gage, "From laboratory to warehouse: Security robots meet the real world," *Int. J. Robot. Res.*, vol. 18. no. 7, pp. 760-768, July 1999.
[6] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," *Autono. Robots,* vol. 8, no. 3, pp. 325-344, 2000.
[7] T. Kajiwara, J. Yamaguchi, Y. Kanayama, S. Yuta, and J. Lijima, "Development of a mobile robot for security guard," in *Proc.15th Int.Symp. on Industrial Robots*, vol. 1, pp. 271-278, Tokyo, Japan, 1985.
[8] A. Kochan. "HelpMate to ease hospital delivery and collection tasks, and assist with security," *Industrial Robot*, vol. 24, no. 3, pp. 226-228, 1997.
[9] M. Mataric, "Behavior-based control: Examples from navigation, learning, and group behavior," *J. Experimental and Theoretical Artificial Intelligence*, vol. 9, no. 2-3, pp. 323-336, 1997.
[10] T. Orwig, "Cybermotion's roving robots," *Industrial Robot*, vol. 20. no. 3, pp. 27-29, 1993.
[11] A. Osipov, V. Kemurdjian, and B. Safonov, "Mobile robots for security," in *Proc.1996 2nd Specialty Conf. Robotics for Challenging Environments*, RCE-II, pp. 290-295, Albuquerque, NM, 1996.
[12] L.E. Parker, "On the design of behavior-based multi-robot teams," *J. Advanced Robotics*, vol. 10, no. 6, pp. 547-578, 1996.
[13] C. Pellerin. "Twenty-first century sentries," *Industrial Robot*, vol. 20, no. 2, pp. 15-17, 1993.
[14] J. Porteous, "Intelligent buildings and their effect on the security industry," In Larry D. Sanson, editor, *IEEE Int. Carnahan Conference on Security Technology*, Sanderstead, Surrey, England, October 1995, pp. 186-188.
[15] D. Pritchard, R. White, D. Adams, E. Krause, E. Fox, M. Ladd, R. Heintzleman, P. Sprauer, and J. MacEachin, "Test and evaluation of panoramic imaging security sensor for force protection and facility security," in *Proc. IEEE Int.Carnahan Conference on Security Technology*, Alexandria, VA, October 1998, pp. 190-195.
[16] M. Saitoh, Y. Takahashi, A. Sankaranarayanan, H. Ohmachi, and K. Marukawa, "Mobile robot testbed with manipulator for security guard application," in *Proc.I EEE Int. Conference on Robotics and Automation*, Nagoya, Japan, 1995, vol. 3, pp. 2518-2523.
[17] C. Weisbin, J. Blitch, D. Lavery, E. Krotkov, C. Shoemaker, L. Matthies, and G. Rodriguez, Miniature robots for space and military missions, *IEEE Robot. Automat. Mag.* vol. 6, no. 3, pp. 9-18, Sept. 1999.

Paul E. Rybski received his B.A. in mathematics/computer science with an interdisciplinary emphasis in cognitive science from Lawrence University in Appleton, Wisconsin. He is currently working towards his Ph.D. in computer science with a graduate minor in cognitive science at the University of Minnesota, in Minneapolis. His research interests include behavior-based robotic control, process management in distributed heterogeneous robotic systems, and probabilistic spatial reasoning techniques.

Nikolaos P. Papanikolopoulos received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1987; the M.S.E.E. in electrical engineering from Carnegie Mellon University (CMU), Pittsburgh, PA, in 1988; and the Ph.D. in electrical and computer engineering from CMU in 1992. Currently, he is an associate professor in the Department of Computer Science at the University of Minnesota. His research interests include robotics, sensors for transportation applications, control, and computer vision. He has authored or coauthored more than 100 journal and conference papers in the above areas (29 refereed journal papers). He was a finalist for the Anton Philips Award for Best Student Paper in the 1991 IEEE Robotics and Automation Conference. Furthermore, he was recipient of the Kritski fellowship in 1986 and 1987. He was a McKnight Land-Grant Professor at the University of Minnesota for the period 1995-1997 and has received the NSF Research Initiation and Early Career Development awards. He has received grants from DARPA, Sandia National Laboratories, NSF, USDOT, MN/DOT, Honeywell, and 3M.

Sascha A. Stoeter obtained his M.S. in computer and information sciences in 1997 from the University of Minnesota. Before entering the Ph.D. program in Minnesota, he was a research assistant at the Institute for Robotics and Process Control in Braunschweig, Germany. He is a member of the Institute for Electrical and Electronic Engineers (IEEE) and Computer Professionals for Social Responsibility (CPSR).

