CONF-890522--6

DE89 008740

SENSOR INTEGRATION USING CONCURRENT COMPUTING ON-BOARD THE ORNL MOBILE ROBOT*

R. C. Mann
J. P. Jones
M. Beckerman
C. W. Glover
L. Farkas
J. R. Einstein

Advanced Computing and Integrated Sensor Systems
Center for Engineering Systems Advanced Research
Oak Ridge National Laboratory
Oak Ridge, Tennessee USA

Paper To Be Presented At:

IEEE Conference on Robotics and Automation
Workshop on "Parallel Algorithms and Architectures for Robotics"
Scottsdale, Arizona

May, 1989

*Research sponsored by the Office of Basic Energy Sciences and the Office of Technology Support Programs, U. S. Department of Energy under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

SENSOR INTEGRATION USING CONCURRENT COMPUTING ON-BOARD THE ORNL MOBILE ROBOT*

R. C. Mann, J. P. Jones, M. Beckerman, C. W. Glover,
L. Farkas, J. R. Einstein

Advanced Computing and Integrated Sensor Systems
Center for Engineering Systems Advanced Research
Oak Ridge National Laboratory
P. O. Box 2008
Oak Ridge, TN 37831-6364
mnn@stc10.ctd.ornl.gov

Abstract

The mobile robot prototypes developed at the Center for Engineering Systems Advanced Research (CESAR) at the Oak Ridge National Laboratory (ORNL) are equipped with sonar sensors, CCD cameras and a laser range camera that are used to support autonomous navigation and inspection tasks in an a priori unknown and unstructured dynamic environment. This paper summarizes work directed at extracting information from data collected with these sensors and integrating it, in order to produce reliable descriptions of the robot's environment. The approach consists in studying different world models and mappings among them, sensor models and parallel algorithms for sensor information processing, and appropriate integration startegies. Specifically, the paper describes the integration of two-dimensional vision and sonar range information, and the integration of laser range and luminance images.

1. Introduction

One of the prerequisites for intelligent behavior in robotic systems is the ability to generate consistent, system-internal representations of the environment. In general, this is impossible on the basis of any single sensor domain. Hence, robotic systems are being equipped with an increasing number of different sensors that supply partly redundant information. Multi-sensor integration (MSI) designates the task of combining data and information from these various sensors so that a consistent world model, i.e., a model free of contradiction, can be generated, on the basis of which decisions concerning

* This research was supported by the Office of Basic Energy Sciences, and the Office of Nuclear Energy, Office of Technology Support Programs, U.S. Department of Energy, under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

navigation, manipulation, etc., can be made. Our efforts in MSI are based on a framework that consists of one or several world models, sensor models, and appropriate integration strategies.

A world model is a geometric representation of items in the robot's environment along with spatial relationships among them. Simple representations can consist of 3-D line segments[1], or points, edges, and surface patches.[2] More elaborate world models can consist of complete 3-D geometric databases.[3] In any case, the world model also contains any a priori information on the robot's environment that might be available. The required complexity of the world model is determined by the robot's mission and diversity of tasks to be accomplished. For navigation with a robot that is constrained to move on the floor, a 2-D rectangular map can be considered appropriate. If that same robot must perform even simple manipulation tasks, a 3-D world model is required. Therefore, it is necessary to consider world models of different complexity and mappings among them.

Each sensor is capable of providing an estimate of some feature, e.g. range, orientation, intensity, of some of the items in the world model. It is necessary to quantify this capability so that the quality of the individual sensor estimates and the resulting overall estimates can be assessed. Therefore, models for each individual sensor are needed that describe how reliably a sensor is capable of providing a piece of information from the collected data. One way to specify the sensor model is through likelihood functions that give the probabilities for the observed data if certain features are present in the environment. Sensor models can also include rules that determine under which environmental conditions a sensor can be relied upon, e.g. maximum range and accuracy specifications for ultrasonic transducers.

Integration strategies are required to combine the individual sensor estimates. Such strategies must accomplish the removal of both random and systematic errors. They can be based on standard statistical criteria, e.g., maximum a posteriori (MAP), minimum square error (MSE), maximum likelihood (ML) estimators, or can include the solution of consistent labeling problems, and can also involve a set of rules that determine when to discard certain sensor estimates. MAP integration methods have recently been receiving much interest in the context of integrating different visual information based on Markov Random Field prior models[4,5] and powerful stochastic optimization algorithms[6] as well as deterministic variants thereof[7].

