Control of Multiple Small Surveillance Robots at AAAI 2000

Paul E. Rybski, Dean F. Hougen, Sascha A. Stoeter, Maria Gini and **Nikolaos Papanikolopoulos**

Center for Distributed Robotics, Department of Computer Science and Engineering

University of Minnesota

{rybski,hougen,stoeter,gini,npapas}@cs.umn.edu

Abstract

We describe our experiences demonstrating a team of miniature robots engaged in a surveillance mission at the AAAI 2000 Mobile Robots Exhibition. We present the architecture and the main features of our system and discuss some of the environmental factors that affect its performance.

Introduction

Reconnaissance and surveillance tasks can benefit from the use of multiple small yet highly capable robots. Such robots 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.

We have developed a set of extremely small robotic systems, called scouts (Hougen *et al.* 2000a), which were designed specifically to meet these requirements. The scout robot, shown in Figure 1, is a cylindrical robot 11 cm in length and 4 cm in diameter.

Scouts have two modes of locomotion designed to transport them over different kinds of terrain and obstacles. In the first mode, the scouts use their wheels to roll over smooth surfaces (even up a 20 degree slope). When confronted with an obstacle taller that itself, the scout employs its second locomotion mode, the jump. The spring-loaded tail of the scout is compressed and released to propel the robot over objects upwards of 20 cm in height. Figure 2 shows a scout jumping over a barrier.

Scouts carry a small video camera and video transmitter which they use to capture information about their environment. They can also transmit and receive digital information over a separate RF communications system which uses a packet-based communications protocol.

The small size of the scouts provides many advantages. They are inexpensive and easily transportable, which makes them ideal for use in large teams. This allows them to be present throughout a wide area, forming a mobile sensor network. It also allows individual scouts to be expendable without jeopardizing an entire mission. scouts are well suited to clandestine operations since they can be concealed easily.

Figure 1: Two scout robots (shown next to a quarter for scale). ©2000 by ACM, appeared in Rybski *et al.* (2000).

Their small size and cylindrical shape also allows them to be deployed by launching or throwing by hand.

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" (Porteous 1995). These may be portable to allow for rapid deployment (Pritchard *et al.* 1998) 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 (Kajiwara *et al.* 1985). Such systems are now available commercially and are in place in, for example, factory, warehouse, and hospital settings (Kochan 1997), and

Figure 2: A scout jumping over a barrier (sequence starts from the upper left corner). ©2000 by IEEE, appeared in Hougen *et al.* (2000a).

research continues along these lines (Dillmann *et al.* 1995; Osipov, Kemurdjian, & Safonov 1996). However, the single mobile agent is unable to be many in places at once—one of the reasons why security systems were initially developed. Further, 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 can do them faster and more reliably, as described in the extensive survey by Cao, Fukunaga, & Kahng (1997). Tasks traditionally accomplished with multiple robots are foraging (Matarić 1997), formation marching (Balch & Arkin 1998), map making (Burgard *et al.* 2000), janitorial service (Parker 1996), security (Everett & Gage 1999), and exploration (Hougen *et al.* 2000b).

Traditionally, mobile robots have ranged in size from roughly dog-sized to somewhat larger than a human. For small robots, Lego blocks and microprocessor boards, such as the Handyboard (Martin 1998), are often used to quickly prototype robots. We have designed multiple Lego-based robots and used them for a variety of navigation tasks (Rybski *et al.* 1998).

Recently a significant interest has arisen in designing even smaller mobile robots for exploration and reconnaissance. These include the popular Khepera robot (Mondada, Franzi, & Ienne 1993) and other prototypes, such as AL-ICE (Caprari *et al.* 1998) and our own scouts (Hougen *et al.* 2000a), developed by various research groups and not yet commercially available.

The major challenge in designing miniature robots is in fitting all the mechanical parts and electronics into a limited volume and in designing an adequate method of locomotion. Most miniature robots have wheels (Baumgartner *et al.* 1998), but some can jump (Halme, Schönberg, & Wang 1996), roll (Chemel, Mutschler, & Schempf 1999), fly (Wu, Schultz, & Agah 1999), or swim (Fukuda, Kawamoto, & Shimojima 1996). Our miniature robots have two rolling wheels but they can also jump to go over small obstacles, as shown in Figure 2.

