PERFORMANCE OF MULTIPLE TASKS BY AN AUTONOMOUS ROBOT USING VISUAL AND ULTRASOUND SENSING*

M. Beckerman, D. L. Barnett and M. Dickens
Center for Engineering Systems Advanced Research
Oak Ridge National Laboratory
P.O. Box 2008
Building 6025, MS-6364
Oak Ridge, TN 37831-6364

and

C. R. Weisbin
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109-8099

CONF-900780--2

DE90 010719

Invited Presentation: 3rd ISRAM International Symposium on Robotics and Manufacturing, British Columbia, Canada, July 18-20, 1990

* Research sponsored by the Engineering Research Program at the Office of Basic Energy Sciences, of the U.S. Department of Energy, under Contract No. DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

PERFORMANCE OF MULTIPLE TASKS BY AN AUTONOMOUS ROBOT USING VISUAL AND ULTRASOUND SENSING

M. BECKERMAN, D. L. BARNETT and M. DICKENS
Oak Ridge National Laboratory, Center for Engineering Systems Advanced Research, Building 6025, MS-6364, P.O. Box 2008, Oak Ridge, TN 37831-6364

C. R. WEISBIN
Jet Propulsion Laboratory, Robotics and Automation Section, 4800 Oak Grove Drive, Pasadena, CA 91109-8099

ABSTRACT

While there have been many sucessful mobile robot experiments, only a few papers have addressed issues pertaining to the range of applicability, or robustness, of robotic systems. The purpose of this paper is to report results of a series of benchmark experiments done to determine and quantify the robustness of an integrated hardware and software system of a mobile robot.

INTRODUCTION

HERMIES-IIB is one of a series of mobile robots developed at the ORNL Center for Engineering Systems Advanced Research (CESAR) [1]. The hardware for this robot includes a mobile platform, a sensor turret equipped with an array of twenty-four ultrasonic range sensors and three CCD cameras, and dual manipulator arms. It also includes a 16-node hypercube computer and an AT host computer.

The integrated system, to be described herein, was designed to perform a number of tasks sequentially. First, the robot had to use its ultrasound sensors to navigate through an obstacles strewn field to a preselected intermediate goal. Next, the robot had to use its CCD cameras to identify a control panel. It then had to navigate up to and dock in front of that control panel. In the last phase, the robot had to read analog meters and lamp conditions, and guided by the vision system, manipulate buttons and slides on the control panel.

By robustness we mean stability against sensor and mechanical errors, uncertainties and limitations for a specified range of environmental conditions. To determine the robustness of the integrated system we examined the robustness of each component subsystem. This was done in a series of benchmark experiments which delineated the environmental conditions under which successful operation could be achieved.

We begin by briefly describing the subsystems and connected tasks, noting the inputs and outputs from each component. We then turn to the experiments and results, summarizing the environmental conditions, problem areas and issues.

SUBSYSTEMS AND CONNECTED TASKS

Ultrasound Data Processing

The HERMIES-IIB ultrasound sensors provided 120 samples of range data covering 360 degrees in three degree steps. These data represented the distance to the nearest object intercepted by the beam. In processing the data, maps were not used. Instead, for each scan lists of paired and unpaired edges were constructed and then used to generate a list of corridors of free space.

Path Planning

The design of a path to a goal had to take into account the incompleteness of the information obtained from any single 360 degree ultrasound scan. Part of the space will lie beyond the effective horizon while other regions will be occluded by objects. Since the robot can only navigate through space known to be clear, the HERMIES-IIB paths were built in several stages using a local path planner. At each stage a new sonar scan was taken, and the next portion of the path was planned.

Panel Identification

The control panel identified by the robot contained a pair of analog meters, a danger light, four buttons and lamps, and two slides. At large distances the analog meters were the only features other than the outline of the panel itself which could be discerned. The function of the vision software during navigation was to identify the panel, and provide distance and angle information for navigation to the docking zone in front of the control panel (Fig. 1).

The CCD cameras mounted on the HERMIES-IIB robot provided 256 x 256 pixels of 8-bit grey-level intensity values from a 60 degree field-of-view. Concurrent processing of the visual data was done on-board the robot using the 16-node hypercube computer. In the low-level parallel processing the grey-scale image was converted to a binary image. A list of the labelled, connected regions in the binary image was then constructed. Two lists were produced, one for black regions and one for white regions. The algorithms for doing this have been described in detail elsewhere [2]. Geometric properties and low-order moments which were then extracted for each region.