This paper is organized as follows. In section 2 we give a brief description of the ORNL mobile robot HERMIES-IIB, which we have been using as a testbed for our research. In section 3 we present results from our work on 2-D robot vision using the on-board hypercube multi-processor, and sonar range data processing, and describe the method we use to integrate 2-D vision and sonar into a 2-D world map and into a vertical edge database generated from vision data. The section also contains a brief description of the integration of registered luminance and laser range images based on Markov Random Field image models.

2. The ORNL Mobile Robot

HERMIES-IIB is one of a series of research robot prototypes designed at ORNL/CESAR for autonomous operation in unknown and possibly hazardous environments.[8] The robot stands 1 m high and weighs 91 kg. Rechargeable batteries supply 20 W power, allowing for 1 h of untethered operating time. Peak movement speed is 0.7 m/s. Sensors include two Sony CCD cameras and an array of Polaroid sonar transceivers mounted on a rotatable turret. The computer architecture (see Fig. 1) consists of a VME rack housing a variety of I/O boards interfaced via a BIT-3 communication link to an NCUBE hypercube computer (NCUBE Corp., Beaverton OR.) featuring 16 nodes with 512 Kbyte RAM memory for each node and an Intel 80286-based I/O processor, which also serves as a host for the hypercube. Each node is a 32-bit microcomputer with on-chip floating point and communications hardware. This gives HERMIES-IIB a computing power of approximately 32 MIPS in the on-board hypercube.

3. Sensor Integration

Our research in intelligent sensor systems at ORNL/CESAR makes use of a unique parallel computing environment built around a 64-node NCUBE hypercube supercomputer with a high-speed parallel interface to a VME-based Motorola 68020 system for fast sensor data acquisition and transfer to the hypercube. Algorithms implemented on this system are extensively tested and then ported to and integrated into the data processing system on-board HERMIES-IIB. In this section we summarize results of our work in robot vision and sonar information processing, and describe the integration of 2-D vision and sonar range information in order to support robot navigation.

3.1 Robot Vision

A 2-D robot vision system has been developed that integrates a series of low- , medium-, and high-level functions executed in parallel on the hypercube[9,10,11] that enable the robot to analyze image frames during navigation as well as to recognize an object of interest (e.g. a control panel) and perform simple inspection tasks (e.g. reading analog gauges).[8,9] The system uses two calibrated cameras. Images used during robot navigation are acquired through a camera with a wide angle lens (4.8 mm), images used for close inspection are acquired through a longer lens (16 mm). For the MSI task discussed here, we use modules of this vision system that detect vertical edges, and generate a list of labeled connected vertical edges and their locations in image coordinates. This database takes the form of a linked distributed list. Each record in the list contains information about a vertical edge (its length, location, orientation, and type, e.g. roof edge, step edge), plus the addresses of records and hypercube nodes containing information on neighboring edges. This information is passed on to the integration module described below. The use of vertical line segments as geometrical primitives is appropriate since the laboratory environment in which HERMIES-IIB navigates contains only box-shaped obstacles that are 1-2 ft wide and 3-4 ft high (see Fig. 2 (a)).

3.2 Ultrasonic Sensing

Ultrasonic sensing for robot navigation has been studied extensively using the HERMIES IIB robot. In this active sensing process ultrasonic beams are sent out from an array of sensors mounted on the robot, and the range to the nearest object intercepted by each sonar beam is determined from the overall time-of-flight of the corresponding returned signal. The ultrasonic beam is approximately 18 degrees wide. It provides a highly efficient means of scanning for empty spaces in which to navigate. However, the data are somewhat sparse, and in many real-world situations they are difficult to interpret correctly (see Fig 2 (c) for a typical sonar scan).

We have described a solution to the problem of how to treat systematic errors which arise in the processing of sparse sonar data.[12] In our methodology, pixels of the world model (a 2-D map of the robot's environment) are assigned one of several labels during the initial processing. One of the labels flags conflict among interpretations from two or more sensor measurements. This happens whenever there are erroneous interpretations of the data. The data are then reinterpreted, making use of pattern analyses and consistent-labeling operations to effect the removal of errors. In their study of sonar mapping and navigation, Elfes and Moravec[13,14] introduce 2-D maps which they call "certainty grids". The contents of our maps differ from those of Elfes and Moravec. Moreover, our methodology is explicitly non-local whereas theirs is local.