Due to their small size, most miniature robots use proxy processing, as in Inaba *et al.* (1996), and communicate via a radio link with the unit where the computation is done. Limited communication bandwidth becomes a problem when robots have to transmit large amounts of data, such as live video, and when many robots need to share the bandwidth. Power consumption is another major problem which limits the mobility, ability to communicate, and long term survivability of miniature robots.

Scout Control Architecture

The scout's small size and limited power supply greatly restricts the kinds of on-board processing that can take place. The only computations done on the scout's two 8-bit microprocessors are what's necessary for communications and actuation. All high-level decision making and video image processing must be relegated to an off-board workstation or a human teleoperator.

Autonomous decision process, such as planners, reactive behaviors, or teleoperation controls, send low-level actuator commands to the scouts through an RF transceiver. Each scout is assigned a unique network ID which allows packets to be routed to the correct robot. By interleaving the transmission of packets destined for different robots, multiple scouts can be controlled simultaneously. The current bottleneck in the system is the number of packets per second that can be transmitted. The frequency at which commands can be broadcast is currently between 5 to 8 per second. The limiting factor is the speed at which the scout can decode the messages from the radio.

Scout video data is captured by a video receiver. This can be either displayed on a monitor or fed into a digitizer. Because the video is a continuous analog stream, only one robot can broadcast at any instant. Signals broadcast simultaneously from multiple robots would interfere with each other, making both signals useless. Scouts running in parallel must interleave their video transmissions and limit their transmissions to short bursts.

We have developed a distributed, CORBA-based (Group

1998) architecture to coordinate all of the individual hardware resources in our system. Access to robotic hardware and other computational resources is controlled through processes called RESOURCE CONTROLLERS (or RCs). Every physical resource is given its own RC process to manage it. Behaviors and decision processes connect to these RCs to send commands and receive data from them. Behavior processes running on different computers can access all of the RCs without even being aware that they may be running on different computers.

Scout Behaviors

The scouts are capable of operating autonomously through the use of simple behaviors. 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 a larger robot or a workstation 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. Figure 3 shows an example image received from the scout's camera. The RF noise degrading this image increases the difficulty of distinguishing the objects from the background.


Figure 3: The world from the scout's point of view. Here the scout is viewing a lab bench and two chairs at a range of 2m. ©2000 by ACM, appeared in Rybski *et al.* (2000)

A number of different visually-based behaviors have been implemented; see Rybski *et al.* (2000) for the full list. The behaviors demonstrated at the AAAI 2000 Mobile Robot Exhibition were:

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 to see that it has moved by computing differences between frames. Other kinds of motion makes use of this check because the scout does not always receive the commands sent for it due to RF interference. Frame differencing is not helpful in this behavior 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 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 the difference threshold until RF noise is filtered out. The scout then subtracts sequential images in the video stream and determines whether the scene changes at all (caused by movement in the image).

Demonstrations at the Exhibition

At the AAAI 2000 Mobile Robot Exhibition, we demonstrated several aspects of our distributed scout control system on two networked laptop computers. The first laptop ran Linux and handled all of the scout control programs. The second ran Windows and handled the video processing hardware and software. The scout's command radio was connected to the Linux laptop while a video receiver was connected to the Windows laptop. In addition to the autonomous behaviors defined above, a simple teleoperation client was used which allowed a human operator to drive a scout and operate its sensor payload and jumping mechanism. Figure 4 illustrates the implementation of the scout control system.
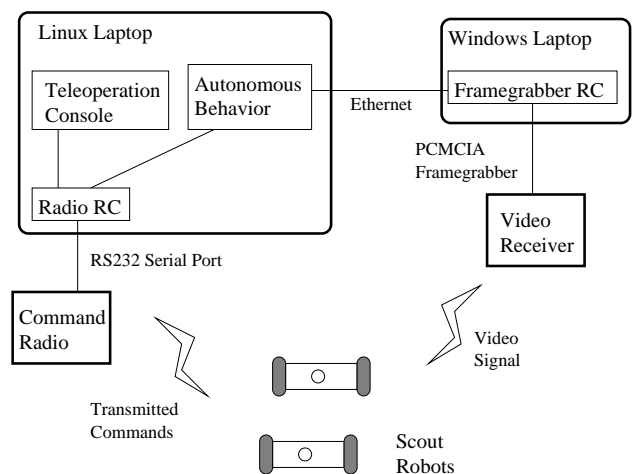

Figure 4: Diagram of the hardware and software architecture used to control the scout robots.