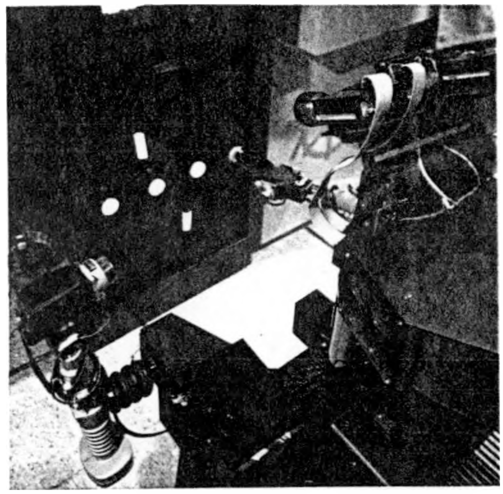

Fig. 1. HERMIES-IIB at the docking position, pushing a button.

In order to identify the control panel with its two meters, the list of connected regions in the image was pruned until only one candidate remained. The entry in the list sought was a dark region (control panel) containing two light regions (meters), situated side-by-side in the upper part of the dark region. A typical pruning procedure is illustrated in Figs. 2-5.

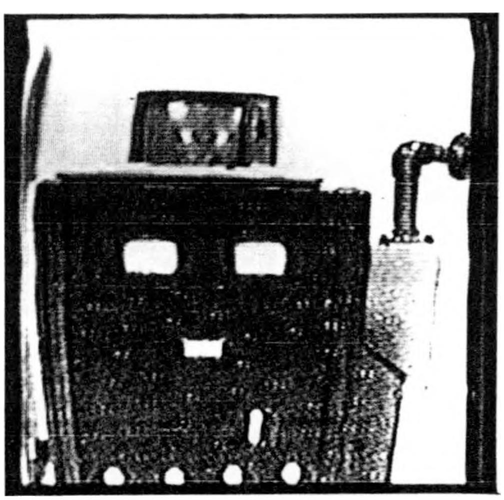

Fig. 2. Visual image of the control panel from a HERMIES-IIB CCD camera.

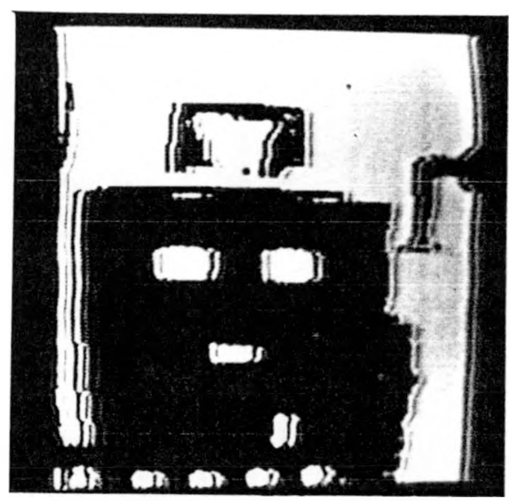

Fig. 3. Early intermediate processing stage: note that a shadow on the right and part of the object on the left are merged with panel.

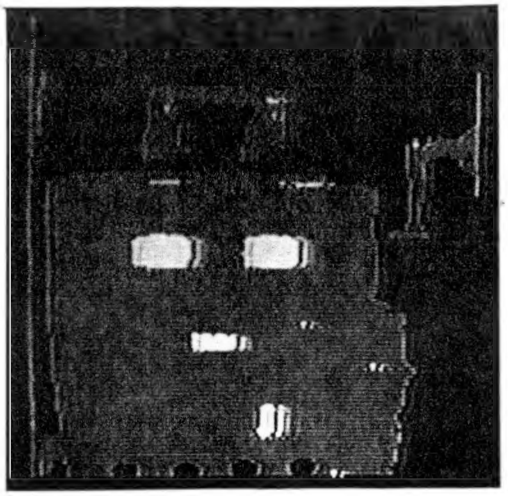

Fig. 4. Later processing stage: the many features which have been pruned are shown as black areas.