3.3 Integrating Sonar and 2-D Vision

Initially, the 2-D map is generated from sonar data and the vertical edge database is built from vision data. We obtain information in polar coordinates (r, theta), with the robot at the origin of the coordinate system, from the sonar sensor (see Fig. 2 (c)), and (theta, phi) information, where phi is the azimuth angle, from the vision sensor. The overlap in theta allows us to establish a correspondence between object features in the two disparate sensor domains. The range information (r) is passed to the vision processing software which uses it to provide an absolute scale for the image coordinates, and to match different edges of an object. Once a set of labeled edge segments is identified as possibly belonging to the same object a weighted least squares linear fit and chi-squared analysis are performed to determine acceptability of the assignment. The range information and chi-square evaluations are used to refine the edge assignment. Sets of edges are then extracted and stored as delineating an object's surfaces in a second linked list. These results are also returned to the sonar processing software. They are combined with existing sonar data, and used to correct and more precisely delineate object boundaries in the 2-D navigation map. Figure 3 shows a block diagram of the operations involved in the integration process. At the end of the integration process, information from both the vision and ultrasonic sensing domains appears in each representation.

3.4 Integrating Luminance and Laser Range Images

The range sensor used in our laboratory is a time-of-flight sensor manufactured by Odetics (Odetics Corp., Anaheim, CA). It returns the distance from the sensor point to the surface being observed as well as measuring the brightness of the returned beam. The range data is coarsely quantized, with a depth resolution of just under one inch per pixel. The brightness data provides a much smoother representation of surfaces, but is corrupted by both random and correlated (scan line artifacts) noise. We are studying an approach which consists in assuming that the underlying surface geometry is correct, and in determining an estimate of that surface by maximizing the a-posteriori probability (MAP restoration) of the particular surface, given the measurements.

The restoration process consists in defining an objective function which describes the problem, and then in minimizing that objective function. Experiments are still being performed to determine the best such function, however, the following approach has given promising results in initial tests[5]:

$$H = H_z + H_b + H_f$$

where

$$H_z = \Sigma_i (z_i - z_{oi})^2 / 2 s_z$$

i.e., the restored range image z should resemble the measured z_o;

$$H_b = \Sigma_i (b_i - b_{oi})^2 / 2 s_b$$

i.e., the restored brightness image b should resemble the measured b_o;

$$H_f = 1/ (T)^{1/2} \Sigma_i \exp (-(bi - R(zi))2 /T)$$

i.e., the restored brightness b, should agree with the reflectivity model R, applied to the restored range values.

In addition, either H_z or H_b may be augmented with a term which incorporates a priori information concerning the local smoothness of surfaces.

4. Summary

The ability to construct accurate, consistent representations of its operational environment from multiple sensors is a key requirement for a mobile robot with advanced autonomous capabilities. We have described a method for integrating sonar range and vision information to support robot navigation. The methodology is based on a framework for MSI that consists of multiple, distributed world models and appropriate mappings among them; sensor models including appropriate, sensor-specific signal and information processing algorithms; and integration strategies. For the specific case presented here, we used a linked list of geometrical features extracted from intensity images as the 2-D vision sensor world model, and a 2-D world map for the sonar sensor domain. Information from both sensors was used to update elements in both world models. The method is implemented on a multi-processor computer architecture on-board the ORNL/CESAR mobile robot HERMIES-IIB.

We briefly outlined our approach to integration of laser range and luminance data in section 3.4. This method is currently undergoing testing with data collected with the Odetics laser camera in order to address some of the issues that can affect the integration in a critical way, e.g., the scene reflectivity model R(z) and the model for noise in the range data (H_z).

Future MSI systems will be used not only for robot navigation but also to provide the intelligent sensing capabilities required for complex robot manipulation tasks.

5. References

[1] N. Ayache, O. D. Faugeras, "Building, registrating, and fusing noisy visual maps", Proceedings of the First International Conference on Computer Vision, London, June 1987, pp 73-82.