Donald G. Krantz received his Ph.D. in computer and information sciences from the University of Minnesota in 1995. Dr. Krantz is a research scientist and engineer at MTS Systems Corp. and is a former Honeywell Engineering Fellow and Alliant Techsystems Fellow. He has helped to design six different autonomous and semi-autonomous large mobile robotic systems, including several for military applications. His research interests also include advanced manufacturing and material processing systems. Dr. Krantz is a winner of the 2000 Federal Laboratories Consortium for Technology transfer Award for his work on laser direct fabrication of titanium structures for aerospace applications.

Kemal B. Yesin received the B.S. degree in mechanical engineering from the Middle East Technical University, Ankara, in 1997 and the M.S. degree in mechanical engineering from the University of Minnesota, Minneapolis, in 2000. He is currently working toward a Ph.D. degree in mechanical engineering at the University of Minnesota. His research interests include meso-scale robotic systems and robotic assembly and manipulation of micro components.

Maria Gini is a professor in the Department of Computer Science and Engineering at the University of Minnesota in Minneapolis. Before joining the University of Minnesota, she was a research associate at the Politecnico of Milan, Italy, and a visiting research associate in the Artificial Intelligence Laboratory at Stanford University. Her research focuses on methods to distribute intelligence among robots or software agents. Her major contributions include robot motion planning, navigation, and negotiation strategies for multiple agents. She is on the editorial board of *Autonomous Robots* and *Integrated Computer-Aided Engineering* and she was the program cochair for the International Conference on Autonomous Agents 2000.

Richard Voyles received the B.S. in electrical engineering from Purdue University in 1983, the M.S. in manufacturing systems engineering from the Department of Mechanical Engineering at Stanford University in 1989, and the Ph.D. in robotics from the School of Computer Science at Carnegie Mellon University in 1997. He is currently an assistant professor in the Department of Computer Science and Engineering at the University of Minnesota. Dr. Voyles' industrial experience includes three years with IBM, where he was a manufacturing/test engineer working on automation projects, and one-and-a-half years with Integrated Systems, Inc., where he was a research scientist working on contracted applications of real-time software. Dr. Voyles founded Trident Robotics and Research, Inc. and cofounded Vigilant Technologies, Inc., both of which develop real-time control hardware.

Dean F. Hougen received his B.S. in computer science from Iowa State University in 1988, with minors in philosophy and mathematics, and his Ph.D. in computer science from the University of Minnesota in 1998, with a graduate minor in cognitive science. With over a decade of experience in all areas of artificial intelligence, his primary research involves robotics focusing on distributed heterogeneous robotic systems and situated learning in real robotic systems, including reinforcement learning, connectionist learning, and evolutionary computation. He has also worked in the areas of expert systems, decision support systems, geographic information systems, and user interfaces, developing several fielded systems in these areas. He is an assistant professor in the Department of Computer Science and Engineering and Associate Director of the Center for Distributed Robotics, both at the University of Minnesota.

Brad Nelson is the Mayhugh Associate Professor of Mechanical Engineering at the University of Minnesota. He received the B.S. degree (Bronze Tablet) in mechanical engineering from the University of Illinois at Urbana-Champaign in 1984, the M.S. degree in mechanical engineering from the University of Minnesota in 1987, and the Ph.D. degree in robotics (School of Computer Science) from Carnegie Mellon University in 1995. He holds a McKnight Land-Grant Professorship and is a recipient of the Office of Naval Research Young Investigator Award and the National Science Foundation Faculty Early Career Development (CAREER) Award. He has been an assistant professor at the University of Illinois at Chicago, has worked as an engineer for Honeywell Inc. and Motorola Inc., and has served as a United States Peace Corps Volunteer in Botswana, Africa. His research interests include biomedical engineering, computer vision, controls, manufacturing, mechatronics, MEMS, and robotics.

Michael D. Erickson did undergraduate work at North Dakota State University (electrical engineering major) and was a sergeant in the US Army for three-and-a-half years. He is currently a senior at the University of Minnesota. His interests include embedded systems design and integration, distributed operating systems, compiler optimizations, computer networks, parallel architectures, cryptography, and computer security.

Address for Correspondence: Nikolaos P. Papanikolopoulos, Center for Distributed Robotics, Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455.