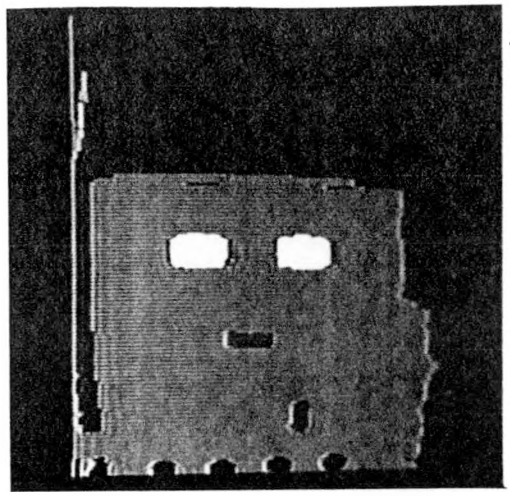

Fig. 5. Final image: the only remaining features are the large connected region (panel) containing two smaller regions (meters) side-by-side.

Docking

Estimates of distance, angle and yaw were made by the robot using the geometric information about the two control panel meters, as extracted from the visual image. Pixel dimensions were converted to spatial dimensions using prior information on the physical dimensions of the meters and the optics of the detection system.

Below 10 degrees, yaw estimates were unreliable. As a result of the inaccuracies in determining yaw at small angles, estimates of this quantity were not used for docking. Navigation to the docking zone was done in several (3-5) steps. In each step, a picture was taken, and a new distance and angle were extracted. The last image was acquired from 3-4 ft., roughly in front of the panel, where highly accurate distance measures could be obtained by the robot.

Vision-Guided Manipulation

The manipulation tasks consisted of pushing one or more of the four buttons, and moving either one or both of the slides. There were two main problems associated with reliable operation of the Heathkit Hero arms used for manipulation, namely: (i) there were no absolute encoders, and (ii) there was considerable backlash in the shoulder and wrist joints. The magnitude of the backlash was reduced by mechanically loading the arms. However, the remaining backlash was still large enough to lead to failures in manipulation unless controlled in some manner. The vision system functioned within a closed-feedback loop to first locate the various control panel features and then guide the movement of the manipulator.

A three-step procedure was used to position the end-effector at the goal location. In each step two CCD cameras mounted on the robot chassis sensed the location of the buttons and slides, and (in the same scene) the position of a light-emitting-diode (LED) mounted near the tip of the two-jaw grippers of the manipulator. Two pictures were taken with each camera, one with the LED on and one with the LED off. The two images for each camera were subtracted from one another to produce images containing only the dot representing the location of the LED. A stereo algorithm [3] was then used to obtain the (x,y,z) coordinates of the end-effector and the features to be manipulated.

EXPERIMENTS AND RESULTS

Environment

1. Ultrasound: Four actions were taken to reduce the effects of ultrasound distortions and specular reflections. First, the boxes which served as obstacles in the ultrasound-guided navigation phase were covered with a material to reduce the specular reflections. Second, a 4.5 m sensing horizon was imposed. Third, there was always at least one corridor whose actual width was equal to or greater than 1.5 m. This was done to take into account both the finite size of the robot (0.9 m) and the distance-dependent distortion errors. Fourth, corridors near the hard, smooth (specular reflecting) walls of the lab were excluded from consideration.
2. Visual: Displayed in Fig. 6 is a photo of a typical setup for the vision-guided phases of the experiments. The control panel is the dark object placed against the wall, as seen in the center of the background of the photo. A monitor plus keyboard has been placed on top of the control panel, and there is clutter to the left and the right. The photo was taken using floodlights. In the experiments only standard overhead fluorescent lighting was used. Some of the fluorescent lamps were situated near air conditioning ducts, and the average illumination varied considerably from place to place in the lab.
3. Manipulation: The length of each slide canal of the control panel was 193 mm. This length was partitioned into low, medium and high segments for the manipulation tasks. The depth accuracy required for the slide operations was given by the 5 cm interior opening of the gripper; the height of the slides was 5 cm. The button radius was 2 cm, and the allowed variation in depth was 1.5 cm about the mean value needed to activate the appropriate lamp. The buttons and levers were

made of plastic, light enough to be manipulable, and no torque was applied to them by the manipulators.

4. Navigation: One to four boxes, tall enough to intercept the sonar viewing plane were used for testing obstacle avoidance. Navigational devices such as beacons were not added to the environment. Instead, the robot moved over the smooth, level floor of the laboratory using dead reckoning. The control panel yaw, however, was fixed for panel identification and docking.

Fig. 6. Photo of the experimental setup used in the vision-guided navigation stages of the experiments.