[2] P. K. Allen, "A framework for implementing multi-sensor robotic tasks", Proceedings of the ASME Conference on Computers in Engineering, New York, August 1987, pp 303-309.

[3] C. Crane, J. Duffy, "An interactive animated display of man-controlled and autonomous robots", Proceedings of the 1986 ASME Computers in Engineering Conference, Chicago, July 1986.

[4] T. Poggio et al., "The MIT vision machine", Proceedings of the 1988 DARPA Image Understanding Workshop, Cambridge, MA, April 1988, pp 177-198.

[5] G. L. Bilbro, H. Hiriyannaiah, W. E. Snyder, "Fusion of range and reflectance image data using markov random fields", Proceedings of the IEEE International Symposium on Intelligent Control, Arlington, VA, August 1988.

[6] S. Geman, D. Geman, "Stochastic Relaxation, Gibbs distributions, and the Bayesian restoration of images", IEEE Trans. Pattern Anal. and Mach. Intell., Vol 6, 1984.

[7] G. L. Bilbro, R. C. Mann, T. K. Miller, W. E. Snyder, D. E. Van den Bout, "Simulated annealing using the mean field approximation", Neural Engineering Technical Report 88-2, North Carolina State University, (1988).

[8] B. L. Burks, G. de Saussure, C. R. Weisbin, J. P. Jones, W. R. Hamel, "Autonomous navigation, exploration, and recognition using the HERMIES-IIB robot", IEEE Expert, Vol. 2, pp 18-27, (1987).

[9] J. P. Jones, R. C. Mann, "Concurrent algorithms for a mobile robot vision system", SPIE Vol. 937, pp 497-504, (1988).

[10] J. P. Jones, R. C. Mann, E. M. Simpson, "A computer vision system for a hypercube concurrent ensemble", ORNL/TM-10679, (1988).

[11] J. P. Jones, K. M. Clinard, "Parallel relaxation matching", abstract to appear in Proceedings of the 1988 Applied Imagery and Pattern Recognition Workshop, Washington D.C., (1988), manuscript in preparation.

[12] M. Beckerman, E. M. Oblow, "Treatment of systematic errors in the processing of wide-angle sonar data for robotic navigation," ORNL/TM-10763, (1988), submitted to the IEEE Journal on Robotics and Automation.

[13] H. P. Moravec, A. Elfes, " High resolution maps from wide angle sonar", Proceedings of the IEEE International Conference on Robotics and Automation, pp 116-121, (1985).

[14] A. Elfes, "Sonar-based real-world mapping and navigation", IEEE Journal of Robotics and Automation, Vol. RA-3, No. 3, pp 249-265, (1987).

[15] M. Beckerman, J. P. Jones, R. C. Mann, L. A. Farkas, S. E. Johnston, "Spatial reasoning in the treatment of systematic sensor errors", Proceedings of the SPIE Conference on Sensor Fusion: Spatial Reasoning and Scene Interpretation, Cambridge, MA, November 6-11, 1988, in press.

[16] M. Beckerman, L. A. Farkas, J. P. Jones, R. C. Mann, C. W. Glover, "World modeling and multi-sensor integration for a mobile robot", Proceedings of the ANS Third Topical Meeting on Robotics and Remote Systems, Charleston, SC, March 13-16, 1989, in press.