Several RCs were used by the behaviors and teleoperation consoles to control the scouts. The first, and arguably the most important, RC is in charge of controlling access to the radio. The user interface consoles and behaviors have to connect to the radio RC's CORBA interface to send packet requests to it. The radio RC supports two kinds of command packet requests. In the first request type, the decision process must explicitly instruct the radio to send each command. If commands are sent to the radio faster than they can be transmitted, they are stored in a queue until the radio can handle them. In the second request type, the decision process gives the radio a single command and instructs it to repeat it as often as it can. This allows the system to keep inter-process communications to a minimum. If multiple decision processes set up repeated commands, the radio will schedule transmission of each in a round-robin fashion. If the queue of non-repeated commands has pending requests, the queue will also be added to the schedule.

The second RC is the framegrabber process, which is required for autonomous behaviors. Images are captured by the framegrabber on the Windows laptop and are processed based on the requests of the decision process. Several different image operators were implemented, including frame differencing, connected region extraction and various statistical operators. To cut down on the amount of network overhead, all image processing was done by the framegrabber RC and only the processed image information was passed back to the decision process.

At the exhibition, the scout robots were demonstrated in both teleoperated and autonomous modes. In the teleoperation demo, we showed the rolling and jumping mobility modes of the robot, as well as demonstrating its usefulness as a remote camera. In the autonomous demos, the scouts served as motion detectors by transmitting video by doing frame differencing operations and notifying a user's console whenever motion was perceived in the image frame. We also demonstrated the scout's ability to servo to a dark object. For this behavior, images were analyzed to find the darkest mean value in the robot's visual field. The robot was commanded to drive itself towards the darkest objects and keep them centered in its visual field.

Discussion & Lessons Learned

Because the scout robots are completely dependent on the two RF communications channels for operation, any interference on those frequencies can seriously degrade performance. Because the command RF channel is packet-based (as opposed to streaming), it is somewhat resilient to local interference. The video channel, on the other hand, transmits continuous analog data which is very sensitive to interference. Corruption in the video signal can make it nearly unusable by autonomous control processes. The frequency that we use for our video transmitters is the 900MHz industrial/medical band and is, unfortunately, extremely popular. Several other groups at the exhibition used this frequency to transmit video (most notably the teleoperated blimp used by the BotBall competition). We had to work out a scheduling arrangement with the organizers of the exhibition to share the communications band.

The small physical size and battery supply of our robots severely limits the kind of RF transceivers that we can use. We are also severely limited in the kinds of video capture devices that we can use. In our current design of the robots, digital cameras were not an option because of lack of available hardware small enough to fit into our robot as well as lack of computational power to process an image even if we had such a camera. More attention will have to be paid to this issue in future scout designs. One solution will be to sacrifice robot size in order to include additional or more powerful computational resources. Some sort of a spread-spectrum or frequency-hopping radio system would be extremely useful as well. Not only would the higher frequencies involved allow for faster data transmission (critical for the transmission of video data), the ability to change frequencies would allow more robots to operate in the same area and would be more robust to interference.

We hope to be able to address some of these issues in the future as advances in technology make better and smaller equipment more available. Until this happens, we are focusing our efforts on algorithms for better video signal processing as well as decision processes which take the inherent fragility of the RF communications channels into account when controlling the robots.

References

Balch, T. R., and Arkin, R. C. 1998. Behavior-based formation control for multiagent robot teams. *IEEE Transactions on Robotics and Automation* 14(6):926–939.

Baumgartner, E.; Wilcox, B.; Welch, R.; and Jones, R. 1998. Mobility performance of a small-body rover. In *Int'l Symposium on Robotics with Applications, World Automation Congress.*

Burgard, W.; Moors, M.; Fox, D.; Simmons, R.; and Thrun., S. 2000. Collaborative multi-robot exploration. In *Proc. of the IEEE Int'l Conference on Robotics and Automation*, 476–481.

Cao, Y.; Fukunaga, A.; and Kahng, A. 1997. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots* 4(1):7–27.

Caprari, G.; Balmer, P.; Piguet, R.; and Siegwart, R. 1998. The autonomous micro robot ALICE: A platform for scientific and commercial applications. In *Proc. of 1998 Int'l Symposium on Micromechatronics and Human Science (MHS'98).*

Chemel, B.; Mutschler, E.; and Schempf, H. 1999. Cyclops: Miniature robotic reconnaissance system. In *Proc. of the IEEE Int'l Conference on Robotics and Automation*, 2298–2303.