Experimental Results

In order to perform the manipulation tasks the robot had to dock within a narrow zone centered in front of the control panel. This zone was determined experimentally to be three inches deep, 19 to 22 inches from the panel, and two inches wide. For both manipulators to operate the control panel the yaw had to be less than 2 degrees.

Distance errors, due to the intrinsic resolution of the CCD arrays and discretization errors of the digitizers, decreased from 10 to 20% at 4 m to 1% at 1 m. By using an incremental algorithm the robot was able to dock successfully from any transition point within the field-of-view of the control panel over distances from 4.5 m to 0.5 m.

At a distance of 60 cm the absolute accuracy of the stereo algorithm used for manipulation was approximately 50 mm with a repeatability of 5 mm. Although the absolute accuracy was marginal, good relative accuracy was achieved by locating the LED and panel feature in the same image. This relative accuracy was ultilized in the three-step positioning.

In the most recent series of experiments the robot was successful in carrying out all its tasks in 14 out of 20 trials. Three failures were attributable to one wheel sticking slightly to produce a 5 degree yaw at docking. Two of the three other failures were manipulation failures; in each case one of the buttons was not pushed in far enough to turn on the lamp. The remaining failure was due to an error in reading the condition of the control panel. These failures were typical of those encountered in other series of experiments done by us over a six month period. In the absence of mechanical problems the robot's success rate for the series of tasks was high, exceeding 90%.

The panel identification algorithm was able to recognize situations where the control panel was not in the field-of-view of the CCD camera. It was also able to handle lighting situations which produced panel shapes such as those illustrated in Figs. 2-5. However, panel identification problems were encountered when glare was produced from highly reflective metal surfaces. Other identification problems were produced by large shadows which were merged with the control panel. To achieve the above-mentioned high success rate some restrictions were imposed upon glare and shadows, and upon background clutter.

Similarly, ultrasound distortions and specular reflections led to failures in path planning for some obstacle configurations such as "picket fences" and "box canyons." In those instances the planner was unable to find either valid edges or free corridors.

CONCLUDING REMARKS

We have attempted to take into account sensor and mechanical errors, uncertainties and limitations in our development of the integrated system described above. The performance of tasks in several steps using sensor feedback in each step overcame many of these problems, and success rates in excess of 90% were achieved. It should be noted that, in large measure, the mechanical problems encountered pertain to an in-house-built prototype, and not to a commercial or production-level system.

Situations were encountered in interpreting the processed binary images and ultrasound data, where there was a lack of sufficient information for the robot to make intelligent decisions. In those situations we need to make better use of the available resources. In work being done in parallel to the present study [4,5] we are investigating ways to combine data from several scans, views, and sensor domains, and to extract more information from the data available.

We have described the factors in the experimental setup which could vary rather freely. With respect to these factors the performance of the robot was robust. We have also indicated those factors which were controlled to some extent and we have described the information about the environment built into the software. From this second viewpoint the performance was not as robust as we would wish. We are now attempting to develop the means to reduce these dependences.

ACKNOWLEDGMENTS

This research was supported by the Engineering Research Program at the Office of Basic Energy Sciences, of the U.S. Department of Energy, under Contract No. DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc. The authors wish to thank S. M. Killough and J. P. Jones for helpful discussions.

REFERENCES

1. B. L. Burks, G. de Saussure, C. R. Weisbin, J. P. Jones and W. R. Hamel, "Autonomous Navigation, Exploration and Recognition Using the HERMIES-IIB Robot," IEEE Expert, Vol. 2, pp. 18-27, 1987.

2. J. P. Jones, "A Concurrent On-Board Vision System for a Mobile Robot," in Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications, G. Fox (ed.), Vol. II, pp. 1022-1032, 1988.

3. S. M. Killough and W. R. Hamel, "Sensor Capabilities for the HERMIES Experimental Robot," Proceedings of the ANS Third Topical Meeting on Robotics and Remote Systems, Charleston, SC, Sec. 4.1, pp. 1-7, 1989.

4. M. Beckerman and E. M. Oblow, "Treatment of Systematic Errors in the Processing of Wide Angle Sonar Sensor Data for Robotic Navigation," IEEE Trans. on Robotics and Automation, Vol. RA-6, April, 1990.

5. M. Beckerman, L. A. Farkas and S. E. Johnston, "Treatment of Systematic Errors II: Fusion of Ultrasonic and Visual Sensor Data," Oak Ridge National Laboratory Technical Memo ORNL/TM-11349 (CESAR -89/38), (in preparation).