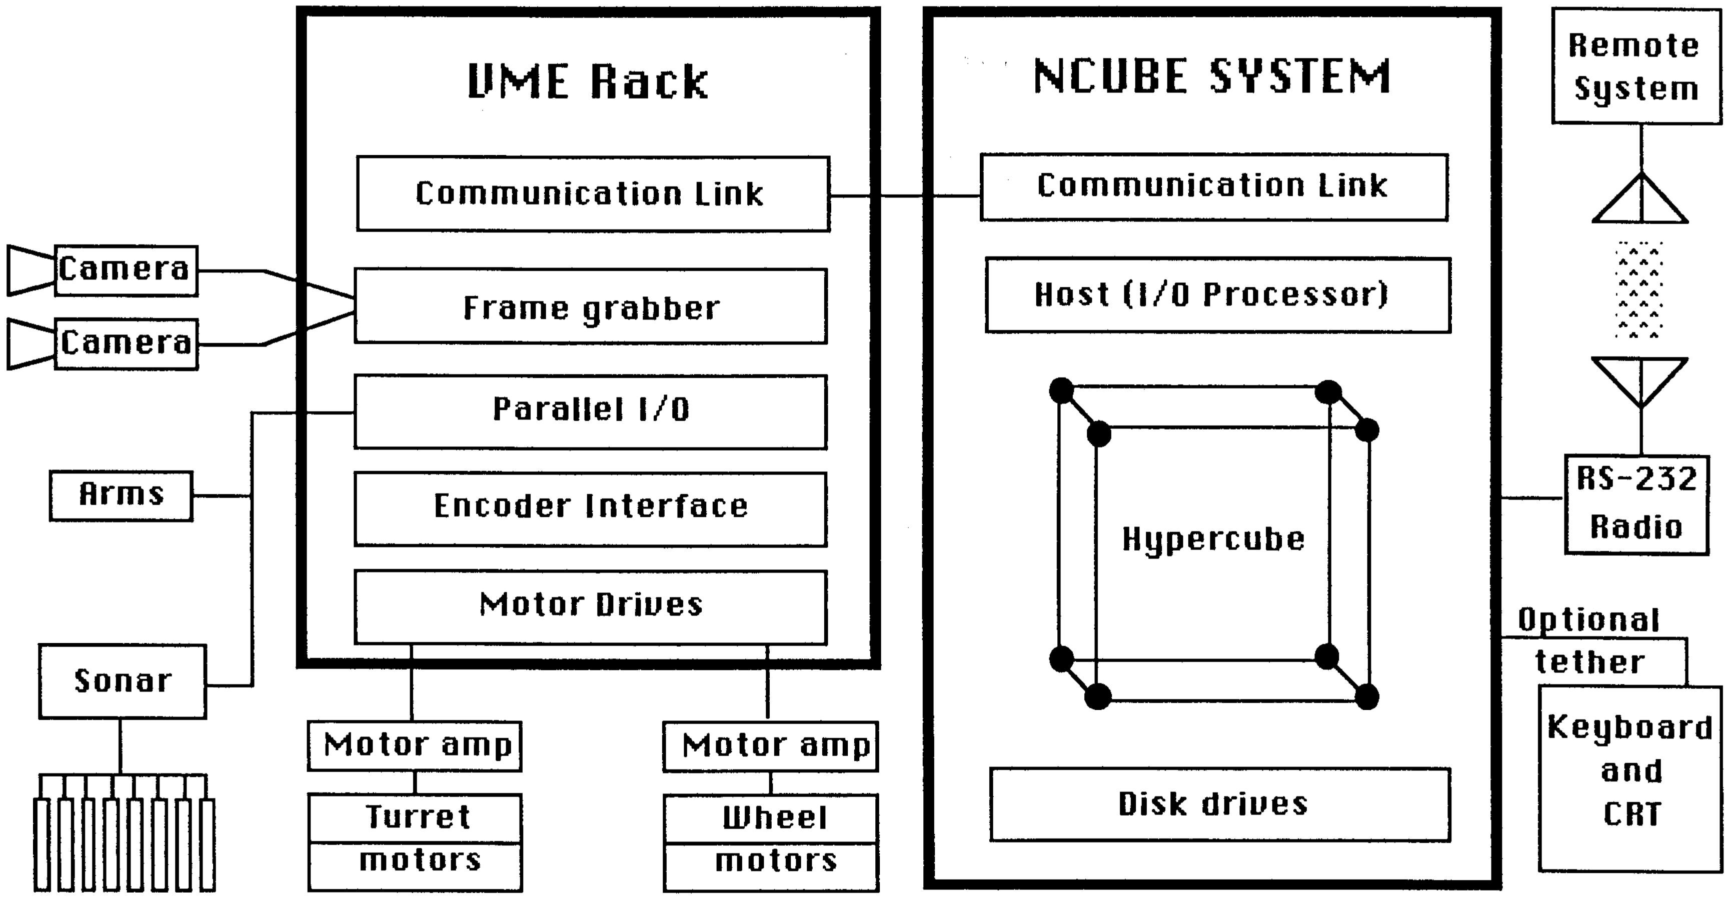

Figure 1: Robot system architecture

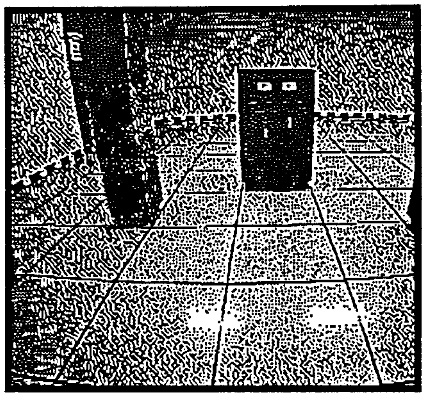

(a)