Dillmann, R.; Kaiser, M.; Wallner, F.; and Weckesser, P. 1995. 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.

Everett, H. R., and Gage, D. W. 1999. From laboratory to warehouse: Security robots meet the real world. *Int'l Journal of Robotics Research* 18(7):760–768.

Fukuda, T.; Kawamoto, A.; and Shimojima, K. 1996. Acquisition of swimming motion by RBF fuzzy neuro with unsupervised learning. In *ALIFE V*, 31–37.

Group, O. M. 1998. *The Common Object Request Broker: Architecture and Specification*. Object Management Group.

Halme, A.; Schönberg, T.; and Wang, Y. 1996. Motion control of a spherical mobile robot. In *4th Int'l Workshop on Advanced Motion Control*.

Hougen, D. F.; Benjaafar, S.; Bonney, J. C.; Budenske, J. R.; Dvorak, M.; Gini, M.; Krantz, D. G.; Li, P. Y.; Malver, F.; Nelson, B.; Papanikolopoulos, N.; Rybski, P. E.; Stoeter, S. A.; Voyles, R.; and Yesin, K. B. 2000a. A miniature robotic system for reconnaissance and surveillance. In *Proc. of the IEEE Int'l Conference on Robotics and Automation*, 501–507.

Hougen, D. F.; Erickson, M. D.; Rybski, P. E.; Stoeter, S. A.; Gini, M.; and Papanikolopoulos, N. 2000b. Autonomous mobile robots and distributed exploratory missions. In *Int'l Symposium on Distributed Autonomous Robotic Systems*.

Inaba, M.; Kagami, S.; Kanechiro, F.; Takeda, K.; Tetsushi, O.; and Inoue, H. 1996. Vision-based adaptive and interactive behaviors in mechanical animals using the remote-brained approach. *Robotics and Autonomous Systems* 17:35–52.

Kajiwara, T.; Yamaguchi, J.; Kanayama, Y.; Yuta, S.; and Iijima, J. 1985. Development of a mobile robot for security guard. In *Proceedings of the 15th Int'l Symposium on Industrial Robots*, volume 1, 271–278.

Kochan, A. 1997. HelpMate to ease hospital delivery and collection tasks, and assist with security. *Industrial Robot* 24(3):226–228.

Martin, F. G. 1998. *The Handy Board Technical Reference*. MIT Media Lab, Cambridge, MA.

Matarić, M. 1997. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence* 9(2–3):323–336.

Mondada, F.; Franzi, E.; and Ienne, P. 1993. Mobile robot miniaturisation: A tool for investigation in control algorithms. In *Experimental Robotics III, Proc. of the 3rd Int'l*

Symposium on Experimental Robotics, 501–513. Kyoto, Japan: Springer Verlag, London.

Osipov, A.; Kemurdjian, V.; and Safonov, B. 1996. Mobile robots for security. In *Proceedings of the 1996 2nd Specialty Conference on Robotics for Challenging Environments, RCE-II*, 290–295.

Parker, L. E. 1996. On the design of behavior-based multi-robot teams. *Journal of Advanced Robotics* 10(6):547–578.

Porteous, J. 1995. Intelligent buildings and their effect on the security industry. In Sanson, L. D., ed., *IEEE Int'l Carnahan Conference on Security Technology*, 186–188.

Pritchard, D.; White, R.; Adams, D.; Krause, E.; Fox, E.; Ladd, M.; Heintzleman, R.; Sprauer, P.; and MacEachin, J. 1998. Test and evaluation of panoramic imaging security sensor for force protection and facility security. In *IEEE Int'l Carnahan Conference on Security Technology*, 190–195. Larry D. Sanson, ed.

Rybski, P. E.; Larson, A.; Lindahl, M.; and Gini, M. 1998. Performance evaluation of multiple robots in a search and retrieval task. In *Workshop on Artificial Intelligence and Manufacturing*, 153–160.

Rybski, P. E.; Stoeter, S. A.; Erickson, M. D.; Gini, M.; Hougen, D. F.; and Papanikolopoulos, N. 2000. A team of robotic agents for surveillance. In *Proc. of the Int'l Conf. on Autonomous Agents*, 9–16.

Wu, A. S.; Schultz, A. C.; and Agah, A. 1999. Evolving control for distributed micro air vehicles. In *Proc. IEEE Int'l Symposium on Computational Intelligence in Robotics and Automation*.