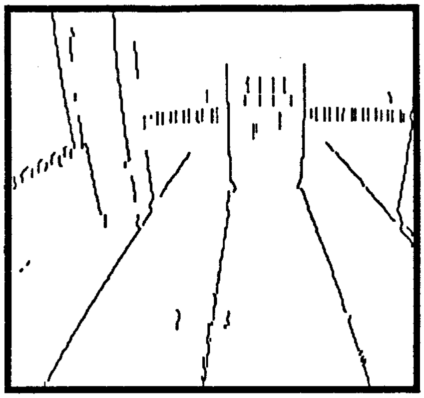

(b)

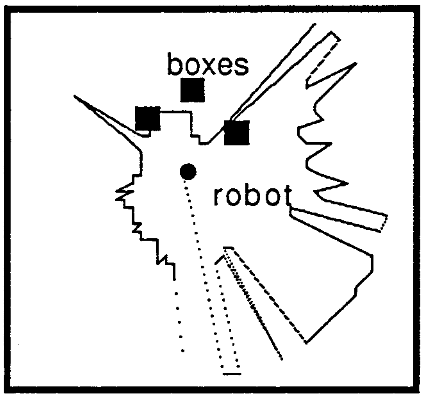

(c)

Figure 2: (a) Picture taken with camera on-board the robot; (b) vertical edges; (c) corresponding 360 deg. sonar scan, unprocessed data.

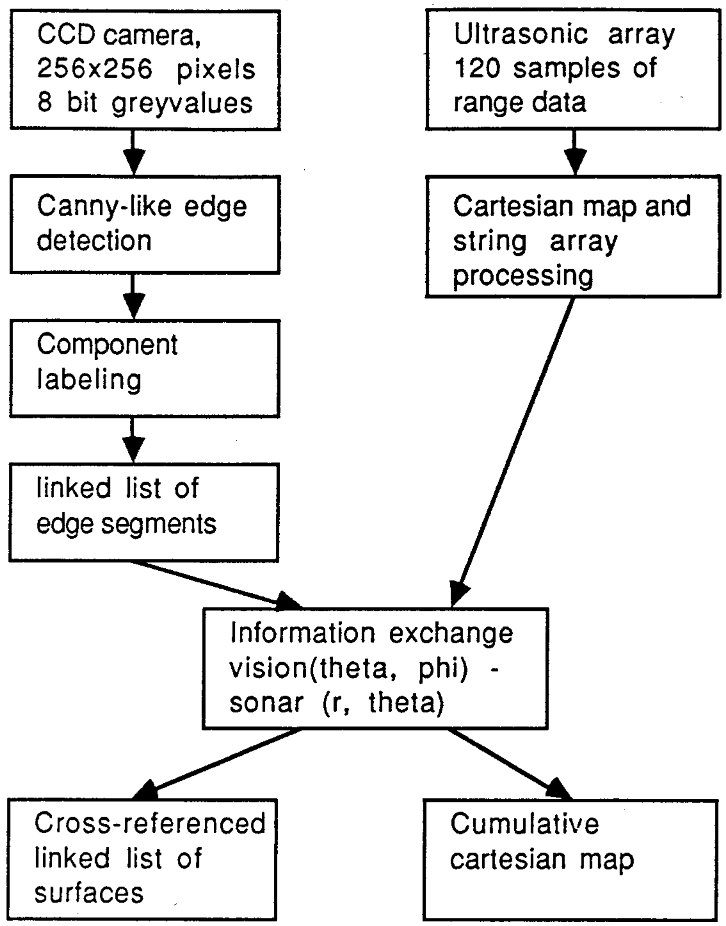

Figure 3: Block diagram of integration procedure.